

A mathematical assessment of the global state of vaccine hesitancy and coverage

Carlos Alexandre O'Doherty de Figueiredo

Department of Mathematics
Imperial College London

January 2018

A thesis submitted in part fulfilment of the requirements for the
degree of *Doctor of Philosophy*

**Imperial College
London**

Declaration of originality

I declare that this thesis is my own work and has not been submitted in for any other degree or diploma. Material which is not my own has been otherwise acknowledged and/or clearly referenced.

Copyright declaration

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

To my Mum and Grandmother.

Acknowledgements

There are many people that have contributed to this thesis either knowingly or unknowingly and I would like to take the time to thank them here.

I would like to extend my sincerest gratitude to my supervisors Nick Jones and Iain Johnston. Nick and Iain – thank you for the valuable scientific lessons you have taught me, for giving me freedom and encouragement to pursue my own research ideas, and for providing me with the materials and support to develop. Nick, I’d also like to thank you for all the interest you have shown in my academic progress and in my wellbeing: it is thoroughly appreciated.

I’d like to thank Heidi Larson for helping me to understand the current challenges in immunisations and for providing invaluable public health expertise: working on providing these public health messages would have been impossible without you. I’d like to thank Sumeet Agarwal and David Smith for taking the time to listen to my questions, providing valuable feedback, and for absorbing scientific discussions.

I’d like to thank my officemates, Alex Bolton, Alex Rush, Kuru Jain, and Marco Agnese (the ‘chuckle brothers’) for thought-provoking office discussions, Carcasonne marathons, and Brewdog-related escapades. Bolton, I’d like to thank you for all your patience with my endless questions and for editing this thesis. To James Kerr, David Blake, Milton, Matt Hitchings, Tess Hellgren, Dario Papavasiliou, Emma Harding, Alex Paddock, Till Hoffman, Kirsten Miller: thank you for your friendship and support. Mark Mullins, thank you for making physics fun: I would not have followed this path without your guidance.

Finally. To my family – thank you for everything.

Abstract

Hesitancy in vaccination programmes has increased over the past few decades and has led to vaccination delays and refusals resulting in epidemic outbreaks and curtailing progress towards disease elimination. Addressing vaccine hesitancy is a public health priority. In this thesis we assess the global state of vaccine coverage and vaccine hesitancy by performing a number of timely analyses, providing key public health messages. These analyses can be used to address vaccine coverage and inform effective immunisation programmes and policies.

In Chapter 2, we assess national-level childhood immunisation programmes by forecasting each nation’s vaccine coverage rates. By these means we are able to provide estimates of likely future coverage levels and derive a quantitative measure of the strength of a country’s immunisation programme – a measure we term the vaccine performance index. In countries where vaccines are freely available, our index captures losses to coverage based on hesitancy issues (non-access barriers that centre on personal beliefs and attitudes such as confidence in vaccines). We perform the first global-scale quantitative analysis of time-varying socio-economic determinants of vaccine coverage.

In Chapter 3, we focus on direct measures of attitudes to vaccines. Using the largest survey ever conducted on attitudes to vaccines (almost 70,000 individuals across 67 countries), we examine global variation in (and socio-economic determinants of) attitudes to vaccines. This is achieved through a four-question survey on individuals’ perceptions of vaccine safety, importance, effectiveness, and religious compatibility. In Chapter 4, we use a latent variable method to project these four attitudinal dimensions onto a single dimension which we claim to be the first quantitative measure of individual’s vaccine hesitancy. This measure is termed the vaccine hesitancy index and we infer hesitancy for each individual. Hesitancy is investigated globally and we perform a case study of hesitancy in Nigeria, where there have been notable recent concerns with vaccinations.

Identifying countries with sub-optimal immunisation rates, vaccine-hesitant groups, and understanding socio-economic correlates of coverage and hesitancy enables us to understand the complex coverage-hesitancy interplay and address gaps in coverage.

Contents

Acronyms and abbreviations	11
1 Introduction	13
1.1 A brief history of the development of childhood immunisations	13
1.2 Causes of incomplete vaccine coverage	15
1.2.1 Vaccine hesitancy	15
1.2.2 Access barriers	18
1.2.3 Quantitative analyses of determinants of incomplete coverage	19
1.3 Why should we care?	20
1.3.1 Disease resurgence	20
1.3.2 Herd immunity	23
1.4 Addressing coverage gaps	23
1.4.1 Mathematical challenges	24
1.5 Outline	26
2 Forecasted trends in vaccination coverage and correlations with socioeco- nomic factors: a global time-series analysis over 30 years	28
2.1 Introduction	28
2.2 Data and data processing	29
2.2.1 COVERAGE	29
2.2.2 GAPMINDER	30
2.2.3 Imputing missing data	31
2.2.4 Exploratory analyses	33
2.3 Statistical methodology	35
2.3.1 Forecasting vaccine coverage rates	35
2.3.2 A Vaccine Performance Index	38
2.3.3 Time-varying correlative study	39
2.4 Results	40

2.4.1	Country trends and forecasts with the VPI	40
2.4.2	Regional trends in determinants of DTP3 coverage	44
2.5	Discussion	46
3	The state of vaccine confidence 2016: global insights through a 67-country survey	48
3.1	Introduction	48
3.2	Data	49
3.2.1	GLOBAL-67	49
3.2.2	Country-level variables	51
3.3	Hierarchical models	52
3.3.1	Hierarchical modelling of vaccination attitude data	52
3.3.2	Modelling assumptions	55
3.4	Results	56
3.4.1	Global and regional trends in attitudes	56
3.4.2	Determinants of attitudes to vaccinations	61
3.5	Discussion	63
4	Deriving a quantitative index of vaccine hesitancy: mapping regional trends and establishing determinants both worldwide and in the Sahel	67
4.1	Introduction	67
4.2	Data	68
4.2.1	SAHEL-6	68
4.2.2	SAHEL-NIGERIA	69
4.3	Item response theory	70
4.3.1	One- and two-parameter IRT models	71
4.3.2	The graded response IRT model for polytomous responses	73
4.3.3	Model specification	74
4.3.4	Structural equation modelling	76
4.4	Bayesian hierarchical item response theory for vaccine hesitancy	77
4.4.1	Bayesian hierarchical IRT	77
4.4.2	Prior distributions	78
4.4.3	Structural models for vaccine hesitancy	79
4.4.4	Model estimation	81
4.5	Results	84
4.5.1	Model implementation and diagnostics	84
4.5.2	Inference of vaccine hesitancy	85

4.5.3	Socio-economic determinants of vaccine hesitancy	88
4.5.4	State-level vaccine hesitancy in Nigeria	91
4.6	Discussion	96
5	Thesis summary and specific future research avenues	99
5.1	Summary	99
5.2	Specific future research	101
5.2.1	Hesitancy inference using non-survey data sources	101
5.2.2	Identifying like-minded behavioural groups	105
5.3	Final Thoughts	106
	Bibliography	107
A	Appendix A	125
A.1	GAPMINDER and COVERAGE Data	125
A.1.1	Socio-economic factors	125
A.1.2	WHO regional classifications	125
A.1.3	Socio-economic factor imputation	129
A.1.4	Imputation results	131
A.1.5	Test correlation study	135
A.2	Coverage correlations with socio-economic factors supporting material	136
A.2.1	Multiple hypothesis correction using meta p -values	136
A.2.2	Selection of SEFs for Figure 2.5	137
A.2.3	Time-varying correlations of BCG, DTP1, MCV, POL3	138
A.3	Vaccine performance index supporting material	141
A.3.1	Gaussian process forecast performance	141
A.3.2	Alternative forms of the Vaccine Performance Index	141
A.3.3	Vaccine performance index values	143
B	Appendix B	146
B.1	Hierarchical linear model assumption checks	146
B.2	Vaccine effectiveness versus vaccine safety	147
C	Appendix C	148
C.1	Sensitivity analysis and simulation study	148
C.1.1	Data generation	148
C.1.2	Choosing individual- and country-level variance hyper-priors	149
C.1.3	Testing model robustness to extreme item response parameters	150

C.1.4	Full parameter inference and MCMC convergence	150
C.2	Verifying model assumptions for Bayesian hierarchical IRT	154
C.2.1	Testing and verifying the conditional independence assumption	154
C.2.2	Testing successful burn-in and convergence	155
C.2.3	Verifying linear model assumptions	155
C.3	JAGS code for Bayesian hierarchical IRT models	158

Acronyms and abbreviations

AFR Africa region

AIC Akaike information criterion

AMR Americas region

BCG Bacillus Calmette-Guérin vaccine

BIC Bayesian information criterion

DIC Deviance information criterion

DTP Diphtheria-tetanus-pertussis

EMR Eastern Mediterranean region

EUR European region

GAVI Global Alliance for Vaccines and Immunization

GP Gaussian process

GRM Graded response model

GVAP Global Vaccine Action Plan

HPD Highest posterior density

IRT Item response theory

LMIC Low and middle income country

LSHTM London School of Hygiene and Tropical Medicine

MAP Maximum a posteriori

MAR Missing at random

MCAR Missing completely at random

MCV Measles-containing vaccine

MMR Measles mumps rubella

MNAR Missing not at random

POL Polio

SAGE Strategic Advisory Group of Experts on Immunization

SEAR South East Asia region

UNICEF United Nations Childrens' Fund

VCP Vaccine Confidence Project

VHI Vaccine Hesitancy Index

VPI Vaccine Performance Index

WHO World Health Organization

WINGIA WIN/Gallup International Association

WPR Western Pacific region

1 | Introduction

PhD questions (3,5,6)

— Philistine, *Guardian*
crossword No. 26,310

1.1 A brief history of the development of childhood immunisations

Childhood immunisation programmes are one of the great public health success stories. From the development of the smallpox vaccine in the late 18th century, vaccinations have prevented death and serious illness for countless people. The World Health Organization (WHO) estimates that 17.1 million lives have been saved by the measles vaccination alone since 2000 [1]. Although smallpox was declared eradicated in 1980 after an intense global effort from WHO and national governments [2], it remains the only vaccine-preventable disease with eliminated status [3]. To understand why other vaccine-preventable diseases have not been eliminated, despite concerted public health efforts, it is instructive to consider the somewhat chequered history of vaccinations.

The development of childhood vaccinations proceeded in two distinct stages owing to the development of bacterial cell cultures followed by viral cultures [4]. The former of these stages (1890-1950) began with the development of the bacillus Calmette-Guérin (BCG) vaccine against tuberculosis (TB) [5]. Early clinical trials of the vaccination were successful [5]. However, in 1930, 72 children died and a further 173 contracted TB in Lübeck, Germany, after the vaccination they received was contaminated with a more virulent strain of the disease [4, 6]. Following this ‘Lübeck disaster’, public acceptance of the vaccination declined until after the Second World War when campaigns by WHO and United Nations (UN) led to a renewed acceptance [4, 7].

During the early 20th century, vaccinations against the bacterial diseases diphtheria, tetanus, and pertussis followed [8]. The introduction of these vaccines greatly reduced the number of cases and resultant deaths. In the USA in the 1920s, for example, there were

approximately 140-150 cases per 100,000 [9], but only six reported cases since 2000 [10]. The introduction of the pertussis vaccine (introduced in the 1940s) performed similarly well, with a reduction in clinical cases of 99% in the USA [11], although many side effects were reported [12]. The substantial gains made by the introduction of the pertussis vaccine (which, since 1948, has been administered as a trivalent vaccine with diphtheria and tetanus) have been eroded in the last few decades through a combination of factors. Partly motivated by fears over side effects [11], many countries changed their pertussis vaccination in the 1990s to an acellular pertussis (aP) vaccine [13, 14]. Initially perceived to be as protective as the previous pertussis vaccine, it was later found that the old vaccine was more efficacious [11]. The aP vaccine is considered to have played a major role in the resurgence of pertussis around the world today [11, 13, 15] (although evolutionary shifts, detection bias, and clusters of non-vaccinated children have also been highlighted as contributing to this increase [11, 15]).

The onset of viral tissue cultures in the late 1940s led to rapid progress in the development of vaccinations against polio, measles, mumps, and rubella over the next few decades [4]. In 1955, after the success of the inactivated polio vaccine (IPV) during a mass clinical trial, the IPV was licensed for use [16]. Despite the success of the IPV, many virologists questioned its benefits relative to the orally-administered polio vaccine (OPV), arguing that only a live virus could confer long-term protection and that passage of the attenuated virus through faecal matter could offer protection to those unvaccinated [16, 17]. Thus the oral polio vaccine (which contains an attenuated live agent) was recommended in the USA and elsewhere [16–18]. Unlike IPV, however, OPV can cause vaccine-associated paralytic poliomyelitis (VAPP) in a small number (1 in 2.7 million) of those who are vaccinated [18]. Although the associated risk of VAPP in the OPV is low, there were an estimated eight cases per year in the USA [19] and an order of magnitude more in India [20]. As of today, the risk of contracting polio across the majority of the world is so low that the continued use of OPV has been hard to justify. Many countries – particularly those where indigenous polio has been eliminated – have included IPV into their immunisation programmes to reduce vaccine refusals on the basis of contracting VAPP [17, 21].

The vaccination against the measles virus was first licensed eight years after the polio vaccine and led to an annual reduction in measles cases in the USA from about half a million to a few thousand [22]. A similar debate about the efficacy of inactive versus attenuated measles ensued, but was settled more convincingly. This was in part due to the lower mortality associated with measles than polio, and so the attenuated vaccine, though associated with a greater risk of minor side effects, required fewer doses and offered greater protection over a sustained period [23]. Introduced in 1971, a combined vaccine for measles, mumps, and rubella (MMR) was introduced that contained the attenuated measles virus [4]. Despite the

safety and efficacy of the MMR vaccine [24], it has arguably been the most controversial of modern vaccinations. In 1998, a link between the MMR vaccine and child autism disorders and colitis was claimed by Andrew Wakefield [25]. A media frenzy followed, which was often misleading [26] and resulted in the UK's MMR coverage rates declining from 91% in 1997 to 80% in 2003 [27]. Subsequent measles outbreaks followed and reached their highest levels since 1997 in 2012 [28]. Although Wakefield's claim proved to be fraudulent (and Wakefield's article and medical registration have since been retracted [25, 29]) it precipitated large-scale mistrust in immunisations which were often worsened by media reporting [30]. In 2008, the UK was declared measles endemic [30]. Vaccination rates eventually recovered to pre-1998 levels in 2012 [27] although hesitancy towards MMR continues, despite numerous studies finding no link between MMR and autism since Wakefield [24, 31–33].

1.2 Causes of incomplete vaccine coverage

1.2.1 Vaccine hesitancy

Vaccine refusal is not a modern phenomenon. As far back as the 19th century, concerns have been expressed over vaccinations. In Victorian England, for example, there was disquiet amongst the working classes over the power of the state to impose compulsory healthcare, viewing the resulting fines for disobedience as a tax on the lower classes [34]. Objections on religious grounds (as the vaccine contained animal virus) and safety fears were also expressed [34, 35]; fears which are still prevalent today [36]. Anti-vaccination movements formed shortly after the introduction of vaccines. The Anti-Vaccination League (later, the London Society for the Abolition of Compulsory Vaccination), for example, opposed the Compulsory Vaccination Act of 1853 and disseminated anti-vaccination books and journals [37, 38]. Similar anti-vaccination groups and opposition towards vaccination grew across Europe [38, 39] and the USA [38] and – combined with government policy – held back progress against smallpox. Germany found more success than the UK by introducing vaccines into law more slowly (initially only requiring compulsory vaccinations during epidemics) and was quicker than England to recognise the importance of re-vaccination (subsequent vaccinations to extend immunity) [40, 41]. Incidence and mortality rates declined in both countries towards the end of the 19th century; however, in part due to mass anti-vaccination demonstrations, England allowed conscientious objectors to be excused from vaccinations, whereas Germany did not [40]. These policy differences led to a divergence in the number of successful vaccinations between the two countries [38, 39]. Although the methods for spreading misinformation have changed over the past century, the underlying objections have largely stayed largely the same. Fears over safety, philosophical or personal objections, and religious compatibility issues have



Figure 1.1: **Anti-vaccination illustration published by the London Society for the Abolition of Compulsory Vaccination** The illustration depicts the vaccinator as death (supported by the police), and questions the safety of vaccines, “Every mother knows of infants, who were well and hearty up to the time of their vaccination and were never the same afterwards ... [vaccination] does not keep off smallpox.” Image from the Historical Medical Library of the College of Physicians of Philadelphia [42].

all featured prominently in vaccine refusal over the last two centuries [38]. As is the case with VAPP, there are a number of real safety concerns with childhood vaccinations (see Table 1 in [43] for a more extensive overview), but vaccine objection citing safety concerns has often been based on unsubstantiated claims.

In 1996, a loss of trust in the polio vaccine in Nigeria can be attributed to a combination of pre-existing safety concerns and the spread of false information. In 1996, the pharmaceutical company Pfizer carried out drug trials during a meningitis epidemic which resulted in the deaths of 11 children in the state of Kano in northern Nigeria [44]. Seven years later, in 2003, religious leaders in three northern Nigerian states built on these pre-existing fears

by claiming the polio vaccine was a Western ploy to sterilise Muslim women [45]. Although progress has been made towards polio eradication in Nigeria, high levels of vaccine hesitancy in Northern Nigeria threatens eradication. Although Nigeria was removed from the list of polio-endemic countries in 2015 [46] (and after a sustained period without a recorded polio case) polio has since returned to Northern Nigeria as a result of sub-standard vaccination uptake [47]. Misinformation from religious leaders is not limited to Nigeria; false information was disseminated by the *Taliban* in Pakistan, who issued *fatwas* discrediting the polio vaccination [48].

Other vaccine scares have occurred across the world which have linked the pertussis vaccine [49] to encephalopathy, the DTP vaccine to sudden infant death [50], and which relate the number of immunisations a child has to diseases including autism [51, 52]. In the past few decades, health professionals – often an individual’s most trusted source for health information [53] – have lent credence to pre-existing fears. Approximately 30% of GPs in the UK advised against pertussis vaccination in the 1970s [4] and, in 2004, only 57% of Scottish GPs describing themselves as ‘very confident’ in discussing the MMR vaccine with patients [54].

Although coverage rates remain high across the world [27], vaccine scares have caused lasting damage. This is made apparent through static or declining immunisation rates in many countries [27], the persistence of pockets of non-vaccinators [55–57], and disease re-emergence.

It is instructive to have a universal definition of the factors associated with vaccine delay and rejection that are not associated with access barriers. These delays and refusals have been termed *vaccine hesitancy* and, in 2014, WHO’s Strategic Advisory Group of Experts on Immunization (SAGE) met to formally define vaccine hesitancy:

“[Vaccine hesitancy is a] delay in acceptance or refusal of vaccines despite availability of vaccine services. Vaccine hesitancy is complex and context specific, varying across time, place and vaccines. It is influenced by factors such as complacency, convenience and confidence.” [58]

This definition is sometimes referred to as the ‘3Cs’ model of vaccine hesitancy: *confidence* is defined as trust in the effectiveness and safety of vaccines and trust in the healthcare system that delivers them; individual *complacency* refers to vaccine rejection due to a low perceived risk of disease; and *convenience* is the extent to which vaccine uptake is hindered by other barriers such as the time and location of vaccinations, perceived quality of service, and language barriers (for the full definitions, see [58]). (Throughout this thesis we will refer

back to this definition.)

As discussed above, vaccine hesitancy in many parts of the world provides a material threat to the success of immunisation programmes. Non-hesitancy barriers – that is, access to vaccinations – also play a key role in driving sub-optimal vaccination rates.

1.2.2 Access barriers

There are stark differences in access to basic healthcare services around the world. Much of the Western world has universal access to healthcare facilities (and corresponding healthcare infrastructure and delivery systems), yet in many low and middle-income countries (LMICs), basic facilities are lacking or are only available to the wealthiest in society [59–61]. In India, for example, out-of-pocket health expenditures have been suggested for moving millions of people into poverty [62].

Distance to healthcare facilities plays a major role in whether healthcare (including vaccinations) is sought [36, 61, 62]. When distance and poverty combine, education and access to relevant healthcare information can become further barriers [63–65]. Access to healthcare information also plays a vital role for immigrant communities in areas where vaccinations are free and easy to access [66]. These access barriers can often place further strain on the poorest in society, making them susceptible to disease and can contribute towards a ‘poverty-disease’ cycle. This poverty-disease cycle has been cited multiple times for trapping people in poverty and reducing life chances [67–69].

Ethnicity, wealth, and economic opportunity are often closely related. For many years access to healthcare (and vaccinations) has been lower for minority ethnicities [70, 71] and this is not limited to LMICs. Until 2005, there was disparity in access to routine MMR and polio vaccines in the USA between white and Hispanic ethnicities [72]. Wealth, ethnicity, and other socio-economic statuses are so closely intertwined that it can obscure a full understanding of the factors at play in access to vaccinations [73].

A coordinated global effort between multilateral organisations, governments, immunisation experts, and stakeholders is required to make vaccines available for every child. In 2012, a framework for improving worldwide vaccination access and uptake received endorsement from 194 member states of the World Health Assembly: this framework is called the Global Vaccine Action Plan (GVAP) [74]. Achieving the goals laid out by GVAP (and the United Nation’s sustainable development goals [75]) has proved challenging for the reasons outlined, but there have also been sizeable gains made to global vaccination rates over the past few decades (and decreases in the recorded cases of associated diseases [27]). This has been due to successful coordination between multilateral organisations such as the Global Alliance for Vaccines and Immunization (GAVI), United Nations Children’s Fund (UNICEF), the Pan

American Health Organization, national governments, vaccine producers, and – for LMICs – procurement teams who aim to lower the cost of vaccinations [76, 77].

1.2.3 Quantitative analyses of determinants of incomplete coverage

Quantitative connections between socio-economic factors, attitudes, and vaccine uptake have been studied in a range of settings. In studies that use only socio-economic data to assess vaccine uptake, the relative contributions of access to vaccinations versus attitudinal barriers can be hard to disambiguate. Research has therefore tended to follow two reasonably distinct analysis strands: quantifying the strength of socio-economic correlates of vaccine uptake and assessing vaccine perceptions through survey or interview data.

Links between socio-economic data and vaccine uptake have been extensively studied and repeated trends have been found. Maternal education level [78–87], household wealth [85, 86, 88–91], birth place of child and distance to healthcare facilities [78, 80, 89, 90, 92, 93], and the administration of pre- or post-natal care [83, 87, 92] appear to be the most consistently (and strongly) informative correlates of complete immunisation status¹. Paternal education has been shown to be informative of complete vaccination status independently of maternal education in a number of countries [94, 95]. Religion, ethnicity, and nationality have also been found to play a strong role in uptake [92, 96, 97].

Strong country-specific trends in determinants have also been established. Gender plays a strong role in India and Pakistan, with males more likely to be immunised reflecting societal and cultural norms [98–100]. Ethnicity and religion play major roles in East African countries [101] and in the UK [97], where MMR uptake during the Wakefield-era scare disproportionately impacted certain ethnicities [102, 103]. Distance to healthcare facilities is a strong correlate in LMIC countries, in particular those with a large relative population living rurally (such as India and many African countries) [101, 104, 105].

When personal beliefs are explicitly studied through survey data, it is repeatedly found that individuals place a high degree of trust in healthcare professionals [106–109]. Trust in the government is mixed, however, as many individuals are sceptical of government intentions [110, 111]. Parental hesitancy towards childhood vaccinations is also strongly associated with the perceived risks and benefits [106, 112–115], which often results in parents rejecting a single vaccination, often on the advice of doctors [4, 106, 116].

Parents’ hesitancy towards vaccines is therefore a complex mix of socio-economics, sources of trust, perceived benefits and side-effects, and weighs decisions from a number of sources including health professionals and the government.

¹Complete immunisation status is defined as whether a child has received all recommended immunisations.

1.3 Why should we care?

It is estimated that routine childhood immunisation prevents between two and three million deaths every year from diphtheria, tetanus, pertussis, and measles alone [117]. Global vaccination rates have made substantial gains over the past few decades, yet, since 2012, global coverage² has decreased for many vaccination programmes (Figure 1.2). Although global vaccine coverage is high, WHO estimates that further gains to vaccination rates could save 1.5 million deaths annually as 19.5 million children remain unvaccinated [117].




Figure 1.2: **Global gains in vaccination coverage levels are threatened by increasing vaccine hesitancy** Global vaccination rates have increased markedly since 1980. However, since 2012, global vaccination rates are declining for vaccinations against tuberculosis (BCG), diphtheria, tetanus and pertussis (DTP) [coverage of three completed doses is shown], measles (MCV), and polio (POL) [coverage for three completed doses is again shown]. Data is from the WHO-UNICEF estimates of vaccination coverage [27].

1.3.1 Disease resurgence

Resurgence of disease consistently follows periods of vaccine refusal, typically occurring in regions with low local coverage rates. Until 2005, for example, measles remained eradicated in the UK, but – following the loss of confidence in vaccines resulting from the Wakefield scandal – large outbreaks followed. The first, between 2006 and 2007, caused large outbreaks amongst the general population but also within communities with low vaccination coverage (notably the Jewish Orthodox and Irish traveller communities [118]). Although reported

²A crude estimate which simply averages over each country’s national immunisation rate.

confidence in the MMR vaccine showed signs of increasing [119], in 2013 another large measles outbreak hit Swansea, Wales [55], which mostly affected 10-14 year olds [118] – adolescents who were born between 1998 and 2003 when parental hesitancy towards the MMR vaccine was at its peak. In addition to the national newspaper coverage of the MMR-autism link [26], the *South Wales Evening Post* ran anti-MMR vaccination campaigns, which likely worsened parents’ fears and resulted in relatively low MMR coverage rates in Swansea [120]. Measles cases have followed vaccine refusal or hesitancy in a number of other settings, most notably the USA [72, 121], France [122], the Netherlands [123], and elsewhere in Europe [124]: these outbreaks resulted in thousands of reported cases and dozens of deaths [72, 124–126].

Increases in reported cases and deaths have also followed hesitancy in other vaccines in a number of high-profile cases: for example, rejection of the polio vaccine in Pakistan, Nigeria, and Afghanistan [57, 127, 128] and refusal of the diphtheria-tetanus-pertussis vaccine in the UK (pertussis controversy [129]) and in Russia (diphtheria rejection [130]).

The link between high coverage rates and protection from disease outbreak is highlighted in Figure 1.3. Annual vaccine coverage rates and reported disease cases are compared between countries with high and low coverage rates for the MMR and polio vaccines. In Figure 1.3 A, MMR vaccination coverage is shown for Sweden (red), the UK (green), and the USA (blue). Reported measles cases for these countries are shown in Figure 1.3 B. In Figure 1.3 C, polio coverage for Botswana (red), Nigeria (green), and Pakistan (blue) is displayed, with respective cases shown in Figure 1.3 D. The importance of maintaining (and sustaining) coverage rates in excess of 95%³ is clear: reported cases of measles in Sweden and polio in Botswana – countries with high coverage for MMR and polio (respectively) – are consistently an order of magnitude less than countries with lower vaccination rates.

When the Wakefield scare originated in the UK in 1998, Sweden was not resistant to the news. The Wakefield article has been suggested to have contributed to Sweden’s decline in MMR rates in the late 1990s [131]. Despite this reduction, Sweden’s high MMR coverage rate likely greatly contributed to limited outbreaks. The UK’s dramatic fall in MMR coverage, however, resulted in thousands of reported cases over the following years (Figure 1.3).

In Africa, Botswana’s polio strategy contrasts markedly with that of Nigeria. As outlined previously, a particularly complex history of vaccine refusal exists in Nigeria (which included a ban on the polio campaign in some northern states in 2003 [57]). In 2004, a polio case was reported in Botswana which was found to have been imported from Nigeria [132]. The response from the government of Botswana was immediate. An emergency vaccination campaign – supported by WHO and UNICEF – was launched to immunise almost 200,000 children against polio [132]. However, despite the resumption of the polio campaign across

³95% is often used by policy-makers as the threshold for herd immunity – see Section 1.3.2.




Figure 1.3: **Decreases in coverage rates are associated with increases in disease incidence** The link between low (or declining) coverage rates is made clear with examples from the MMR and polio vaccinations. MMR coverage rates for Sweden, UK, and USA are shown (A) with associated reported cases (B). Measles cases are highest in the UK which has had consistently lower coverage rates than the USA or Sweden. Coverage rates for the completed polio series are shown (C) with corresponding outbreaks (D) for Botswana, Nigeria, and Pakistan. Botswana. Botswana has maintained coverage in excess of 95% for the past decade and has had only had one reported case since 2000. Polio issues in Nigeria and Pakistan are discussed in Section 1.2.1. The black horizontal line represents 95% coverage. (Data from [27].)

Nigeria, vaccination levels remain low today (Figure 1.3).

Pakistan – which remains polio endemic – has similar polio concerns to Nigeria. It has been estimated that whilst only 1% of children were missed through parental refusal over vaccine safety in 2015, a further 20% of children were missed due to parental concerns over security [127]. These security concerns include death threats from militant groups (the *Taliban* in Pakistan and *Boko Haram* in Nigeria [128]) who have killed a number of vaccine campaign workers [128, 133]. *Taliban* opposition to the polio vaccination in Pakistan stems from the USA’s Central Intelligence Agency’s plot to gather information on the whereabouts of Osama bin Laden under the guise of a hepatitis vaccination drive [134].

1.3.2 Herd immunity

High vaccination rates stifle the propagation of a disease. The vaccination coverage level at which a disease outbreak can no longer cause an epidemic is called the epidemic threshold [135]. Once vaccination levels are above this threshold, protection against an epidemic is conferred by vaccinated to unvaccinated individuals: this is known as *herd immunity* [136]. The amount required for herd immunity varies by disease, though 95% is often taken as the value of coverage required to confer protection via herd immunity [136].

Despite many countries around the world holding non-endemic status for diseases [28], high immunisation rates are still required to prevent outbreaks from imported cases. Recent measles outbreaks from imported diseases in Swansea and California (and the lack of outbreak of imported polio in Botswana) highlight emphasise this fact [119, 121]. It is therefore crucial, while diseases are not globally eliminated, that vaccine coverage rates remain above the levels required for herd immunity. This will require systems to identify early signals of increasing hesitancy and mechanisms to address it. As of late, many countries have addressed this issue by demanding that parents' refusal of vaccines is noted by the school (so that public health agencies are aware of non-vaccinated children) or by withholding public resources from parents [137, 138].

1.4 Addressing coverage gaps

Transnational public health bodies such as WHO and UNICEF assist governments, stakeholders, and public health teams to increase access to vaccinations and to coordinate responses to vaccine-related issues [76]. In December 2017 (at the time of writing), WHO is coordinating an emergency response to an outbreak of diphtheria at a Rohingya refugee camp in Bangladesh [139, 140]. The Myanmar state of Rakhine – in which the majority of Myanmar's Rohingya reside – has amongst the lowest DTP coverage in the country, and rates have been declining over the past decade [141].

A number of guidelines, targets, and recommendations to improve national resilience to diseases have been laid out and include specific immunisation rate or disease reduction targets [142, 143] and other frameworks such as the Tailored Immunisation Programme (TIP) which aims to reach under-vaccinated groups in Europe [144]. TIP – a series of guidelines to bolster immunisation rates – has been used by numerous countries, including the UK [145] and Sweden [131] to establish and address sub-optimal coverage within specific communities. These public health interventions have evolved in light of scientific research into vaccine acceptance and intervention strategies now include specific targetting of socio-economic groups [144], improving education levels amongst mothers with low health literacy [146, 147],

and disseminating more information to communicate vaccine benefits and risks more transparently [148].

Complicating these intervention strategies is, however, individual reluctance to change opinion. In many cases, individuals outright refuse to engage with immunisation programmes [149, 150] and intervention strategies that provide individuals with pro-vaccine sentiment has even been found to further entrench anti-vaccination sentiment [151]. These results suggest that some individuals may not be shifted into pro-vaccination behaviour, emphasising the need to understand the location and dynamics of hesitancy, so that appropriate intervention strategies can be designed.

Understanding the relationship between hesitancy and coverage uptake is key to developing novel methods to address the issues caused by pockets of non-vaccinators. Efficient identification of these clusters through monitoring of local immunisation rates has been identified as a key public health challenge [152]. Monitoring of trust in immunisation programmes at finer spatial scales has been proposed to better understand how vaccine hesitancy arises [153, 154]. Coupling this information with coverage rates, socio-economic data, and network structure can – through applied epidemiological models in real-world settings (for example) – enable us to develop a more proactive approach to tackling diseases spread through confidence and coverage forecasting.

Despite concentrated efforts to address coverage gaps, global targets for vaccination rates and eradication of diseases have consistently been missed. Examples include the failure to meet WHO’s regional target of European measles elimination by 2015 [155], the failure to end neonatal tetanus and wiping out measles from three world regions by 2015 (set out by GVAP [142]), and the failure to meet specific immunisation rate goals [142]. In this thesis we identify a number of mathematical challenges in the topic area of vaccine coverage and vaccine hesitancy: answering these challenges can provide timely public health messages to inform interventions to address gaps in vaccine coverage.

1.4.1 Mathematical challenges

The following key challenges in assessing gaps in vaccine coverage have been identified:

Challenge 1: Forecasting national vaccine coverage rates and developing a quantitative marker of the performance of immunisation programmes;

Challenge 2: Providing a quantitative assessment of historic and current trends in factors modulating coverage;

Challenge 3: Assessing the state of worldwide attitudes to vaccines;

Challenge 4: Developing an individual index of vaccine hesitancy to identify hesitant clus-

ters and to re-assess herd immunity levels;

Challenge 5: Predicting future levels of vaccine hesitancy and modelling the interaction between hesitancy and coverage.

Despite the GVAP’s immunisation goal of 90% national vaccination coverage (and at least 80% in every sub-national district) [142], there do not exist national- (or sub-national) level forecasts of vaccine coverage data: **Challenge 1** address this shortfall. Coverage forecasts can enable an early assessment of immunisation goals (such as those laid out by GVAP [142]) and the reallocation of resources to countries falling behind in these goals. Forecasting coverage rates in countries with universal access to immunisations can also establish future levels based on the level of vaccine hesitancy as changes in coverage will not be related to access barriers. To meet this challenge, a flexible time-series forecasting tool – capable of modelling the wide range of dynamics in national level immunisation coverage (see Figure 2.1 C) – is required.

Challenge 2 concerns establishing global historical patterns in the socio-economic correlates of immunisation. Challenge 2 has been identified as no such global-scale, time-varying analysis of national-level determinants of coverage has been performed (although we note that numerous studies probe an individual’s determinants at isolated time points – see Section 1.2.3). In performing such an analysis, there are notable statistical hurdles, such as data gathering (to obtain a large set of variables putatively linked to coverage) and data cleaning (involving the careful handling of any missing data). In seeking time-varying correlates of vaccine coverage, an appropriate correlation metric will be required. *Challenges 1 and 2 are met in Chapter 2.*

As outlined in Section 1.2.3, many studies also probe attitudes to vaccinations in specific contexts. However (as with challenge 2), a global-scale study on attitudes to vaccines has not been undertaken. This is an ambitious goal and constitutes **challenge 3**. Probing attitudes on this grand scale will naturally require a large dataset and appropriate mathematical modelling tools. *Challenge 3 is the focus of Chapter 3.*

Challenge 4 is based on calls by SAGE [58] and by national policy makers [156] to understand vaccine hesitancy. Vaccine hesitancy is defined through a complex mix of factors ([58] and 1.2.1) and so quantifying it in this way requires a mathematical approach that can appropriately summarise individuals’ vaccine attitudes. Obtaining a quantitative measure of hesitancy unlocks many other analyses: we can assess spatial trends in hesitancy; we can link hesitancy to socio-economic factors; and – if attitude data becomes routinely-collected – then hesitancy dynamics can be modelled (and incorporated into epidemiological models). These analyses have obvious advantages for identifying and predicting hesitant groups (these

analyses constitute **Challenge 5**). *The goal of deriving a quantitative measure of hesitancy and understanding its determinants is the goal of Chapter 4.* Although not analysed further in this thesis, the re-assessment of herd-immunity levels and the analyses methods outlined proposed above for Challenge 5 are discussed throughout this thesis (in particular, in Chapter 5).

An outline of this thesis is now provided.

1.5 Outline

This thesis is arranged into four further chapters:

- In Chapter 2, **Challenge 1 and 2** above are considered. Forecasts of national vaccine coverage rates are obtained via Gaussian process regression. These forecasts are used to derive a summary measure of the strength of a country’s immunisation programme. This summary measure is termed the **vaccine performance index**. This index quantifies variability in national coverage rates induced by hesitancy and access issues. Index values are compared between countries and interpreted in view of GVAP’s target of attaining 90% vaccination coverage in all countries by 2015. To find temporal associations between socio-economic factors and coverage, we perform a large-scale correlative study with publicly-available socio-economic data. Comparisons are drawn between world regions to reveal spatio-temporal trends.

The results presented in this chapter were published in August 2016 as **Forecasted trends in vaccination coverage and correlations with socioeconomic factors: a global time-series analysis over 30 years** in *The Lancet Global Health* [157] with commentary by Dr. Michiel van Boven and Dr. Alies van Lier [158].

- In Chapter 3 worldwide variation in attitudes to vaccinations are explored using the largest dataset ever collected on attitudes to vaccines (**Challenge 3**). This survey was conducted in 2015 by WIN/Gallup International Association (WINGIA) in collaboration with the London School of Hygiene and Tropical Medicine (LSHTM) with whom we collaborate. This survey examines individuals’ views across 67 countries on four dimensions of vaccine attitude: the importance, safety, effectiveness, and religious compatibility of vaccines. Hierarchical linear models are employed to investigate relationships between socio-economic factors and these four dimensions of vaccine attitudes.

The results presented in this chapter were published in August 2016 as **The state of vaccine confidence 2016: global insights through a 67-country survey** in *EBioMedicine* [159] with commentary by Dr. Steven Black [160].

- In Chapter 4, we attempt to quantify SAGE’s definition of vaccine hesitancy by defining a **vaccine hesitancy index** (**Challenge 4**). Survey data from the previous challenge is used and augmented with an additional survey conducted in the Sahel region by WINGIA a year later in 2016. Individual-level hesitancy is obtained via a latent variable model which projects the four dimensions of attitude obtained in the two surveys onto a latent dimension which we propose to be a measure of vaccine hesitancy. This inference analysis – still to be placed in final form for journal submission – therefore explores the treatment of vaccine hesitancy as a single quantitative latent summary and establishes its determinants. Country-level trends – in particular, those countries in the Sahel region – are explored in depth. Sub-national hesitancy and its drivers are assessed in Nigeria – a country with standout vaccine hesitancy concerns. The results presented in this chapter are currently in preparation for publication.
- In Chapter 5, we discuss the relevance of our findings for addressing coverage gaps and outline future mathematical and statistical methods which could be deployed to further progress in closing coverage gaps. We discuss how the latent variable models used in Chapter 4 may be used to infer vaccine hesitancy *in the absence of vaccine hesitancy surveys*. This could have profound significance for using publicly-available public health datasets to infer and forecast vaccine hesitancy in a number of LMICs. Moreover, such historic trends in hesitancy can be linked to vaccine coverage to understand the interplay between these two important factors.

2 | Forecasted trends in vaccination coverage and correlations with socioeconomic factors: a global time-series analysis over 30 years

Novel piece of theorem
mathematician produced (4)

— Brendan, *Guardian*
crossword No. 27,003

2.1 Introduction

Monitoring attitudes towards immunisation programmes and progress towards immunisation goals has been cited as a public health priority [58, 142, 153, 154]. Indices that record trust in immunisation programmes and assess progress towards immunisation goals have been urged to aid policy makers evaluate immunisation goals [153, 154].

In this chapter we determine the strength of national childhood immunisation programmes by developing a measure of the relative ability of national immunisation programmes to successfully vaccinate their populations. This index is derived from forecasts to national-level vaccine coverage data and rewards countries that are able to obtain and maintain high vaccination coverage. Conversely, our index penalises countries for either not being able to obtain high coverage, or not being able to sustain it. We term this index the vaccine performance index. Although forecasts of national-level coverage are useful in their own right, our vaccine performance index summarises recent trends and ranks countries by the strength of their immunisation programmes.

We investigate our index in light of the Global Vaccine Action Plan’s target of 90% coverage of the third dose in the diphtheria-tetanus-pertussis immunisation programme [142]. This indirect approach of assessing a country’s immunisation strength also allows vaccine hesitancy to be captured: when a country’s health system provides universal access to vaccines, then the primary cause of depleted coverage rates is *vaccine hesitancy*. However, in countries with non-universal access to vaccinations, our index measures vaccine losses due to both access and hesitancy barriers.¹ In both settings, forecasts of coverage can enable policy makers to plan for future coverage possibilities.

In addition to vaccine coverage forecasting, a large-scale correlative study of the socio-economic correlates of immunisation coverage is performed. This study allows the changing strength of possible socio-economic drivers of vaccine coverage to be assessed. We study correlations between coverage and 190 socio-economic factors that include healthcare, education, and infrastructure variables. Countries are grouped by regional classifications according to the World Health Organization to understand how factors vary in their association with coverage across world region. This work extends the current literature by performing the broadest analysis of socio-economic links to vaccine coverage yet available and allowing insights into the how determinants of coverage vary over time and by world region. This identification of modulating factors can inform effective immunisation programmes and policies.

This chapter proceeds as follows: in Section 2.2 we provide an overview of the data sources used in this thesis and detail the data cleaning processes used: these data suffer from missing data and so a sensible missing data treatment is required (Section 2.2.3); in Section 2.3, the mathematical methods used to forecast vaccine coverage and perform a large-scale time-varying correlative study of determinants associated with coverage are presented; in Section 2.4.1 we introduce our vaccine performance index; findings from all analyses are presented in Section 2.4; and, finally, findings from our application of our vaccine performance index and correlative study are discussed in Section 2.5.

2.2 Data and data processing

2.2.1 COVERAGE

Childhood vaccination coverage data are obtained from WHO-UNICEF estimates of national immunisation rates [161]. We denote this dataset **COVERAGE**. Coverage is obtained for the following childhood vaccinations: BCG (tuberculosis), DTP1 and DTP3 (representing the first and third doses, respectively, of the diphtheria-tetanus-pertussis vaccine), MCV (measles-containing vaccine), and POL3 (the third dose of the polio vaccine programme). Vaccination

¹See SAGE’s definition of vaccine hesitancy in Section 1.2.1 and [58].

coverage data are collected between 1980 to 2013 across 190 countries around the world². (WHO also recommends three doses of the hepatitis B and Haemophilus influenzae vaccines; however, due to incomplete data records – partly due to these vaccines not being offered to many children in high-development regions (see, for instance, [162]) – we do not consider these programmes.) The level of DTP3 coverage and the difference between DTP1 and DTP3 coverage are often used as indicators of strength in a country’s immunisation programme due to the ability of national immunisation programmes to reach children with three doses in a series in a short time [10, 163].

Coverage is reported as the percentage of those immunised in a specified target group. For example, the BCG vaccination is typically reported as the percentage of children who survive birth and are immunised against BCG at, or soon after, birth (see [161] for details). Coverage is therefore expressed as a percentage between 0 and 100. When forecasting vaccination coverage in Section 2.3.1, we transform coverage using the logit transform,

$$y'(t) = -\log\left(\frac{100}{y(t)} - 1\right), \quad (2.1)$$

where $y'(t)$ denotes transformed coverage in year t . This transform maps original coverage values $y(t) \in [0, 100]$ onto the real line to avoid forecasting coverage to outside the range $[0, 100]$ as these transformed values lie in the range $(-\infty, +\infty)$. (Note that $y'(t)$ does not denote a derivative.)

2.2.2 GAPMINDER

The website Gapminder [164] draws socio-economic data from sources including the World Bank [165], the International Labour Organization [166], and WHO [167]. A broad range of quantitative factors are obtained that span the fields of economics, health care, industry, demographics, communications, infrastructure, physical geography, trade, and education.

A total of 513 socio-economic factors (SEFs) are initially obtained. After examining the quality of these data, we filter this initial set down to 190 SEFs. In this curation process 219 SEFs are removed because they have fewer than three data points for 136 countries. Unnormalised versions of SEFs are either normalised or removed (if their normalised version already exists). For example, **gross domestic product** and **national electricity use** are removed in favour of **GDP per capita** and **electricity use per person**. Unnormalised SEFs often have values spanning multiple orders of magnitude and are prone to outliers which can cause issues in correlative studies [168]. Five SEFs are normalised and kept; 104 outlier-prone SEFs are

²In Section 2.2.2, we introduce socio-economic data which is only available from 1980 to 2010. We obtain coverage data up to 2013 for the purposes of time-series forecasting – see Section 2.3.1.

removed. This final set of 190 SEFs is tabulated with missing data summary statistics in Appendix Table A.1. (We refer to this dataset as **GAPMINDER**.)

Countries are grouped into six world regions according to WHO classifications: Africa (AFR; 47 countries), the Americas (AMR; 35 countries), the Eastern Mediterranean (EMR; 21 countries), Europe (EUR; 53 countries), South-East Asia (SEAR; 11 countries), and the Western Pacific (WPR; 27 countries) [169] (see Figure 2.1 F and Appendix A.1.2). In Section 2.3.3, time-varying associations between national-level socio-economic factors and vaccine coverage for each world region are sought. It is therefore instructive at this stage to define the set of vectors $\mathbf{x}_{rst}$, where $r = 1, \dots, 6$ denotes world region, s denotes a socio-economic factor, and t denotes year. These vectors have length equal to the number of countries in a given world region (see Appendix A.1.2 for a full country list). Similarly, vaccination coverage for vaccine v in year t is given by $\mathbf{y}_{rvt}$.

Reducing the set of socio-economic factors from 513 to 219 (as described above) lowers the overall missing data fraction from 67.7% to 48.4%. This percentage varies substantially across SEFs, between countries, and over time (see Appendix Table A.1). This missing data fraction must be substantially reduced.

2.2.3 Imputing missing data

There are a number of approaches to replacing missing data values [170, 171]. Understanding the mechanism by which missing data are generated can motivate the optimal imputation strategy [172]. There are three mechanisms through which missing data can arise: missing data can be *missing completely at random* (MCAR), *missing at random* (MAR), or *missing not at random* (MNAR) [171]. Data are MCAR when the process that results in missing data is independent of the values of other non-missing data *or* to the values of missing data that *should* have been observed. When missing data depends on other variables in the data, but not to the values of missing data that should have been observed, these data are said to be MAR. When the missing data mechanism is related to the missing data that should have been observed, the missing data mechanism is MNAR.

As an example of each of these mechanisms, consider the following toy example using the socio-economic variable **aid received per person** (reported in US dollars; see Appendix Table A.1). Suppose that in every year a country decides whether to report the aid it receives based on the toss of a fair coin. Any missing data arising would be MCAR, as the process which generates the missing data is random and unrelated to other observed data and to the value of **aid received per person** which *should* have been observed. Suppose instead, that a country that receives no aid does not officially report **aid received per person**, then the fact that these data are missing is related to missing value itself: this is an example of data which is MNAR.

If, however, countries with a gross domestic product above a threshold were neither entitled to aid nor required to report it, then the missing data value is related to an auxiliary variable and the missing data mechanism is MAR.

When data are MAR or MNAR, ignoring the missing data mechanism (for example by deleting all missing data values) can result in biased statistics estimated from these data [172] and so a sensible imputation strategy is required.

Imputing socio-economic factors

There are a total of 36,100 time-series of socio-economic data ($190 \text{ countries} \times 190 \text{ variables}$). These time-series have different ‘patterns’ of missing data; for example, data might have only been collected since a certain year, or might only be collected and reported in every n^{th} year. As the number of socio-economic variables is large, an imputation method based on a parametric model between missing and non-missing data is challenging and so we seek an alternative.

We believe most time-series missing data to be MAR as other data (such as adjacent time-series points) are typically highly informative of the missing data value. When all data points within a time-series are missing, it might be supposed that these data are MNAR. However, in this case, the missing time-series may be well approximated from other countries’ time-series for this factor. These missing data patterns – which include the occasional missing time-series datum and entire missing time-series – suggest a combined imputation approach that accounts for time-series imputation methods and nearest-neighbour methods which can ‘borrow’ missing data based on similarities between countries’ socio-economic data records.

The combined imputation strategy which we develop (the full details of which are provided in Appendix A.1.3) uses linear interpolation and extrapolation to impute time-series data points when there are at least two data points in the time series; mean imputation when there is a single time-series datum [173]; and, in the case when there are no time-series data points, uses a K -nearest neighbours algorithm [174–176]. Our combined imputation strategy searches over different parameterisations of these individual strategies to find an optimum (see Appendix A.1.3 for imputation details and evidence supporting the robustness of this approach).

Imputing vaccine coverage data

Missing data are less prevalent in **COVERAGE** than **GAPMINDER**: only 14% of data are missing across all vaccines. BCG has the highest missing data fraction (26.6%) followed by MCV (13.1%), POL3 (11.1%), DTP3 (10.9%), and DTP1 (10.4%). Coverage data are typically missing between 1980-1990 (when some campaigns were newly-introduced). Coverage time-series are imputed using linear interpolation and extrapolation; however, as vaccine coverage

data tends to be more volatile than socio-economic data, it is best extrapolated based on a higher number of time-series data points than socio-economic time-series (see Figure 2.1 C and Appendix A.1.3).

2.2.4 Exploratory analyses

After data cleaning and imputation, exploratory data analyses are performed to understand the structure and characteristics of the `GAPMINDER` and `COVERAGE` datasets.

Trends in immunisation coverage

Vaccine coverage rates have increased across all six WHO regions between 1980 and 2010. In 2010, coverage for all vaccines except BCG is highest in the European region. The Americas, South-East Asia, and Western Pacific regions have a higher overall BCG coverage than Europe in 2010 (Figure 2.1 A). In 2010, vaccine coverage is lowest in the Africa region for all vaccines. High mutual correlations exist between the DTP1, DTP3, MCV, and POL3 vaccines, but BCG is not highly correlated with this group (Figure 2.1 B). Time-series of DTP3 coverage across individual countries show large contrasts in both coverage magnitude and variability (Figure 2.1 C).

Trends and structure in socio-economic factors

Trends in socio-economic time-series are illustrated with two examples. In Figure 2.1 D and E trends in countries' female education and rural access to water, respectively, are shown. With the exception of naturally volatile socio-economic factors, such as those based on natural disaster incidents and deaths, most SEFs are smoothly time-varying. In Figure 2.1 G, a scatter plot of countries' 1982 DTP3 coverage rates versus their education levels (in the same year) for three world regions is shown. This plot illustrates the regional correlations we perform in Section 2.3.3: correlations are obtained between a single SEF (across all countries in a given region) and DTP3 coverage for each year and averaged over time (see Section 2.3.3).

Structure in socio-economic data is investigated using *t*-distributed stochastic neighbour embedding (*t*-SNE) [177]: a dimensionality reduction method used for data visualisation. Although other dimensionality reduction methods for data visualisation exist, such as principal components analysis [178], multi-dimensional scaling [179], and Sammon mapping [180], *t*-SNE is used for its improved data visualisation quality of high-dimensional data over similar algorithms [177]. A rough overview of *t*-SNE in the context of visualising `GAPMINDER` and `COVERAGE` data is now provided.

t -SNE projects high-dimensional data into a lower dimensional space through minimisation of a cost function defining the quality of the mapping [177]. A similarity matrix $S_{ss'}$ representing the ‘closeness’ between each socio-economic factor in the high-dimensional space is first obtained,

$$S_{ss'} = 1 - \frac{1}{31} \left| \sum_{t=1}^{31} \rho(\mathbf{x}_{st}, \mathbf{x}_{s't}) \right|, \quad (2.2)$$

where $\mathbf{x}_{st}$ is a vector of all countries’ socio-economic factor s in year t ($= 1, \dots, 31$), and $\rho(\cdot, \cdot)$ is the Spearman’s rank correlation coefficient³. (Vaccination coverage data $\mathbf{y}_{vt}$ is also included when calculating the similarity matrix). The probability of two factors being closest neighbours is assumed to be proportional to the similarity matrix,

$$p_{ss'} \propto S_{ss'}. \quad (2.3)$$

The aim is to map these probabilities onto a two-dimensional vector space $\mathcal{Y}$. Another probability kernel is used in this low dimensional space to model the similarity of vectors $\mathbf{y}_s, \mathbf{y}'_s \in \mathcal{Y}$. In t -SNE, the t -distribution is used for this probability distribution,⁴

$$q_{ss'} \propto (1 + \|\mathbf{y}_s - \mathbf{y}'_s\|^2)^{-1}. \quad (2.4)$$

The low dimensional vectors $\mathbf{y}_s$ are obtained by minimising a cost function between the high- and low-dimensional space. Stochastic neighbour embedding uses the Kullback-Leibler (KL) divergence [182] to measure the data faithfulness between these high- and low-dimensional spaces [181]. The cost function is a sum over all KL divergences,

$$C = \sum_{ss'} p_{ss'} \log \frac{p_{ss'}}{q_{ss'}}. \quad (2.5)$$

When $p_{ss'} \approx q_{ss'}$ this cost function is small. When $p_{ss'}$ is large then there is a large penalty unless $q_{ss'}$ is also large. This KL distance is therefore appropriate for a local structure-preserving map. The heavy tail of the t -distribution further accentuates large distances in the high-dimensional space, whilst maintaining short distances [177].

Application of t -SNE shows strong links between socio-economic factors in similar categories⁵ and reveals the mutual correlation structure within the vaccine coverage data (Figure 2.1 H). BCG, for example, is closer in the reduced mapping to factors relating to health spending, whereas the other four vaccines seem to be more closely aligned with schooling

³See Equation (2.12).

⁴In the original stochastic neighbour embedding paper [181] a Gaussian distribution was used.

⁵These categories are heuristically chosen.

variables.

In accordance with GAVI’s use of DTP3 as a marker of strength in immunisation programmes [163] (and because of the high mutual correlations between vaccines), the analyses in this chapter henceforth focus on DTP3 coverage.

2.3 Statistical methodology

2.3.1 Forecasting vaccine coverage rates

Vaccine coverage time-series show wide-ranging characteristics between countries that include smooth and highly volatile time-series (Figure 2.1 C). Forecasting such data therefore requires a versatile method that takes into account this range of characteristics. Further, this method should provide a way to account for uncertainty in the data (which might be time-varying) and should return probabilistic distributions over predicted values for interpretability.

In this chapter, we consider fitting *Gaussian processes* (GPs) [183] to coverage time-series data. Gaussian processes offer many advantages: for example, probabilistic predictions, model flexibility, and a generalised modelling framework [183, 184]. Whilst *parametric* models may require a parametrised class of functions to be selected *a priori*, GPs are *non-parametric* as they involve inference over an infinite-dimensional function space [183, 185]. GPs are therefore appropriate for national vaccine coverage forecasting, as little is known about the functional forms of coverage time-series data, which seem to differ markedly between countries.

Gaussian processes

Gaussian processes are defined as a collection of random variables with the specific property that any finite subset of these random variables has a multivariate Gaussian distribution [183]. A Gaussian process is then written as a draw from a multivariate normal distribution,

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K}) \quad (2.6)$$

where $\{f_t, t \in T\}$ is the Gaussian process, $\boldsymbol{\mu} = \mathbb{E}[\mathbf{f}]$ is the mean of the GP and $\mathbf{K} = \text{cov}(\mathbf{f}, \mathbf{f}')$ is a covariance matrix. The covariance matrix can be obtained through a *covariance function* $k(\cdot, \cdot)$, such that $K_{ij} = k(t_i, t_j)$. Temporal trends in the Gaussian process are modelled via the mean and covariance functions.

Many choices exist for the covariance function [183]. We choose the squared-exponential (SE) kernel as it allows a wide range of time-series to be modelled (see Figure 2.2). The SE




Figure 2.1: Exploratory analyses of socio-economic factors and illustrative dynamics of SEFs and vaccine coverage. (A) Vaccine coverage levels have increased across WHO regions since 1980. Leftmost circle represents 1980 coverage levels; rightmost is 2010. Coverage is averaged across all countries in a region. (B) Spearman correlation between vaccines across all countries and years. (C) Diversity of behaviours between countries' time-series. (D) Time-series of the SEF rural water access. (E) time-series of the SEF female education. (F) World regions as defined by WHO. (G) Scatter plot illustrating the link between DTP3 coverage and maternal education in 1982, coloured by WHO-region. (H) Visualisation of high-dimensional SEF matrix via *t*-SNE (described in text) reveals pronounced structure throughout the set of SEFs considered. Heuristic categories are assigned before using *t*-SNE.




Figure 2.2: **Diversity of Gaussian processes using the squared-exponential kernel** (A) Squared-exponential (SE) kernel (Equation (2.7)) with $\tau_c = 0.01$ and $s = 0.5$. (B) SE kernel with $\tau_c = 0.1$ and $s = 0.5$. (C) SE kernel with $\tau_c = 0.5$ and $s = 0.5$. The SE kernel is able to capture a wide range of time-series properties, including quickly-varying processes that can capture sudden drops in coverage (A) and slowly-varying long-term trends (B) including near-constant time-series (C).

kernel is defined as,

$$k(t_i, t_j) = s^2 \exp \left(- \frac{(t_i - t_j)^2}{2\tau_c^2} \right), \quad (2.7)$$

where τ_c is the characteristic time scale of the GP and s^2 is a scale factor. As SE kernels are stationary [183], a non-zero mean function is specified to model long-term time trend in the coverage time-series. We opt for a linear mean function $\mu(t) = mt + c$, with m and c the gradient and intercept, respectively. This linear function should capture the broadly linearly increasing time-series trends in transformed coverage space.

Gaussian process regression

Vaccine coverage is modelled as,

$$\mathbf{y}' = \mathbf{f} + \boldsymbol{\epsilon}, \quad (2.8)$$

where $\mathbf{y}'$ is a time-ordered vector of logit-transformed coverage (as defined in Section 2.2.1), $\mathbf{f}$ is a Gaussian process, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{1})$ is observation noise [183].

Fitting GPs to observed vaccination coverage data proceeds in a Bayesian framework by maximising the posterior density,

$$\begin{aligned} p(\boldsymbol{\theta} | \mathbf{y}', \mathbf{t}) &\propto p(\mathbf{y}' | \mathbf{t}, \boldsymbol{\theta}) p(\boldsymbol{\theta}) \\ &= \int p(\mathbf{y}' | \mathbf{t}, \mathbf{f}) p(\mathbf{f} | \boldsymbol{\theta}) d\mathbf{f} \times p(\boldsymbol{\theta}), \end{aligned} \quad (2.9)$$

where the integral in Equation (2.9) is the (marginal) model likelihood (integrated with respect to the Gaussian process $\mathbf{f}$) and $p(\boldsymbol{\theta})$ is the prior over the parameters to be inferred, $\boldsymbol{\theta} =$

$\{l, \sigma^2, s^2, m, c\}$. The likelihood integrand comprises $\mathbf{y}'|\mathbf{t}, \mathbf{f} \sim \mathcal{N}(\mathbf{0}, \sigma^2)$ and $\mathbf{f}|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K})$.

Maximum a posteriori estimates of the parameters are obtained through the GPML package in MATLAB [186]. The performance of Gaussian process regression is assessed with respect to a linear regression benchmark in Appendix A.3.1; GPs are shown to be predictively superior.

Prior selection

The same prior distributions are chosen for each country and the choice of priors are motivated by the scales involved. All prior distributions are specified for transformed coverage $y'(t)$. The time-scale τ_c over which coverage data interacts is likely to be a few years, and so we use the prior $\tau_c \sim \mathcal{N}(0, 2^2)$. Due to the large range of coverage starting values in 1980, we centre the intercept at 95% (which is $-\log(100/95 - 1) \approx 3$ in transformed space) but use a large variance to account for the wide variability in values here, thus we choose $c \sim \mathcal{N}(3, 4^2)$. The prior of the gradient m is centred at 0 to account for both growing and declining coverage and a large variance is chosen to account for large growths or decays in coverage, thus we specify $m \sim \mathcal{N}(0, 1^2)$ (the variance here is chosen with the range of time – 31 time points – and coverage extremes (which range from approximately -5 to +5 in transformed space) in mind.

Reasonably uninformative prior variances are chosen to allow for a large variability in both observation noise and scale factor: we choose $\log \sigma^2 \sim \mathcal{N}(0, 2^2)$ and $\log s^2 \sim \mathcal{N}(0, 2^2)$. (We note that $p(\boldsymbol{\theta})$ is thus expressed through independent priors for each parameter.)

2.3.2 A Vaccine Performance Index

Gaussian processes provide a probabilistic distribution over values forecasted from the mean function. This distribution is known as the predictive distribution and is given by,

$$y'_*|t_*, \mathbf{t}, \mathbf{y}' \sim \mathcal{N}(\bar{y}'_*, \text{cov}(y'_*)), \quad (2.10)$$

where $\bar{y}'_* = \mu(t_*) + \mathbf{k}_*^T [\mathbf{K} + \sigma^2 \mathbf{1}]^{-1} (\mathbf{y}' - \boldsymbol{\mu})$, and $\text{cov}(y'_*) = k(t_*, t_*) - \mathbf{k}_*^T [\mathbf{K} + \sigma^2 \mathbf{1}]^{-1} \mathbf{k}_*$ with $k_{*i} = k(t_*, t_i)$ [183].

This predictive distribution over forecasted values is used to define a summary measure of the strength of nation's immunisation programmes. We term this measure the **vaccine performance index** (VPI). The VPI thus uses previous trends in mean and variability (via the GP model) to balance the desirable probability of high future coverage levels against the undesirable probability of experiencing drops in coverage. We define the VPI as the probability coverage in a future year t_* , $y_* = y(t_*)$, will exceed a threshold value ν less the probability that this future year's coverage will drop more than a value d . The VPI is

therefore expressed as,

$$\text{VPI}_{\nu,\tau,d}(t_*) = P(y_* > \nu) - P(y_* < y(t_N) - d), \quad (2.11)$$

where $y(t_N)$ is the most recent year for which coverage was used in the forecast and we note that we have mapped transformed coverage back to $[0,100]$.⁶ The index τ on the index refers to how far in the future the forecast is, $t_* = t_N + \tau$. Under all parameterisations, the VPI is between -1 ($P(y_* > \nu) = 0$ and $P(y_* < y(t_N) - d) = 1$) and +1 ($P(y_* > \nu) = 1$ and $P(y_* < y(t_N) - d) = 0$), representing the two extremes of dynamic behaviour.

This expression is re-weightable if the reader or policy maker wishes to focus on a different combination of strategic properties. There are more natural parameterisations of this index than others. For example, to assess the GVAP goal of 90% coverage in all countries by 2015, it is reasonable to consider $\tau = 2$ (since coverage data are available until 2013 – see Section 2.2.1), $\nu = 90$, and $d \rightarrow \infty$, thus using the index to forecast the probability that coverage in 2015 will exceed 90% (as the final term in Equation (2.11) becomes zero). Alternative forms of the VPI are discussed in Appendix A.3.2.

2.3.3 Time-varying correlative study

Spearman’s rank correlation coefficient is a non-parametric measure of correlation that is robust to outliers, model-free, and can compute correlations with continuous and discrete variables [187, 188]. As we consider a large number of socio-economic factors – which may be discrete or continuous and whose relationship form with coverage is unknown – these properties are appealing. Spearman’s rank correlation coefficient between two vectors $\mathbf{x}$ and $\mathbf{y}$ is given by

$$\rho(\mathbf{x}, \mathbf{y}) = \frac{\text{cov}(\mathbf{r}_x, \mathbf{r}_y)}{\sigma_{\mathbf{r}_x} \sigma_{\mathbf{r}_y}}, \quad (2.12)$$

where $\mathbf{r}_x$, and $\mathbf{r}_y$ are ranked vectors for $\mathbf{x}$ and $\mathbf{y}$, respectively (this equation is equivalent to the Pearson correlation between ranked data) [189], and $\sigma_{\mathbf{r}_x}$ denotes the standard deviation of the ranked vector $\mathbf{r}_x$ (and similarly for $\mathbf{r}_y$).

Temporal trends in association between a coverage programme v and an SEF s are mod-

⁶This is done via the equations,

$$P(y(t_*) > \nu) = \int_{-\infty}^{-\log(\frac{100}{\nu}-1)} p(y'_*|t_*, \mathbf{t}, \mathbf{y}') dy'_*,$$

and

$$P(y(t_*) < y(t_N) - d) = \int_{-\log(\frac{100}{y(t_N)-d}-1)}^{\infty} p(y'_*|t_*, \mathbf{t}, \mathbf{y}') dy'_*.$$

elled by taking time-averaged Spearman coefficients between all countries within a region,

$$\bar{\rho}_{rsv} = \frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} \rho(\mathbf{x}_{rst}, \mathbf{y}_{rvt}), \quad (2.13)$$

$\bar{\rho}_{rsv}$ is the time-averaged Spearman's correlation for socio-economic factor s with vaccine coverage programme v for WHO region r , and where t_2 is the most recent year and t_1 the most distant. In Section 2.4.2, we compare historic trends to recent ones by considering $t_2 = 2010$, and $t_1 = 2001$ (recent) and $t_1 = 1980$ (historic).

As we consider a large number of correlations, we must guard against the increased chance of committing a type I error⁷ by performing multiple hypotheses corrections [190]. Each correlation between coverage and a socio-economic factor in a given year returns a p -value. An overall p -value is required across the multiple averages calculated (31 when $t_1 = 1980$, and 11 when $t_1 = 2001$). Many approaches to multiple hypothesis testing assume independent p -values [191] which may result in an overly-penalised corrected p -value. To combat this issue, we develop a meta p -value from the 31 ($t_1 = 1980$) or 11 ($t_1 = 2001$) p -values obtained in each time-average based on an approach introduced by Delongchamp *et. al* [192]. These meta p -values are then corrected using the Bonferroni method [190]. The meta p -value derivation and correction method is outlined in Appendix A.2.1.

As we consider a total of 190 SEFs, a mechanism is required to focus on the strongest signals. To obtain the most robust signals, the top-ten correlating socio-economic factors within each WHO region across both time intervals are selected. In this top-ten list, many factors will not be significantly correlating with vaccine coverage. To address this issue, SEFs are replaced from this top-ten list if they have a corrected meta p -value less than 0.01 (see Appendix A.2.1). In addition to these top ten SEFs, other factors are included that are related and useful for comparison: we include **primary school completion rate (male)**, for example, for a direct comparison with **primary school completion rate (female)** (see Appendix A.2.2 for full details on this replacement process).

2.4 Results

2.4.1 Country trends and forecasts with the VPI

To compare countries' national immunisation programme strength we parameterise the VPI in Equation (2.11) with $\nu = 95\%$ and $d = 2$. These choices are motivated by the 95%

⁷A type I error is the incorrect rejection of a null hypothesis [190].

vaccination rate required for herd immunity⁸ and because a drop of $d = 2$ weighs the desirable property of high coverage against an undesirable drop.

In Figure 2.1 A, world maps of the VPI values under the parameterisation above in 2001, 2005, 2009, and 2013 are shown. Consistently high-performing countries include Saudi Arabia, Russia, Hungary, and Slovakia (Figure 2.1 A; and see the table of VPI values in Appendix A.3.3). The UK, USA, Canada, Australia, and New Zealand, have historically low VPI values; whilst VPI values have improved in the UK, USA, and Canada, they remain low in Australia and New Zealand. More generally, North America, Europe, and much of Asia have shown steady progress since 2005; by contrast, South America, sub-Saharan Africa, and the Indian subcontinent have shown little improvement. Many European countries have been unable to increase coverage levels to values in excess of 95%. Romania (89%), Iceland (91%), Norway (83%), and Ukraine (94%) have VPI values far from +1, signifying coverage rates still below 95% and showing little sign of gaining above this threshold. Africa shows mixed results. North African countries (in the EMR region) have high VPI values. Many other non-African Eastern Mediterranean countries show clear improvement in the past decade and high current VPI values (such as Iran, Saudi Arabia, Oman), although Syria, Iraq, Pakistan, Afghanistan, and Yemen are notable exceptions. Sub-Saharan Africa (countries within the AMR definition) have low VPI values. These values are due to high coverage variability and low mean coverage rates. There are some exceptions, however. Rwanda, Burundi, Botswana, and Zimbabwe have high VPI values representing high and stable DTP3 coverage rates. South-East Asia performs similarly to Africa. Many countries have vaccine performance index values close to -1 (Figure 2.1 A and Appendix A.3.3) suggesting a high probability of declining coverage in future years.

In Figure 2.1, Gaussian process forecasts for the five countries with the highest increases (B) and decreases (C) in VPI values between 2009 and 2013 are displayed. Malta has the most improved VPI values between 2009 and 2013 (Figure 2.1 B): DTP3 coverage decreased sharply between 2004 and 2007 (falling to a low of around 70%) though has since recovered to over 95% in 2013. Romania has the most deteriorated VPI values, as high and steady values in excess of 95% have decreased to sub 90% in 2013.

Assessing GVAP's 90% coverage goal

An alternative form of the VPI is used to assess the GVAP goal of 90% coverage in all countries by 2015. To do this we forecast vaccine coverage data into 2015 (2 years after the most recent available datapoint, in 2013) and calculate the probability that DTP3 coverage will exceed 90%. We therefore take $\nu = 90$, $\tau = 2$, and $d \rightarrow \infty$ in Equation (2.11)). VPI

⁸Although we choose this threshold as this value is often cited as the level required to maintain herd immunity, we note that this is disease specific [136] and policy makers can choose ν accordingly





Figure 2.3: **A Vaccine Performance Index highlights regions with a large probability of variable coverage far from the 95% threshold.** (A) The performance index (VPI; Equation (2.11)) for all countries for $t_N = 2001, 2005, 2009$, and 2013 . Sub-Saharan Africa, parts of Southern and South-East Asia and some European countries have low VPI (despite many having high coverage; in these cases, volatile coverage history decreases the VPI). Black represents missing data. (B) Countries with the largest improvement in VPI between 2009 and 2013 (Malta is most improved). (C) Countries with the largest decrease in VPI over the same period (Romania has the most deteriorated performance). WHO-UNICEF coverage estimates are in blue. For 2013's VPI value, coverage data from 1980 to 2012 is used. The mean of the Gaussian process in black, and shaded 95% confidence intervals for training (blue) and forecasted (red) data. We show 2013 as the red vertical line. VPI values can be found in Appendix A.3.3



$0.5 \leq \text{VPI} < 0.9$: Angola, Argentina, Armenia, Australia, Bosnia and Herzegovina, Brunei Darussalam, Burkina Faso, Cameroon, Chile, Colombia, Costa Rica, Cuba, Denmark, El Salvador, Honduras, Iceland, Jamaica, Kiribati, Malawi, Nauru, Nepal, Norway, Palau, Paraguay, Peru, Romania, Samoa, Senegal, Trinidad and Tobago, Tuvalu, United Republic of Tanzania, Viet Nam, Yemen, Zimbabwe.

$0.1 \leq \text{VPI} < 0.5$: Barbados, Comoros, Korea, Democratic People's Republic of, Djibouti, Dominican Republic, Lao People's Democratic Republic, Liberia, Mauritania, Micronesia, Federated States of, Namibia, Panama, Republic of Moldova, San Marino, Solomon Islands, Suriname, Togo, Ukraine.

$\text{VPI} < 0.1$: Afghanistan, Austria, Benin, Central African Republic, Chad, Congo, Cote d'Ivoire, Democratic Republic of the Congo, Equatorial Guinea, Ethiopia, Gabon, Guinea, Guinea-Bissau, Haiti, India, Indonesia, Iraq, Kenya, Lebanon, Madagascar, Mali, Marshall Islands, Mozambique, Burma, Niger, Nigeria, Pakistan, Papua New Guinea, Somalia, South Africa, Syrian Arab Republic, Timor Leste, Uganda, Vanuatu, Venezuela, Zambia.

Figure 2.4: **Extrapolation of current trends based on our Gaussian process model reveal countries which may fail to reach GVAP targets** (A) 54 countries have more than a 50% chance of failing to meet GVAP 2015 targets if current DTP3 trends continue. (B) Histogram of VPI values illustrating a gulf in vaccine performance.

values for these parameters are mapped in Figure 2.4 A.

A histogram of VPI values yields a bimodal distribution (Figure 2.4 B). This bimodality suggests many countries are set to either comfortably reach this goal or set to fall considerably short. There are a few dozen countries falling within these peaks and include many South American and countries in southern Africa. European countries generally perform better than other regions (with the notable exceptions of Denmark, Iceland, Romania, Austria, Moldova, San Marino, and Ukraine) and have – under our model – an almost certain chance of reaching the GVAP goal. Countries in sub-Saharan Africa and the Indian subcontinent (with only few exceptions) seem likely to drastically miss the GVAP target.

2.4.2 Regional trends in determinants of DTP3 coverage

In Figure 2.5 A, the most strongly correlating socio-economic factors with vaccine coverage over the periods 1980 to 2010 and 2001 to 2010 are shown. Correlation strengths differ between WHO region and vary over time. The two most consistently-correlating SEFs (p -values < 0.05 over the largest number of regions) are **births attended by skilled health staff**⁹ and **access to sanitation**¹⁰. In AFR, AMR, and EUR, correlation magnitude is high between 1980 to 2010, but lower recently. In contrast, SEFs have remained correlated over time in the EMR and the WPR regions.

In AFR and AMR, national levels of **access to water/sanitation**¹¹ and **births attended by skilled health staff** are informative of a country's DTP3 immunisation rate. In AFR and EMR, educational variables, such as **primary completion rates**¹² and **education ratio**¹³, and **access to water** are highly informative of a country's DTP3 level. In EMR alone, **government health spending**¹⁴, income metrics (such as **GDP per capita**¹⁵ and **GNI per capita**¹⁶), **access to sanitation**, **phone use**¹⁷, and **CO2 emissions**¹⁸ are strongly linked with coverage. Historically, socio-economic correlates are low (across both time periods) in EUR, and are low between 2001 to 2010 AMR. In SEAR, **education ratio** and **access to sanitation** are the only two significant correlates at the $p < 0.05$ level between 1980-2010 and only **rural sanitation access**¹⁹ is significant between 2001 to 2010. In WPR, **births attended by skilled health staff** is the strongest correlate. EMR is the only region where **gross domestic product per capita** is the strongest correlate of DTP3 coverage. In all other regions the top correlating factor is typically related to education or access to water or sanitation. This correlation study was repeated for other childhood vaccines; see Appendix A.2.3 for figures and a discussion.

In our correlative analysis, a significant effect was noted in the difference between the correlation strength of male and female education level in AFR and SEAR. This effect is shown in Figure 2.5 B. We also note the decreasing strength of the tie between a country's DTP3 coverage and its education level in AFR, AMR, EMR and EUR.

⁹Percentage of births attended by a skilled health personnel from the World Bank [193] accessed via Gapminder [164].

¹⁰Percentage of population with access to sanitation facilities [194], via [164]

¹¹Percentage of population with access to improved water sources [195] via [164].

¹²Percentage of those completing primary education [196] via [164]

¹³Ratio of girls to boys in primary and secondary education [197] via [164]

¹⁴Central government expenditure on health as a percentage of total government expenditure [198] via [164]

¹⁵Gross domestic product per capita [199] via [164]

¹⁶Gross national income per capita [200] via [164]

¹⁷Fixed line and mobile phone subscribers (per 100 people) [201] via [164]

¹⁸Per capita carbon dioxide (CO2) emissions (metric tonnes of CO2) [202] via [164]

¹⁹Percentage of rural population using improved sanitation facilities [203] via [164]




Figure 2.5: Historic and recent socioeconomic correlates with DTP3 coverage (A) Strength of time-averaged correlation as outlined in Equation (2.13) for 1980-2010 (dark colours) and 2001-2010 (light colours) with meta p -values (black for 1980-2010, red for 2001-2010: *** denotes $p < 0.001$, ** denotes $p < 0.01$, * denotes $p < 0.05$, and $\diamond$ denotes $p < 0.10$). Error bars represent 95% CIs. The selection process of SEFs is described in Section 2.3.3 (B) Time-series of Spearman's ρ between DTP3 coverage and female (blue) and male (red) education. GDP=gross domestic product. GNI=gross national income. AFR=Africa. AMR=Americas. EMR=Eastern Mediterranean. EUR=Europe. SEAR=South-East Asia Region. WPR=Western Pacific region. DTP3=diphtheria, tetanus, and pertussis third dose.

2.5 Discussion

In this chapter, we introduced the vaccine performance index: a quantitative measure of the strength of a nation’s immunisation programme based on previous trends and volatility in DTP3 coverage. In countries with advanced healthcare systems where there is universal access to vaccines, this index can be interpreted as a measure of vaccine hesitancy. Monitoring hesitancy and trust in immunisation programmes has been recommended by the public health community [153, 154].

In 2016, a WHO report [204], noted that 65 countries had failed to achieve either “90% coverage nationally or 80% coverage in every district or equivalent administrative unit against [DTP] by 2015”. This report is in line with predictions from our vaccine performance index which identifies 36 countries with a VPI < 0.1 (that is, who have less than a 10% chance of achieving more than 90% coverage in 2015) and a further 17 with a less than 50% chance (see Figure 2.4). We thus suggest that our index can be used to identify countries failing to reach these targets at least two years ahead of the goal’s deadlines, highlighting the importance and utility of quantitative forecasts introduced in this chapter.

Review of historical trends and variability in vaccination rates can aid policy makers’ assessments of the strength and resilience of immunisation programmes. Although there have been some recent attempts to compare determinants of vaccine coverage between countries (for example, across East African countries [92] and sub-Saharan Africa [88]), there has been no attempt to identify time-varying correlates on a global scale. Our global-scale correlative study finds that countries with high fractions of people with access to skilled healthcare providers tend to have high coverage rates – a trend which is significant for four world regions (AFR, AMR, EMR, and WPR). This result appears to be in agreement with a systematic review of determinants of coverage [205] which finds a strong link between out-of-hospital birth (likely highly correlated with access to skilled health staff) and coverage. Moreover, this review also finds a strong tie between private health-care services with low immunisation levels; a result in line with our finding of the link between government health spending and coverage in the European and Eastern Mediterranean regions.

In country-specific contexts, literature from the past few years has identified several socioeconomic and attitudinal barriers to an *individual’s* vaccination records, such as parental education level [78–82], age [80, 101, 206], employment status and workplace [78, 80, 81], poverty [207–210], and distance to health-care facilities [78, 80, 101, 207, 210, 211]. Our time-varying analysis also finds many of these factors to be informative of a *nations’* immunisation rate, with a country’s overall education level, amount of access to sanitation and water, and fraction of births attended by skilled health staff all informative of a country’s immunisation

levels. Many of these factors point to access-related barriers: **access to sanitation** and **water** are likely proxies for poverty; similarly, births attended by skilled health staff could also be a poverty indicator or an indicator of the average distance people have to go for healthcare. Many factors, however, paint a more nuanced picture: wealth-based correlates in EMR suggest a financial barrier to vaccines and education-based correlates in multiple regions suggest that education could influence personal decisions (though of course, these could also be access – rather than hesitancy – barriers).

The declining strength of socio-economic correlates in AMR and EUR leads us to speculate that there is an emerging trend in regions with high access from access-based coverage losses to hesitancy-based coverage losses. In these regions, we suggest that socio-economic factors are less important in determining vaccine coverage as personal rejection becomes the primary barrier. Conversely, when strong correlations between coverage and SEFs exist, we believe that these determinants are informative of regional-access barriers to vaccination. In EUR, for example, access to vaccines is high, and there are non-infrastructural barriers to vaccination uptake that centre on religious beliefs and perceived risk of vaccines [36, 53]. In AMR, where socio-economic links to coverage are low and declining, we suggest the beginning of this transitional stage wherein vaccines in the region are broadly accessible, but that some factors (including attendance of medical staff at births) are restricting uptake of DTP3. In regions where correlation strength is increasing, we suggest that socio-economic determinants can be targeted for intervention. For example, the ratio of girls to boys in primary education, public health spending, and sanitation in SEAR. In AFR, the presence of skilled health staff could be targeted (a potential surrogate, children born in a public or private health-care facility, is consistently linked to completion of routine vaccinations [101]). We note, however, that further analyses would be required to infer causation between factors and coverage.

Although regional variability reveals an overall picture of trends at the WHO region, there is likely to be large sub-national variability in these correlates: this is a key limitation of the GAPMINDER dataset.

Further developments to the work in this chapter can be made: improvements to the accuracy of our GP model could be sought through the inclusion of socio-economic or media-sentiment data, for example, and forecasts could be extended to sub-national levels would allow an evaluation of regional vaccination targets [142].

In the next Chapter, we seek to obtain a more concrete handle on global attitudes to vaccines by investigating large-scale vaccine attitude survey data from 67 countries around the world.

3 | The state of vaccine confidence 2016: global insights through a 67-country survey

3.1 Introduction

A measure of fluid injected
into vein, possibly (7)

— Hawthorn, *Times* quick
crossword No. 853

In Chapter 2, trends in immunisation coverage rates are linked to trends in socio-economic determinants and are forecasted to provide probable ranges of future national coverage levels. These forecasted levels are used to evaluate the performance of national immunisation programmes. Although this index may provide key insights into the levels of vaccine hesitancy in countries where access to vaccines is not a barrier, this measure is indirect and cannot yield information on individual perceptions towards a range of vaccine-related topics. Obtaining and assessing vaccine-specific views is the focus of this chapter.

Individuals' perceptions of vaccines can greatly affect vaccine uptake rates but, despite virtues of routine monitoring of attitudes to immunisation programmes, there are only a small number of ongoing quantitative monitors of attitudes to vaccines [119, 212–214]. This lack of monitoring obfuscates a rigorous understanding of hesitancy and hinders our ability to forecast coverage based on attitudinal factors. Regular monitoring of attitudes can inform and optimise proactive public health interventions and can enable more sophisticated vaccine coverage forecasts. It is thus unsurprising that regular attitude monitoring has thus been cited as a public health priority [58, 154, 215].

The London School of Hygiene and Tropical Medicine's Vaccine Confidence Project¹

¹www.vaccineconfidence.org/

(VCP) is assuming the mantle of establishing routine attitude monitors. The central aim of the VCP is to “...monitor public confidence in immunisation programmes by building an information surveillance system for early detection of public concerns around vaccines.” In 2014, research led by the VCP surveyed almost 20,000 respondents across Georgia, India, Pakistan, Nigeria, and the UK to monitor attitudes to vaccines and perceptions of health services amongst parents of children under five and the general population as a whole [216]. This survey was the first of its kind to probe a broad range of attitudes to vaccine programmes at the national scale across multiple countries. In this chapter, we present findings from a more recent survey – conducted by WINGIA in collaboration with LSHTM – which greatly extends the scope the of this previous study.

We believe that this survey is the largest ever conducted on vaccine attitudes in terms of both the number of participants and the number of countries surveyed: 70,000 individuals across 67 countries are interviewed on their perceptions towards the importance of childhood vaccines, the safety and effectiveness of vaccines, and the extent to which vaccines are compatible with their religious beliefs. In addition to these questions, a range of socio-economic data are collected for each respondent. We use these socio-economic data to explore differences in attitudes to vaccinations between individuals and – using socio-economic factors from Chapter 2) – we also explore differences between countries.

At this stage it is worth highlighting the difference between independent attitudes to vaccines – which explore perceptions towards a range of vaccine related topics (and which are the focus of this chapter) – and vaccine *hesitancy*, which is a complex combination of attitudinal factors ([58] and see Section 1.2.1). In the following chapter, we implement a timely extension to the research presented here by inferring an individual’s *vaccine hesitancy* based on their responses to survey data in this chapter.

We begin this chapter with an introduction to the survey data collected by WINGIA. Hierarchical regression models – which enable variation between individuals and between countries to be explored – are introduced in Section 3.3. Results are then presented: worldwide attitudes to vaccinations are mapped (Section 3.4.1) and global determinants of attitudes are established (Section 3.4.2). Finally, the results are discussed in Section 3.5.

3.2 Data

3.2.1 GLOBAL-67

Every year since 1977, WIN/Gallup International Association (WINGIA) have conducted their *End of Year* global survey monitoring the public’s view on beliefs, happiness, and economic and social expectations [217]. In 2015, in partnership with the Vaccine Confidence

Project at the London School of Hygiene & Tropical Medicine, WINGIA extended this 14-question survey to include four questions on attitudes to vaccinations:

1. Overall I think vaccines are important for children to have;
2. Overall I think vaccines are safe;
3. Overall I think vaccines are effective; and,
4. Vaccines are compatible with my religious beliefs.

Each vaccination question is answered on the five-point Likert scale: **strongly disagree**, **tend to disagree**, **do not know** (and no response), **tend to agree**, or **strongly agree**. No response is coded with **do not know**. Only 1.9% of respondents provided this response to all four vaccination statements: as most of these 1.9% also provide non-missing responses to sensitive income and religion questions, we assume that this percentage is likely to be a strict upper bound on the missing vaccine survey questions. We therefore henceforth assume that these responses – jointly coded as **do not know/no response** – therefore represent **do not know** and are not missing.

In addition to vaccination data, a number of socio-economic characteristics for each individual are collected (factor levels are stated and baseline categories are underlined):

1. Sex: male, and female (SEX);
2. Age: 18-24, 25-34, 35-44, 45-54, 55-64, and 65+ (AGE);
3. Education: none, primary, secondary, university, and masters/PhD (EDU);
4. Income: first, second, third, fourth, and fifth (quintiles) (INC);
5. Work type: employed (full and part-time), unemployed, student, retired/disabled, and housewife (EMP); and,
6. Religion: Roman Catholic, Protestant, Russian/Eastern-Orthodox, Muslim, Jewish, Buddhist, Hindu, atheist/agnostic, other Christian, and other (REL).

A total of 65,819 persons were interviewed globally across 67 countries. Roughly 1,000 individuals are interviewed in each country (a maximum count of 2,080 in Afghanistan, a minimum of 398 in Algeria, and a median of 1,000) comprising roughly 50% males and females. Surveys were conducted either face-to-face (30,230 respondents in 28 countries), by telephone (13,062 in 16), or online (22,527 in 23). Refused or otherwise missing socio-economic data are also coded with **do not know** thus complicating the data cleaning process: dealing with this data ambiguity is discussed in Section 3.2.2.

3.2.2 Country-level variables

Country-level variables are used to account for variance in attitudinal response data between countries. Key country-level factors from Chapter 2 are therefore augmented with the GLOBAL-67 dataset. These variables are:

1. Health Expenditure (% of GDP) (HEX);
2. Births attended by skilled health staff (% of total) (BIR);
3. GDP per capita (US\$) (GDP);
4. Under-five mortality rate (per 1000 births) (U5M);
5. Access to sanitation (% with access) (SAN);
6. Access to water (% with access) (WAT); and,
7. Mean years of schooling (years) (SCH).

Country-level variables BIR, GDP, U5M, SAN, and WAT are obtained from the World Bank [218]; HEX is obtained from WHO [219]; and SCH through the UN Development Programme [220]. Wherever possible, 2015 records for each variable are used; however, whenever unavailable, the closest available record after and including 2010 is used.

Missing data

The missing data percentage of the socio-economic elements of the GLOBAL-67 dataset is slightly higher than that of the vaccination survey responses. Counting *do not know / refused / no answer* responses as missing data, the overall missing data percentage of the socio-economic data is 6.4% and this varies by country and socio-economic variable: SEX, AGE, EMP, and EDU have well below 1% missing data, whereas REL and INC are higher at 11.4% and 11.3% respectively. Of course, these missing data estimates are again likely to be tight upper bounds, as they include *do not know* responses.

Given that individual data have low missing data fractions, a complete-case analysis is opted for. Although complete-case analyses are often not preferred by statisticians as they provide unbiased estimates only when the missing data are MCAR [171], it can be reasonably expected that the missing data are MCAR as income and religion non-responses are likely independent of a respondent's income, religious affiliation, or of other variables present. (It is noted that adjacent analyses in Chapter 4, in which religion and income are discarded – and thus the missing data fraction is much lower – the associations of the other explanatory data with vaccination beliefs are in good agreement with the analyses in this chapter.)

Country-level data has 3% missing data. As full country-level data are required for the hierarchical model, a KNN procedure is opted for (as described in Appendix A.1.3). Following imputation, country-level data are z -scored to more easily compare data on different scales.

3.3 Hierarchical models

3.3.1 Hierarchical modelling of vaccination attitude data

Hierarchical (also known as *multilevel*) generalised linear models (HGLMs) extend regression methods when data falls into natural hierarchies. Multiple model assumptions are prone to be violated when regression methods are applied to data with hierarchical structure, thus invalidating model inferences. The most obvious model violation is that of independence, as the likelihood of many regression models (such as generalised linear models) invokes observational independence: $\mathcal{L}(\boldsymbol{\theta}|\mathbf{y}) = \prod_i P(y_i|\boldsymbol{\theta})$ for observations y_i and parameters $\boldsymbol{\theta}$. For hierarchical data, within-group observations are more likely to be similar to those in the same group than in others; without taking this dependency into account, the independence assumption is at risk, and this can lead to a cascade of failures for other model assumptions [221].

To highlight the issues caused when models do not account for hierarchical structures, consider applying a classical linear regression model to the hierarchical data (black circles) in Figure 3.1. In the absence of a hierarchical formalism, there are two typical ways of analysing these data: pooling data together with or without the higher-level labels (disaggregated analysis), or summarising within-cluster data at the top-level (aggregated analysis). In the disaggregated analysis in which the group label is kept (faint red lines, Figure 3.1 A), group-level variation cannot be modelled and the disaggregated data within each group is prone to being overfit by the model [222]. In the latter disaggregated case, in which a linear regression model is applied to completely disaggregated data (thick red line Figure 3.1 A), the assumption of independence of observations (and thus also independence of residuals, in this case) and residual normality are violated, and homoscedasticity of residuals would not be guaranteed either. Importantly, this disaggregated fit would lead to a false conclusion of negative relationship between response and explanatory data: this is an example of Simpson’s paradox [223] which a hierarchical treatment can protect against. In an aggregated analysis – which summarises group-level data with, for example, the mean (filled red points, Figure 3.1 B) – the correct global trend is observed (that is, that groups with lower mean explanatory variable have higher mean response), however variance at the within-cluster level has not been accounted for, and caution must be exercised to not commit an ecological fallacy in the interpretation of results (that is, to suggest the group-level trend applies at the individual level). Hierarchical modelling eliminates these issues (and can provide superior predictive




Figure 3.1: **Linear model violations and logical traps in non-hierarchical modelling approaches to hierarchical data** A disaggregated analysis (A) can lead to gross violations of linear model assumptions, in particular independence of observations. The linear fit (red thick line) can lead to flawed inferences suggesting a negative relationship when the true relationships within each cluster is positive (faint red lines). An aggregated analysis (B) will lead to a correct group-level trend, but a large loss of information and care needs to be taken not to suggest this trend at the individual-level.

performance when compared to classical regression methods [222]) by *partially pooling* information from both hierarchical levels. Hierarchical models are now introduced in the context of the GLOBAL-67 survey data introduced in Section 3.2.1. (Snijders and Bosker provide a discussion of model violations in non-hierarchical and hierarchical models in [221].)

Hierarchical logistic model for GLOBAL-67 data

The five-point Likert vaccination survey data are dichotomised into agree (tend to agree or strongly agree) versus non-agree (do not know, tend to disagree, or strongly disagree) responses. Vaccination attitude response data are therefore modelled as the random variable $Y_{ij} \in \{0, 1\}$ for individual i in country $j = 1, \dots, J$, where 1 is a response in agreement. The discrete individual-level socio-economic data are specified as ‘one-hot’ matrices: for a variable p with L_p levels the one-hot matrix $\mathbf{X}$ has $x_{lij} = 1$ if individual i in country j is in level $l = 2, \dots, L_p$ and $x_{lij} = 0$ otherwise. A baseline level for each variable ($l = 1$) is selected and removed from the one-hot matrix to prevent collinearity; instead, an intercept term (a vector of ones) is included to model the baseline category. Thus model coefficients are interpreted relative to the baseline group. The number of random-effect parameters is then given by $n_X = 1 + \sum_p (L_p - 1)$. Country-level data are specified through vectors $\mathbf{Z}_j$, where Z_{qj} is the value for country-level variable q for country j (there are n_Z country-level variables including an intercept term).

An individual's vaccination attitude response is related to their socio-economic data via the hierarchical model,

$$\begin{aligned} \text{logit}(P(Y_{ij} = 1|\boldsymbol{\beta})) &= \beta_{0j} + \boldsymbol{\beta}_{\text{SEX},j} \cdot \mathbf{x}_{\text{SEX},ij} + \boldsymbol{\beta}_{\text{AGE},j} \cdot \mathbf{x}_{\text{AGE},ij} + \boldsymbol{\beta}_{\text{EDU},j} \cdot \mathbf{x}_{\text{EDU},ij} \\ &+ \boldsymbol{\beta}_{\text{INC},j} \cdot \mathbf{x}_{\text{INC},ij} + \boldsymbol{\beta}_{\text{EMP},j} \cdot \mathbf{x}_{\text{EMP},ij} + \boldsymbol{\beta}_{\text{REL},j} \cdot \mathbf{x}_{\text{REL},ij}, \end{aligned} \quad (3.1)$$

where the intercept, β_{0j} , captures country-level variability in the baseline group, the parameters $\boldsymbol{\beta}_{[\cdot],j} \in \mathbb{R}^{L_p \times 1}$ model variability in gradients between countries, the vectors $\mathbf{x}_{ij}$ are indicator vectors of length $L - 1$ from the one-hot matrices $\mathbf{X}$, and where the logit transform is used to relate the dichotomous variable Y_{ij} to explanatory data.² The appropriate probability distribution over dichotomised response data is the Bernoulli distribution,

$$P(Y_{ij}|\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}) \sim \text{Bern}(P(Y_{ij} = 1|\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\Sigma})). \quad (3.2)$$

Variance in the intercept and gradient parameters is through a regression of random-effect parameters against country-level data,

$$\boldsymbol{\beta}_j | \mathbf{Z}_j, \boldsymbol{\gamma}, \boldsymbol{\Sigma} \sim \mathcal{N}(\boldsymbol{\gamma} \mathbf{Z}_j, \boldsymbol{\Sigma}), \quad (3.3)$$

where $\boldsymbol{\beta}_j = (\beta_{0j}, \boldsymbol{\beta}_{\text{SEX},j}, \boldsymbol{\beta}_{\text{AGE},j}, \boldsymbol{\beta}_{\text{EDU},j}, \boldsymbol{\beta}_{\text{INC},j}, \boldsymbol{\beta}_{\text{EMP},j}, \boldsymbol{\beta}_{\text{REL},j})^T$, $\boldsymbol{\gamma} \in \mathbb{R}^{n_X \times n_Z}$ is a matrix of fixed-effect parameters, and $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_{n_X}^2)$ is the covariance matrix. The fixed-effect parameters represent global ‘averages’ about which variance is captured through the random-effect terms.

The form of Equation (3.3) dictates how country-level variables are used to explain the random-effects. In the form presented, country level data are used to explain variance in both the intercept and the gradients: this is known as the *intercepts-and-slopes-as-outcomes* model [224] and allows cross-level interactions so that different country-level variables can have different effects in different countries.

A simpler form of Equation (3.3) uses country-level variables to explain unaccounted variance in the intercepts β_{0j} ,

$$\begin{aligned} \beta_{0j} &= \gamma_{00} + \gamma_{01} Z_{\text{HEX},j} + \gamma_{02} Z_{\text{BIR},j} + \gamma_{03} Z_{\text{GDP},j} + \gamma_{04} Z_{\text{U5M},j} \\ &+ \gamma_{05} Z_{\text{SAN},j} + \gamma_{06} Z_{\text{WAT},j} + \gamma_{07} Z_{\text{SCH},j} + \nu_{0j}, \\ \beta_{pj} &= \gamma_{p0} + \nu_{pj}, \end{aligned} \quad (3.4)$$

where $\nu_{pj} \sim \mathcal{N}(0, \sigma_p^2)$ for $p = 1, \dots, n_X$. This form is known as the *intercepts-as-outcomes*

²The logit function is defined as $\text{logit}(x) = \log\left(\frac{x}{1-x}\right)$.

model [224] as variability in the intercepts (the baseline groups) are explained using country-level data (but variability in the slopes is not).

3.3.2 Modelling assumptions

Hierarchical models involve a number of key assumptions which are catalogued here in the context of the hierarchical logistic regression model given by Equations (3.1) and (3.4).

Hierarchical regression addresses the issue of observation non-independence in data with a hierarchical structure. Independence of individual-level vaccine survey data in GALLUP-67 is guaranteed through the data collection process, which deploys representative random surveys within each country [225].

In the hierarchical treatment of data, classical regression model assumptions apply within each hierarchy. For linear regression, these model assumptions are: residual normality, residual homoscedasticity, and residual uncorrelatedness [226]. These assumptions, though needing verification for the (linear) random-effects model in Equation (3.3), do not apply within the logistic regression model in Equation (3.1). For the random-effects model, residual normality can be verified through a QQ-plot, and a plot of residuals against predicted value can verify the assumptions of homoscedasticity and uncorrelatedness of residuals [227].

For regressions at both hierarchies, a further two assumptions are required. Firstly, no predictor variable can perfectly predict the response variable (when this occurs *perfect separation* is said to occur) and, secondly, there can be no multicollinearity between explanatory variables. Perfect separation is easily verified through examination of the data, whilst multicollinearity can be detected using the condition number of the matrix of explanatory data. The condition number of a matrix A quantifies how sensitive the linear equation $\mathbf{y} = A\mathbf{x}$ is to small changes in $\mathbf{x}$ and is defined as,

$$\text{cond}(A) = \frac{\max_{\mathbf{x} \neq \mathbf{0}} (\|A\mathbf{x}\|/\|\mathbf{x}\|)}{\min_{\mathbf{x} \neq \mathbf{0}} (\|A\mathbf{x}\|/\|\mathbf{x}\|)}, \quad (3.5)$$

where $\|\cdot\|$ denotes the a vector norm (typically the Euclidean norm) [228]. The condition number thus quantifies the relative amount of stretching to shrinking of a non-zero vector by the matrix A . If A is ill-conditioned (for example through multicollinearity) then this will induce a large condition number. When the condition number is above 30 there is likely to be some strong multicollinearity within the data [229].

3.4 Results

Global and regional trends in attitudes to vaccinations are now explored by investigating individual response data across vaccination survey questions, country, and world region. This is followed by the implementation of the hierarchical logistic regression model introduced in Section 3.4.2 to explore the individual- and country-level determinants of attitudes. We implement the *intercept-as-outcomes* model to obtain a simple picture of the country-level factors which affect baseline variation in attitudes across countries. The logistic hierarchical regression model is implemented using the R package **lme4** [230]. Fixed-effect parameters in the hierarchical model are reported as odds ratios. Odds ratios are defined slightly differently depending on whether the explanatory variable is discrete or continuous. For discrete explanatory variables (such as the individual-level survey responses), the odds ratio is given by,

$$\begin{aligned}\text{OR}(\gamma_p) &= \frac{P(Y_{ij} = 1|x_{pij} = 1, X_{\bar{p}ij})/P(Y_{ij} = 0|x_{pij} = 1, X_{\bar{p}ij})}{P(Y_{ij} = 1|x_{pij} = 0, X_{\bar{p}ij})/P(Y_{ij} = 0|x_{pij} = 0, X_{\bar{p}ij})} \\ &= \exp(\gamma_p),\end{aligned}$$

where $X_{\bar{p}ij}$ are explanatory data for every variable except p , and where $P(\cdot)$ is the hierarchical model defined by Equation (3.3). For discrete data (individual-level data), the odds ratio is interpreted as the increase in odds of positive vaccination view for variable p relative to the odds of the baseline group. For continuous variables (such as the country-level predictors), the odds ratio is still given by $\text{OR}(\gamma_p) = \exp(\gamma_p)$,

$$\begin{aligned}\text{OR}(\gamma_p) &= \frac{P(Y_{ij} = 1|x_{pij} = x + 1, X_{\bar{p}ij})/P(Y_{ij} = 0|x_{pij} = x + 1, X_{\bar{p}ij})}{P(Y_{ij} = 1|x_{pij} = x, X_{\bar{p}ij})/P(Y_{ij} = 0|x_{pij} = x, X_{\bar{p}ij})} \\ &= \exp(\gamma_p).\end{aligned}$$

However, the interpretation for a continuous variable is slightly different: an odds ratio greater (less) than one means the probability of the outcome variable increases (decreases) with a unit increase in explanatory variable (holding all other parameters constant). The model assumptions outlined in Section 3.3.2 are verified for these models in Appendix B.1.

3.4.1 Global and regional trends in attitudes

An overall summary of vaccination survey response data stratified by world region is shown in Figure 3.2. As in Chapter 2, these regions are as defined by WHO.³ Within each vaccine

³For reference, these are the African (AFR), Americas (AMR), European (EUR), Eastern Mediterranean (EMR), South East Asian (SEAR), and the Western Pacific (WPR) regions.

survey question variability in vaccine sentiment across countries and WHO regions is investigated by considering the fraction of respondents who either agree or disagree with the four statements on vaccinations. Countries are ranked by the percentage of negative responses (responding either **strongly disagree** or **tend to disagree**): these results are shown in Figure 3.2 (red denotes a negative response; blue positive; and the vertical grey bar shows the region’s mean-averaged negative sentiment for a relative comparison).

Overall differences in responses between the six WHO regions are shown in Figure 3.3 A and vaccine safety scepticism is mapped in Figure 3.3 B. (In this exploratory analysis, **do not know / no response** responses are omitted to more clearly show the level of vaccine trust and scepticism. This requires a renormalisation after removal of this category. Survey weights provided by WINGIA are utilised to appropriately account for over- or under-represented individuals within a population.)

Regional trends

EUR reports the highest mean-averaged negative responses for vaccine importance, safety, and effectiveness (8.0%, 17.0%, and 11.3%, respectively) and seven of the ten most negatively-reporting countries to vaccine safety are in EUR (France, Bosnia & Herzegovina, Russia, Ukraine, Greece, Armenia, and Slovenia) (Figures 3.2 and 3.3 A). t -tests⁴ between regional means reveal that EUR has significantly more importance- and safety-related scepticism than AFR, AMR, EMR, and SEAR, and that WPR has more safety-related scepticism than SEAR. SEAR and WPR report the most religious incompatibility (25.7% and 24.3% respectively) each of which is significantly more than the other four regions. (All results stated above for pairwise t -tests are significant at $p < 0.05$.)

Worldwide vaccine-scepticism is mapped in Figure 3.3 C. With exceptions of France and Italy – which have notably high levels of safety-scepticism – southern and eastern European countries express more safety concerns than northern and western European countries. The USA, Mexico, and Canada have higher levels of safety concerns than south American countries which have relatively strong vaccine safety confidence. WPR has high levels of vaccine safety concerns, notably in Japan, Mongolia, Vietnam, and Hong Kong.

Although vaccine safety and importance are concerns in EUR, there is relatively little reported religious incompatibility with vaccinations: in WPR and SEAR, however, substantial concerns are expressed about religious compatibility in a number of countries including China in WPR and India in SEAR: two countries with approximately a third of the world’s

⁴A t -test in this context examines whether the difference between two means is statistically different from zero. The test statistic is given by $T = \frac{\bar{\mu}_1 - \bar{\mu}_2}{\bar{s}}$, where $\bar{\mu}_{1/2}$ are estimates of the mean distribution and are assumed to be normally distributed, and $\bar{s}$ is an estimate of the standard error of the sampling distribution and is assumed to be chi-square distributed. Then T is t -distributed with $n_1 + n_2 - 2$ degrees of freedom under the null hypothesis [190], where $n_{1/2}$ are the number of data points in estimating $\bar{\mu}_{1/2}$.




Figure 3.2: Ranking by overall level of vaccine scepticism and stratifying by world region reveals high levels of vaccine safety concerns in many European countries and religious incompatibility concerns in the South East Asian and Western Pacific regions. Weighted responses to survey questions by country and WHO region across the four survey statements ranked by overall negative sentiment (strongly disagree or tend to disagree). EUR performs poorly for vaccine importance, safety, and effectiveness-related concerns. SEAR and WPR have the highest levels of religious incompatibility-based scepticism. Black vertical lines show the mean-averaged negative response rate in each region.




Figure 3.3: **Variation in attitudes to vaccine safety shows spatial contiguity and differences between importance and safety responses reveals higher belief in the importance of vaccinations than in their safety** (A) Summary of Likert responses by world region. (B) Pearson correlations between percentage of respondents across all countries agreeing (strongly agree or tend to agree) with each statement. (Pearson correlations are given by Equation (2.12) but for unranked data.) (C) World map of percentage negative (tend to disagree or strongly agree) survey responses to the statement overall I think vaccines are safe. (D) Differences in the proportion of survey respondents within each country reporting that vaccines are important and that vaccines are safe (with 95% confidence intervals).

population.

Country-specific trends

The lowest levels of *vaccine importance*-related scepticism were reported in Bangladesh (0.3%), Ecuador (0.4%), and Iran (0.5%) whilst respondents from Italy (15.4%), Azerbaijan (15.7%), and Russia (17.1%) report the highest (global average: 5.8%). The lowest levels of *vaccine safety* scepticism were reported in Bangladesh (0.2%), Saudi Arabia (1.2%), and Argentina (1.3%), whilst respondents from Japan (31.0%), Bosnia & Herzegovina (38.3%), and France (45.2%) have the highest (global average: 13.0%). With regards to *vaccine effectiveness*-scepticism, the least sceptic countries were Argentina (1.3%), Ethiopia (2.1%), and Ecuador (2.2%), whilst Italy (18.7%), Russia (20.1%), and Bosnia & Herzegovina (27.3%) had the highest levels of scepticism. *Religious compatibility* issues were of least concern in Saudi Arabia (2.3%), Finland (2.6%), and Brazil (3.2%) and of most concern in Vietnam (31.8%), Thailand (44.4%), and Mongolia (50.5%).

Trends between vaccine responses

Taking Pearson⁵ correlations between the fractions of negative responses across countries for each question reveals strong correlations among vaccine importance, safety, and effectiveness than with religious compatibility (Figure 3.3 B; also see latent variable treatment of these data in Chapter 4).

A striking observation is the difference in the number of respondents reporting positively to vaccine importance relative to vaccine safety; a trend which is significant for almost all of the countries surveyed (Figure 3.3 D). The reported difference between vaccine importance and safety suggests that intent to vaccinate could be buffered by the perceived importance of vaccinations, implying that people are willing to take a risk against an apparent danger to guard against disease. As suggested by the mutual correlation in responses to vaccine importance and safety, individual responses between vaccine effectiveness and vaccine safety are very similar, though we note a higher fraction of countries with no significant difference and some countries with a higher fraction of respondents reporting that vaccines are safe then effective (Nigeria, Ghana, Pakistan, Indonesia, and Bangladesh). See Appendix B.2 for country-wide variability in the differences between individuals' responses to vaccine safety and effectiveness.

⁵Pearson correlations are Spearman correlations on unranked data – see Equation (2.12).

3.4.2 Determinants of attitudes to vaccinations

The results of the hierarchical regression are shown in Table 3.1. Although significant variables are highlighted for p -values at the 0.10 level, only p -values lower than 0.05 are considered in this analysis. Since multiple models are evaluated, consideration of corrections for multiple comparisons is required [190]; however, due to interdependency of many explanatory variables, choice of correction is not straightforward. As this analysis is somewhat exploratory in nature – since factors plausibly linked to vaccine attitudes are sought – rejecting the null hypothesis of no association when there is no link is slightly less of a concern than vice versa (that is type II errors⁶ are of a greater concern). Uncorrected p -values are therefore provided with 95% confidence intervals. Credibility to the significance of factors associated with attitudes is gained through the consistency of their link across the four models.

Individual-level trends

Sex: Males are less likely to report that vaccines are important than females (odds ratio [OR] 0.86 [0.80, 0.94] [95% confidence interval]), but there is no significant difference found between how the sexes report to vaccine safety, effectiveness, or compatibility with their religious views. Age: The oldest age group (65+) is more likely to report that vaccines are effective and religion-compatible than 18-24 year olds (OR 1.39 [1.11, 1.76] and OR 1.27 [1.05, 1.53] respectively). Those aged 25-34 are less likely to believe vaccines are safe compared to 18-24 year olds (OR 0.88 [0.77, 1.00]). There are no significant differences in age groups relative to the baseline group for vaccine importance. Education: Any amount of education appears to elevate vaccination views to recognise the importance and effectiveness of vaccinations but, importantly, not the safety of vaccinations. Those with high levels of education (undergraduate, masters, or PhD degree) have fewer religious compatibility issues than those with low levels of education. Income: Income has one of the strongest effects on vaccine responses: all income groups above the first quintile (low income) hold stronger pro-vaccination beliefs for the importance, safety, and effectiveness of vaccines; and those in the third and fourth quintiles hold stronger views for religious compatibility. Work type: There are few associations with work type suggesting that work status has little impact on vaccination beliefs. The only significant association found is that those unemployed are more likely to report vaccine sceptic views for the safety and effectiveness of vaccinations (OR 0.77 [0.67,0.90] and OR 0.79 [0.67, 0.92], respectively) than the baseline group. Religion: With the exception of Hindu and Jewish respondents, every religious group (and atheists/agnostics), holds less positive vaccination importance, effectiveness, and safety views than Roman Catholics. With respect to religious incompatibilities of vaccines, other Christians and atheists/agnostics are

⁶A type II error is falsely accepting the null hypothesis [190]

	importance	safety	effectiveness	rel. compatibility
individual-level				
sex				
female (baseline)				
male	0.86 (0.80-0.94)***	1.01 (0.94-1.09)	0.96 (0.89-1.03)	0.96 (0.90-1.02)
age				
18-24 (baseline)				
25-34	0.89 (0.77-1.04)	0.88 (0.77-1.00)**	0.97 (0.86-1.10)	0.94 (0.85-1.05)
35-44	0.99 (0.84-1.17)	0.95 (0.83-1.07)	1.02 (0.89-1.16)	1.06 (0.94-1.20)
45-54	0.99 (0.83-1.19)	0.97 (0.85-1.12)	1.16 (1.00-1.34)*	1.05 (0.93-1.20)
55-64	0.99 (0.80-1.23)	0.98 (0.84-1.13)	1.07 (0.90-1.28)	1.14 (0.98-1.32)*
65+	0.94 (0.69-1.27)	1.07 (0.88-1.31)	1.39 (1.11-1.76)***	1.27 (1.05-1.53)**
education				
no education (baseline)				
primary	1.46 (1.08-1.97)**	0.90 (0.71-1.15)	1.20 (0.99-1.46)*	1.12 (0.94-1.35)
secondary	1.39 (1.05-1.83)**	0.95 (0.74-1.22)	1.31 (1.08-1.57)***	1.19 (0.97-1.44)*
university	1.44 (1.09-1.92)**	0.98 (0.76-1.26)	1.39 (1.15-1.68)***	1.35 (1.08-1.68)***
masters/PhD	1.33 (0.94-1.89)	0.98 (0.73-1.33)	1.33 (1.05-1.70)**	1.43 (1.12-1.83)***
income				
first quintile (low) (baseline)				
second quintile	1.19 (1.04-1.38)**	1.17 (1.04-1.31)***	1.21 (1.07-1.36)***	1.08 (0.98-1.20)
third quintile	1.36 (1.16-1.58)***	1.25 (1.10-1.43)***	1.31 (1.13-1.50)***	1.16 (1.04-1.29)***
fourth quintile	1.53 (1.26-1.85)***	1.32 (1.15-1.53)***	1.43 (1.22-1.68)***	1.19 (1.05-1.36)***
fifth quintile (high)	1.43 (1.21-1.70)***	1.28 (1.10-1.48)***	1.33 (1.14-1.55)***	1.12 (0.96-1.30)
work type				
housewife (baseline)				
retired/disabled	1.13 (0.88-1.45)	0.96 (0.81-1.14)	0.91 (0.77-1.08)	1.00 (0.87-1.15)
student	0.99 (0.77-1.27)	0.96 (0.80-1.15)	1.04 (0.85-1.26)	1.00 (0.86-1.17)
unemployed	0.95 (0.77-1.17)	0.77 (0.67-0.90)***	0.79 (0.67-0.92)***	0.90 (0.78-1.05)
part-time	0.86 (0.70-1.07)	0.89 (0.75-1.06)	0.91 (0.77-1.07)	0.91 (0.78-1.06)
full-time	1.03 (0.86-1.25)	1.01 (0.88-1.16)	1.00 (0.88-1.15)	0.98 (0.87-1.11)
religion				
Roman Catholic (baseline)				
Protestant	0.57 (0.43-0.77)***	0.72 (0.58-0.90)***	0.83 (0.67-1.02)*	1.04 (0.85-1.26)
Russian/Eastern-Orth	0.41 (0.30-0.57)***	0.56 (0.46-0.70)***	0.64 (0.51-0.82)***	0.87 (0.71-1.07)
other Christian	0.52 (0.41-0.67)***	0.74 (0.61-0.90)***	0.72 (0.58-0.89)***	0.81 (0.68-0.97)**
Muslim	0.63 (0.43-0.91)**	0.72 (0.55-0.96)**	0.67 (0.51-0.89)***	0.82 (0.62-1.08)
Jewish	1.48 (0.50-4.39)	0.64 (0.34-1.20)	2.43 (0.74-7.98)	0.97 (0.45-2.06)
Buddhist	0.30 (0.17-0.51)***	0.48 (0.32-0.71)***	0.55 (0.35-0.87)**	0.83 (0.55-1.26)
Hindu	0.55 (0.28-1.08)*	1.40 (0.66-2.97)	0.69 (0.41-1.16)	1.07 (0.61-1.88)
other	0.29 (0.21-0.38)***	0.45 (0.35-0.58)***	0.44 (0.34-0.57)***	0.47 (0.38-0.59)***
atheist/agnostic	0.49 (0.38-0.64)***	0.73 (0.61-0.86)***	0.75 (0.63-0.91)***	0.45 (0.38-0.54)***
ine				
country-level				
HEX (% of GDP)	0.89 (0.74-1.08)	0.71 (0.60-0.85)***	0.81 (0.70-0.93)***	1.00 (0.87-1.14)
BIR (% of total)	0.66 (0.49-0.88)***	0.55 (0.44-0.68)***	0.75 (0.62-0.91)***	0.79 (0.65-0.96)**
GDP (US\$ per capita)	1.56 (1.25-1.95)***	1.67 (1.38-2.02)***	1.68 (1.43-1.97)***	0.99 (0.86-1.15)
U5M (per 1000 births)	0.76 (0.55-1.06)	0.93 (0.71-1.23)	0.94 (0.75-1.18)	1.16 (0.93-1.45)
SAN (% with access)	0.82 (0.58-1.17)	1.35 (1.00-1.82)**	1.21 (0.96-1.54)	1.29 (1.01-1.65)**
WAT (% with access)	1.31 (0.99-1.75)*	1.23 (0.98-1.54)*	0.94 (0.78-1.14)	1.05 (0.88-1.26)
SCH (years)	0.56 (0.43-0.72)***	0.75 (0.60-0.93)***	0.63 (0.52-0.75)***	0.80 (0.66-0.97)**

*p<0.1; **p<0.05; ***p<0.01

Table 3.1: **Odds ratios and 95% confidence intervals for the fixed-effect parameters in the *intercepts-as-outcomes* hierarchical logistic regression model.** Responses to vaccination survey data are dichotomised so that odds ratios (OR) greater (less) than one represent a positive (negative) association between explanatory data and each vaccination attitude question. Individual's income level is strongly associated with all four statements, with second to fifth income quintiles holding more positive vaccine views than the first quintile. Any level of education appears to elevate pro-importance, effectiveness, and religious compatibility views but notably education levels have no bearing on whether an individual replies more positively to vaccine safety.

more likely to have issues (though it is unclear – due to the phrasing of this question assuming the respondent holds religious beliefs – whether sensible inferences can be made about agnostics/atheists).

Country-level trends

Country-level variables are used to explain variation in the baseline group parameter β_{0j} . Countries with higher levels of births attended by skilled health staff (BIR) and mean years of schooling (SCH) have baseline groups with more positive responses than those with lower levels of health assistance at birth or lower levels of schooling (OR 0.66 [0.49, 0.88] and OR 0.56 [0.43, 0.72] respectively). Countries with higher levels of GDP per capita have higher values of β_{0j} than countries with lower levels of GDP per capita (OR 1.56 [1.25, 1.95]).

Similar results hold for vaccine safety and effectiveness. BIR: OR 0.71 [0.60, 0.85] and OR 0.81 [0.70, 0.93]; SCH: OR 0.75 [0.60, 0.93] and OR 0.63 [0.52, 0.75]; GDP: OR 1.67 [1.38, 2.02] and OR 1.68 [1.43, 1.97] (all respectively). In addition, higher levels of health expenditure are also associated with more negative baseline beliefs for importance and effectiveness (OR 0.71 [0.60, 0.85] and OR 0.81 [0.70, 0.93], respectively).

With regard to religious compatibility, countries with higher levels of BIR and SCH have individuals with more sceptic vaccination beliefs (0.79 [0.65, 0.96] and 0.80 [0.66, 0.97], respectively), whilst countries with a higher percentage of access to improved sanitation facilities are more likely to have fewer religious compatibility issues (1.29 [1.01, 1.65]).

3.5 Discussion

In this chapter worldwide attitudes to vaccinations are explored by probing individuals' views on the importance of childhood vaccines, the safety and effectiveness of vaccines, and the compatibility of vaccines with religious beliefs. Attitudes on these four measurement items are mapped and regional variability is explored. A large majority of respondents report that vaccines are important, safe, effective, and compatible with religious beliefs; however, there are notable concerns too. Interestingly, individuals in almost every country are more likely to report that vaccines are important than they are safe. This perception of relative importance to safety may mitigate losses in vaccination uptake. If this difference lessens over time, it could suggest an increase in vaccine *complacency* (as defined by SAGE's '3Cs' definition in Section 1.2.1) and could lead to further vaccine refusals.

Sentiment towards the safety of vaccines is most negative in the European and Western Pacific regions, wherein nine of the ten most vaccine-safety-sceptic countries are situated (France (67), Bosnia & Herzegovina (66), Russia (64), Ukraine (62), Greece (61), Armenia (60), Slovenia (58), Japan (65), and Mongolia (63)). (Numbers here, and below, represent a

rank out of 67 countries.) In contrast to EUR and WPR, individuals surveyed in the South East Asia region, including Bangladesh (1), Indonesia (6), Thailand (12), largely perceive vaccinations to be safe. Respondents in India (29), Pakistan (19), and Nigeria (18) also have higher confidence in safety than the global average. The contrast between these countries is particularly noteworthy given the high degree of scepticism in regions with advanced healthcare systems where vaccinations are readily available (notably most European countries, North America, and Japan) and higher safety-confidence in areas which lack universal access to healthcare facilities (Bangladesh, Nigeria, Pakistan, and India). Sentiment towards the effectiveness of vaccines is again lowest in the European and Western Pacific regions. This suggests that vaccine *confidence* hesitancy could be lowest in EUR and WPR.

Many countries are found to have strong faith-based vaccination issues. High proportions of individuals with religious objections are found in WPR and SEAR: with Mongolia (67), Vietnam (65), China (63), Thailand (66), Bangladesh (60), and India (52) all expressing high levels of religious incompatibility. In countries where polio rejection has been related to religious leaders, religious incompatibility is moderate: Nigeria (19), Pakistan (17), Afghanistan (18). Of the individuals who reported to be Muslim in these countries, 12%, 14%, and 3% (respectively), had religious objections themselves, highlighting the complex contextual nature of hesitancy in these countries.

Overall global trends in socio-economic correlates of vaccine attitudes reveal that countries with higher levels of births attended by skilled health staff (BIR) and mean years of schooling (SCH) have lower baseline levels of pro-vaccine sentiment across all four survey questions. Countries with higher levels of gross domestic product (GDP) tend to have higher baseline confidence. At the individual level the global trend is for those with some level of education to hold more positive views on vaccine importance, effectiveness, and religious compatibility, though not on safety. Globally, individuals with high levels of income are associated with more positive beliefs across all four survey questions.

In Section 1.3.1, safety concerns were highlighted for a number of countries. According our findings, Russia is the 4th most safety-sceptic country in the world; yet, vaccine coverage of childhood vaccinations BCG, MCV, DTP3, and POL3 in Russia are high (in excess of 96% for the past decade [27]). Similarly, Japan – with a history of vaccine scepticism [110] – is the 3rd most safety-sceptic country in the world. Interestingly, despite Russia and Japan’s low safety rankings, both countries have a higher fraction of individuals reporting more positively to the importance of vaccines. This strongly suggests that the perceived importance of vaccines is acting to maintain high coverage rates. Despite historic concerns with the MMR vaccine, the UK is the 29th most safety-confident country from the countries surveys and Pakistan and Nigeria (with strong polio concerns) both rank higher than the UK.

We also find that France is the most sceptic country towards the safety of vaccines. This high percentage of negative sentiment around safety-based vaccine scepticism reported in France is likely due to multiple recent vaccine controversies in the past decade [231, 232]. Italy also has low confidence in the safety of vaccines: the MMR vaccine has been a concern in Italy and sentiment has possibly been weakened by a court’s ruling in 2014 that the MMR vaccine *had* caused autism in their child (although this was overturned in 2015 by a higher court) [233, 234].

The results presented in this chapter highlight the importance of routinely-collected national and sub-national vaccine attitude surveys. For example, although the origins of the MMR scandal and polio boycotts are quite different (as are the healthcare provisions within each country), underlying both is a population with latent distrust and a susceptibility to misinformation. Currently, Nigeria, India, and Pakistan have smaller amounts of vaccine-safety confidence than the USA, UK, and Japan, despite undergoing similarly high-profile vaccine-scepticism-inducing events in the past two decades [20, 22, 57, 110, 127, 235]. This suggests that relatively short time-scales are involved with opinion dynamics and emphasises the need to monitor attitude levels to understand the dynamics, time-scales, and determinants at play. Higher frequency hesitancy survey data can enable us to investigate the causal link between an individual’s socio-economic attributes, their social-network ties, and their ability to be influenced by media and their social network. This would allow identification of sceptic clusters (in spatial and socio-economic dimensions) and an understanding of how vaccine sentiment spreads. Of particular importance in this regard is understanding when and where clusters of non-vaccinated individuals form, as clusters of non-vaccinated individuals can act to drastically lower herd immunity levels, even if coverage rates averaged over whole population are above levels thought sufficient for herd immunity [236]. Such clustering of non-vaccinated individuals has been recently established in a number of East African countries [237]; and has been recognised by the US Centers for Disease Control as an important phenomenon to address [238]. As access to vaccinations becomes improved across many of these areas, focus will shift into how clusters of non-vaccinated individuals may form through negative vaccine-sentiment spreading on social networks. Understanding this link between attitudes and coverage uptake is crucial if we are to predict future coverage levels and understand how, in the case of Russia (for example), relatively high negative sentiment can still provide robust immunisation coverage.

There are a number of limitations in this study. Survey questions do not provide information on an individual’s view of specific vaccines or an individual’s vaccination history or intent to vaccinate. Of course, a further limitation is this inability to track temporal changes to attitude: in Chapter 5, we suggest a method to infer hesitancy through an approach which

circumvents the need for attitude survey data altogether, thus providing a platform for the analyses proposed above.

In the next chapter, the **GLOBAL-67** data are augmented with more recently-collected data from the Sahel region for a wider global analysis. In this analysis, the focus is switched from these four dimensions of vaccine attitude into obtaining a single quantitative measure of vaccine hesitancy.

4 | Deriving a quantitative index of vaccine hesitancy: mapping regional trends and establishing determinants both worldwide and in the Sahel

Reluctant bunch beginning to
hesitate (4)

— Arachne, *Guardian*
crossword No. 27,150

4.1 Introduction

According to SAGE’s ‘3Cs’ model, vaccine hesitancy is a combination of vaccine *confidence*, *complacency*, and *convenience* ([58] and see Section 1.2.1). In order to understand the mathematical characteristics and properties of vaccine hesitancy, we require a suitable quantitative measure.

In this chapter, we propose modelling vaccine hesitancy as a latent variable derived from the four dimensions of attitudes to vaccines introduced in Section 3.2.1: vaccine importance, safety, effectiveness, and religious compatibility. This approach is reconcilable with the ‘3Cs’ definition of vaccine hesitancy as the safety and effectiveness of vaccines contribute to vaccine *confidence*, whilst the perceived importance of vaccines may be closely related to *complacency*, as those who perceive vaccines as less important might be more complacent in positive vaccine health-seeking behaviours. Religious compatibility is harder to place; however, it likely spans a combination of confidence, convenience, or complacency [58].

A latent variable inference method called *item response theory* [239] is used to derive the summary measure of vaccine hesitancy. For each person, this mathematical tool projects

the four dimensions of vaccine hesitancy onto a single latent value which we deem to be vaccine hesitancy. This projection implicitly weights the relative contributions from each vaccine attitude dimension allowing us to investigate how these four questions on attitudes to vaccines contribute to vaccine hesitancy. This method can be readily extended to include additional survey questions that might provide a more rigorous understanding of hesitancy.

We term this quantitative measure of individual vaccine hesitancy the **vaccine hesitancy index**, which we believe to be the first such measure of vaccine hesitancy. This index permits a regional and national assessment of vaccine hesitancy (through summary statistics of individuals' indices); allows variation in hesitancy to be explained through socio-economic factors; and – when future survey data becomes available – will also permit a mathematical treatment of the spatio-temporal dynamics of hesitancy, including hesitancy forecasts and the modulating role of the media.

We perform this analysis on an augmented dataset: we combine the **GLOBAL-67** dataset introduced in Section 3.2.1 with vaccine attitude data collected a year later in 2016. This newer survey probes the same dimensions of vaccine attitude in the Sahel region in Africa (we label this dataset **SAHEL-6**). A major focus of this chapter is therefore the Sahel region, and – due to recent concerns over vaccinations here – Nigeria [44, 47, 240]. An enhanced set of socio-economic data available in Nigeria allows an investigation of the effect of religious upbringing, ethnicity, and wealth on hesitancy.

We begin this chapter with an overview of the **SAHEL-6** dataset and a description of how it is combined with **GLOBAL-67** (Section 4.2). An introduction to item response theory and its hierarchical generalisations are then presented in Section 4.3. Parameter estimation and model fit assessment of item response models are discussed in Section 4.4.4. In Section 4.5 findings are presented for hesitancy inference and the socio-economic correlates of hesitancy. A full interpretation and discussion of findings is given in Section 4.6. In this final section, we outline further research directions and, in particular, we propose a novel method for hesitancy inference in the *absence of vaccination attitude survey data*. This proposed analysis method could have profound consequences, as it would allow the historic inference of vaccine hesitancy using publicly-available non-survey datasets, and would thus allow historic dynamics to be quantified and forecasted: this is fully discussed in Chapter 5.

4.2 Data

4.2.1 SAHEL-6

In 2016, WINGIA, in association with the LSHTM, extended the **GLOBAL-67** survey by interviewing survey respondents in the Sahel region in west Africa. (The Sahel comprises

Burkina Faso, Mali, Mauritania, Niger, Nigeria, and Senegal and we denote this Sahel dataset **SAHEL-6**.) Although the four Likert dimensions measured on attitudes to vaccines remain the same, there are three differences between the two surveys: 1) income and religion data are not collected in **SAHEL-6**; 2) age groups are coded differently (18-24, 25-34, 35-44, 45-54, 55-64, 65+ in **GLOBAL-67** versus 18-19, 20-29, 30-39, 40-49, 50-59, 60-64, 65+ in **SAHEL-6**); and 3) the survey wording for the vaccination responses **tend to agree** and **tend to disagree** in **GLOBAL-67** is changed to **somewhat agree** and **somewhat disagree** in **SAHEL-6**.

As we wish to ‘match’ these datasets to perform analyses on combined data, these issues need to be resolved. This first issue is resolved by simply removing income and religion from the **GLOBAL-67** dataset; the second by reclassifying ages into three groups, 18-45, 45-64, and 65+ (with 40-49 year olds in **SAHEL-6** grouped into 18-45; this classification acts to minimise the classification error and decrease the number of groups to seek broad association between age and vaccine attitude); and the third by assuming that respondents who replied with **somewhat dis/agree** would have replied **tend to dis/agree** in the original survey. (We also note that we rename primary, secondary, and post-graduate education in Section 3.2.1 to low, medium, and high education, respectively.)

As data for Nigeria is collected in both surveys, it is assumed that any temporal differences between the data are small in comparison to the overall patterns in the full data and that global trends obtained will not be adversely impacted.

A total of 13,266 respondents are surveyed across the Sahel countries (Burkina Faso: 1,907; Mali: 1,856; Mauritania: 1,130; Niger: 1,957; Nigeria: 4,032; Senegal: 2,348) using the same procedure as outlined in Section 3.2.1. There are no missing data entries for sex and age and the missing data fractions for education and employment are very small (0.09% and 0.2%, respectively). **Do not know / no response** replies to vaccination survey questions are also small at 0.8%, 2.3%, 1.6%, and 3.3% for the importance, safety, effectiveness, and religious compatibility of vaccines, respectively.

4.2.2 SAHEL-NIGERIA

A focus of this chapter is a national and regional analysis of Nigeria. In addition to the variables previously described, the **SAHEL-6** survey data contains further variables likely linked to hesitancy that are not featured in the **GLOBAL-67** data. In Nigeria – in which 4,032 respondents were surveyed across 37 states – the following variables are available (variable recodes are given in parentheses and the baseline groups are underlined):

1. **Ethnic group**: Fulani/Peulh/Haalpulaar, Ibibio (other), Idoma (other), Igbo, Hausa, Kanuri/Beri Beri (other), Yoruba, other (ETH)

2. Religious schooling: Yes, no (REL)
3. Household economic situation: Unable to meet basic needs without charity (low), able to meet basic needs (med), able to meet basic needs with some non-essential goods (med), able to purchase most non-essential goods (high), plenty of disposable money (high) (EC)
4. Change in economic situation: Worsened, stayed the same, improved (DEC)
5. Perceived local economic situation: Much lower (lower), a little lower (lower), about average (the same), a little higher (higher), much higher (higher) (PEC).

Some factors have been recoded: for example, ethnicities with low counts are recoded to **other**. There are approximately 100 individuals sampled within each Nigerian state (a maximum count of 130 in Benue, a minimum of 100 in Ebonyi, and a median of 106).

4.3 Item response theory

The *raison d'être* of latent variable models is to relate observed (or *manifest*) data to unobserved, *hidden*, or *latent* variables. Latent variable models allow for data correlation and dependency structures to be better understood in terms of unobserved variables. Latent variable models are prevalent in the scientific literature, with methods such as principal components analysis and hidden Markov models commonly deployed to explore and explain unobserved patterns and dynamics in data [241–243].

Initially developed for the purposes of characterising the statistical properties of examinations and examinees [244], item response theory (IRT) has since developed into a more general framework to relate individuals' multivariate data to descriptive latent attributes. Item response theory infers an individual's (or, more generally, an observation's) position along a continuous latent scale: in the context of examination testing, for example, this would place an individual along an ability (or intelligence) scale based on their responses to a number of *items* (examination questions) [245].

F.M. Lord set out the main practical benefits of IRT in 1980 in the context of educational testing [245]. We restate these benefits in a completely generalised setting as: *a framework for estimating item-specific parameters to describe the variation in probability of individual response data [for example examination question responses] as a function of an individual's latent ability or latent hesitancy*. IRT therefore provides a means to perform simultaneous inference on the characteristics of each item (for example how difficult examination questions are) and how they relate to the latent dimension (for example student ability). In the context of examination testing, IRT allows standardised examinations to be designed, students'

abilities to be inferred, and new examination questions to be designed based on their characteristics [245, 246]. IRT offers advantages over ability inference through a sum of correct responses to examination questions (for example), as students may be awarded unduly high scores through guessing or through other sources of noise. IRT, however addresses these effects by awarding more weight to harder questions that smaller proportions of the sample answer correctly (indeed, some IRT models include a guessing parameter to further account for this variability [246, 247]).

We propose utilising item response theory in the context of vaccination survey response data to infer an individual’s *vaccine hesitancy* (analogous to student’s ability in an examination context) and to examine how the probability agreement with the vaccination survey questions varies with vaccine hesitancy. These goals are best illustrated through examples. (Unless otherwise explicitly stated, we henceforth switch notation and terminology to talk about IRT models in the context of an individual’s responses to vaccination survey data (items) and vaccine hesitancy (the latent variable).)

4.3.1 One- and two-parameter IRT models

The central tenet of IRT is that responses to items can be linked to a latent variable through a probabilistic link function. The simplest of the item response models is arguably the one-parameter Rasch model [248] which neatly illustrates many of the features of item response theory outlined above. Consider M vaccination survey questions to which individual i provides a response $Y_{im} \in \{0, 1\}$, where 0 corresponds to a negative response and 1 a positive response. (In our case $M = 4$ for questions on vaccine importance, safety, effectiveness, and religious compatibility.) The probability of a positive response to vaccination question m is related to vaccine hesitancy through an *item response curve* (IRC),

$$P(Y_{im} = 1 | \eta_i, b_m) = \frac{1}{1 + \exp(-(\eta_i - b_m))},$$

where η_i is an individual’s vaccine hesitancy (the latent variable), and b_m are item-specific parameters called *difficulty* parameters. Difficulty parameters are so-called because they represent a horizontal shift in IRCs and thus, in an educational setting, tell us how difficult an examination item is. In the vaccine hesitancy setting, they provide information on which survey questions illicit most sceptic responses. Item response curves for the one-parameter model are shown in Figure 4.1 (left-hand plot).

The one-parameter Rasch model is limited to basic modelling scenarios; in particular, it is unable to model items with varying rates of change with the latent variable. A natural

generalisation of the one-parameter Rasch model is the two-parameter logistic model,

$$P(Y_{im} = 1|\eta_i, a_m, b_m) = \frac{1}{1 + \exp(-a_m(\eta_i - b_m))}, \quad (4.1)$$

where the parameter a_m has been included, allowing IRCs with increased modelling complexity (as gradients can now differ). The a_m parameters are known as *discrimination* parameters, as IRCs (in an examination context) with steeper slopes are better able to discriminate between low and high ability levels. In the design of examinations or surveys, items that are unable to discriminate between individuals (that is, $a_m \approx 0$) may be removed, as they provide no information about a student's ability. As a side note, we suggest that – if this methodology is adapted for a quantitative index for vaccine hesitancy by others – survey questions for which $a_m \approx 0$ should be dropped from future surveys. Such questions are not informative of vaccine hesitancy and deployment of other questions may yield more profitable information on the determinants of hesitancy: this becomes important when, for example, budgetary requirements place a restriction on the number survey questions.

IRC for various difficulty and discrimination parameters are shown in Figure 4.1. A family of curves with fixed discrimination parameters and varying difficulty parameters is shown on the left ($b_m = -1, 0$, and 1), and a family of curves with fixed difficulty and varying discrimination is shown on the right ($a_m = 0.2, 1$, and 2). The higher the difficulty parameter the lower the probability of a positive vaccination survey response at all hesitancy levels. The




Figure 4.1: **Item response curves illustrating a range of difficulty and discrimination parameters** (Left) IRCs with the same discrimination parameter but different difficulty parameters (which is an example of the one-parameter IRT model when $a_m = 1$ in Equation (4.1)), and (right) IRCs from the two-parameter logistic IRT model with the same difficulty parameter but different item discriminations.

lower the discrimination value (that is, the shallower the gradient), the less able the item is to distinguish between vaccine hesitancy.

In Chapter 3, a global analysis was performed by dichotomising the polytomous Likert response data. This approach, however, results in a loss of statistical information and – in the context of IRT – provides less variation in the latent variable, which has consequences when seeking associated variables [249]. In this chapter – in which a more thorough analysis of these data is performed than in Chapter 3 – a polytomous IRT model is required.

4.3.2 The graded response IRT model for polytomous responses

The graded response model (GRM) is designed to model Likert-scale response data [239]. Responses to the four vaccination survey data items can take one of five values, $Y_{im} \in \{1, \dots, k, \dots, K\}$ (where $K = 5$ in this case). These five values correspond to each of the five possible Likert responses: $k = 1$ (strongly disagree), $k = 2$ (tend to disagree), $k = 3$ (do not know), $k = 4$ (tend to agree), or $k = 5$ (strongly agree). The objective is to model the response probabilities $P(Y_{im} = k|\eta_i, \boldsymbol{\theta})$, where η_i is vaccine hesitancy index of individual i and $\boldsymbol{\theta}$ are the graded response model parameters. GRM achieves this by modelling the cumulative probabilities,

$$P(Y_{im} \geq k|\eta_i, \boldsymbol{\theta}_{m,k}) = \psi(\eta_i; \boldsymbol{\theta}_{m,k}),$$

where ψ is some cumulative distribution function and $\boldsymbol{\theta}_{m,k}$ are the GRM parameters for item m and category k . A common choice for the cumulative distribution function is the cumulative logistic distribution (or logistic function),

$$P(Y_{im} \geq k|\eta_i, a_m, b_{m,k}) = \frac{1}{1 + \exp(-a_m(\eta_i - b_{m,k}))}, \quad (4.2)$$

where $k = 2, \dots, K$, and $\boldsymbol{\theta}_{k,m} = \{a_m, b_{m,k}\}$ are the graded response model parameters. It is clear that $P(Y_{im} \geq 1|\eta_i, a_m, b_{m,k}) = 1$, so the probabilities for each individual item response can be easily derived as,

$$\begin{aligned} P(Y_{im} = k|\eta_i) &= P(Y_{im} \geq k|\eta_i, a_m, b_{m,k}) - P(Y_{im} \geq k+1|\eta_i, a_m, b_{m,k+1}) \\ &= \frac{1}{1 + \exp(-a_m(\eta_i - b_{m,k}))} - \frac{1}{1 + \exp(-a_m(\eta_i - b_{m,k+1}))}, \end{aligned} \quad (4.3)$$

for $k = 1, \dots, K$, and where $P(Y_{im} \geq K+1|\eta_i) = 0$. The graded response model is so called because the parameter set $\{b_{km} : m = 1, \dots, M; k = 2, \dots, K\}$ represents ordered boundaries of the cumulative distribution (we have that $b_{2,m} < b_{3,m} < \dots < b_{K,m}$, for all m). In an educational-testing context, these boundaries are often selected to assign a given

proportion of students particular grades. If the probability of obtaining higher scores on items decreases with an individual's vaccine hesitancy, then the constraint $a_m > 0$ for all m must be applied. This constraint enforces ordinality with respect to latent vaccine hesitancy.

4.3.3 Model specification

In Section 4.3.2 an individual's response to vaccination survey data was related to their latent vaccine hesitancy through Equations (4.2) and (4.3). In order to perform parameter inference, we need to establish the model likelihood.

Letting $\mathbf{Y}_i$ denote a 4-dimensional column vector of polytomous item responses to vaccination survey data for individual i , the matrix of multivariate item response data can be written as $\mathbf{Y} = [\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_N]^T$ (N is the number of individuals). The likelihood is the probability $P(\mathbf{Y}|\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the model parameter set. The likelihood is obtained by writing the joint density for observed and latent data and marginalising over the latent variables,

$$P(\mathbf{Y}|\boldsymbol{\theta}) = \prod_{i=1}^N \int P(\mathbf{Y}_i, \eta_i|\boldsymbol{\theta}) d\eta_i. \quad (4.4)$$

Applying the multiplication rule and invoking conditional independence between the item responses yields,

$$\begin{aligned} P(\mathbf{Y}|\boldsymbol{\theta}) &= \prod_{i=1}^N \int P(\mathbf{Y}_i|\boldsymbol{\theta}, \eta_i) P(\eta_i) d\eta_i, \\ &= \prod_{i=1}^N \int \prod_{m=1}^M P(Y_{im}|\boldsymbol{\theta}, \eta_i) P(\eta_i) d\eta_i, \end{aligned} \quad (4.5)$$

where $P(Y_{im}|\boldsymbol{\theta}, \eta_i)$ is the probability density over the item responses and $P(\eta_i)$ is a specified distribution for the latent hesitancy values [246]. The assumption of conditional independence often simplifies the statistical modelling process (verification of this assumption is discussed in Appendix C.2.1).

For polytomous item response data, the categorical distribution specifies a probability distribution over the vaccination response data,

$$P(Y_{im}|\eta_i, \boldsymbol{\theta}) = \prod_{k=1}^K P(Y_{im} = k|\eta_i, \boldsymbol{\theta})^{\mathbb{1}[Y_{im}=k]},$$

where $\mathbb{1}[Y = y]$ is an Iverson bracket ($\mathbb{1}[Y = y] = 1$ if $Y = y$ and 0 otherwise). The GRM folds into this model through $P(Y_{im} = k|\eta_i, \boldsymbol{\theta})$ which was introduced in the preceding section.

In a Bayesian setting, the object of focus is the posterior distribution over the unknown model parameters $\boldsymbol{\theta}$. Invoking Bayes' theorem yields the following posterior distribution over these parameters,

$$P(\boldsymbol{\theta}|\mathbf{Y}) \propto \prod_{i=1}^N \int P(\mathbf{Y}_i|\boldsymbol{\theta}, \eta_i) P(\eta_i) d\eta_i \times P(\boldsymbol{\theta}), \quad (4.6)$$

where $P(\boldsymbol{\theta})$ is the prior distribution over item response parameters, which remain to be specified [246] (see Section 4.4.1 for further details).

Model identification

Model identification issues occur when more than one set of model parameter values gives the same probability of item response. A number of restrictions must be applied to the graded response model parameters to prevent parameter estimation issues caused by unidentifiable parameters.

Taking the two-parameter IRT model set out in Equation (4.1) as an example, consider a transformation of the latent variable $\eta_i \rightarrow c_1\eta_i + c_2$ (where c_1 and c_2 are constants); the logistic function can be re-expressed as,

$$\begin{aligned} P(Y_{im} = 1|\eta_i, a_m, b_m) &= \frac{1}{1 + \exp(-\tilde{a}_m(\eta_i - \tilde{b}_m))} \\ &= P(Y_{im} = 1|\eta_i, \tilde{a}_m, \tilde{b}_m), \end{aligned}$$

where $\tilde{a}_m = c_1 a_m$ and $\tilde{b}_m = \frac{b_m c_1 - c_2}{c_1}$, thus two sets of parameters have identical response probabilities. This causes issues in model inference; in iterative inference schemes, for example, successive estimates of parameters can ‘drift’ in parameter space, yet still correspond to areas of high likelihood or high posterior density.

Although identification can be achieved through post-hoc rescaling of parameters [250], there are two simpler approaches with equivalent outcomes [246]. The first is to fix parameters of the distribution over the latent variable $P(\eta_i)$, for example, by choosing $P(\eta_i) \sim \mathcal{N}(0, 1)$. This simple identifying restriction is not always possible, however, and this is especially true where structural models are introduced to explain variability in the latent variable since the mean is not fixed (see Section 4.3.4). A more universally robust method to identify the model is to place constraints on the item parameters; in the case of the two-parameter model and graded response model, fixing a single discrimination and difficulty parameter will identify the scale and position of the latent variables [246]. Another method – often preferred as it fixes the latent variables close to the origin when there are extreme (or relatively extreme) values of difficulty and discrimination – is to constrain the sum of the difficulty parameters (for a single item) to zero and the product of the discrimination parameters to one [246, 251].

These sum and product constraints have the effect of fixing the location of the latent variables close to zero, with a fixed variance. This latter procedure is used in this thesis (see Section 4.5), where an additional rescaling of hesitancy (and thus model parameters) is used to fix the vaccine hesitancy onto a scale that is easily-interpreted scale (see Section 4.5.2).

4.3.4 Structural equation modelling

In addition to inference of item response parameters and individuals' latent hesitancy, we wish to understand how variance between individuals' hesitancy can be explained using socio-economic explanatory data. This relationship between the latent trait and explanatory data is often obtained through a two-step approach, wherein the latent variable is first inferred and then used as a dependent variable in subsequent regression. This approach suffers from some flaws, notably inefficient regression parameter estimates (such as an underestimation of parameters' standard errors) and biased parameter estimates resulting from an assumption that the inferred latent variables are errorless [252]. Item response theory allows us to simultaneously regress hesitancy against explanatory data *and* perform item parameter and latent variable inference: this is achieved by including a *structural equation model* within the item response theory framework [246].

To incorporate this structural model, the latent vaccine hesitancy variable η_i is conditioned on explanatory data. The item response theory likelihood specified in Equation (4.5) then becomes,

$$P(\mathbf{Y}|\boldsymbol{\theta}, \mathbf{X}) = \prod_{i=1}^N \int \underbrace{P(\mathbf{Y}_i|\eta_i, \boldsymbol{\theta}^{\text{meas}})}_{\text{measurement model}} \times \overbrace{P(\eta_i|\mathbf{X}_i, \boldsymbol{\theta}^{\text{struct}})}^{\text{structural model}} d\eta_i, \quad (4.7)$$

where $\mathbf{X} \in \mathbb{R}^{N \times (n_X)}$ is a design matrix of n_X explanatory variables, and where we make explicit the *measurement* model (that is, the graded response model) and the *structural model* which is our hesitancy regression model. The parameter set is now decomposed: $\boldsymbol{\theta} = \{\boldsymbol{\theta}^{\text{meas}}, \boldsymbol{\theta}^{\text{struct}}\}$, where $\boldsymbol{\theta}^{\text{meas}}$ is the set of item response parameters in the measurement model, and $\boldsymbol{\theta}^{\text{struct}}$ are the parameters in the structural regression model.

In IRT the latent variable is continuous. A linear regression model is an appropriate structural equation model for a continuous response variable,

$$P(\eta_i|\mathbf{X}_i, \boldsymbol{\theta}^{\text{struct}}) = \mathcal{N}(\mathbf{X}_i\boldsymbol{\beta}, \sigma^2), \quad (4.8)$$

where $\boldsymbol{\theta}^{\text{struct}} = \{\boldsymbol{\beta}, \sigma^2\}$, $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_{n_X-1})^T$ is a set of regression parameters, and σ^2 is unexplained variance in the regression [246].

4.4 Bayesian hierarchical item response theory for vaccine hesitancy

4.4.1 Bayesian hierarchical IRT

Item response theory is now generalised for hierarchical data. The aims and benefits of modelling within a hierarchical paradigm were previously stated in Section 3.3, so they are not recapitulated here. Details of the following outline of hierarchical item response theory model are found in a number of sources [246, 253, 254].

The model likelihood (omitting the dependence on observed covariate data for notational convenience) is given by

$$P(\mathbf{Y}|\boldsymbol{\eta}, \boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\gamma}, \Sigma) \propto \prod_{j=1}^J \prod_{i=1}^{n_j} P(\mathbf{Y}_{ij}|\eta_{ij}, \boldsymbol{\theta}^{\text{meas}}) P(\eta_{ij}|\mathbf{X}_{ij}, \boldsymbol{\beta}_j, \sigma^2) P(\boldsymbol{\beta}_j|\mathbf{Z}_j, \boldsymbol{\gamma}, \Sigma), \quad (4.9)$$

where $\boldsymbol{\theta}^{\text{meas}} = \{a_k, b_{k,m}\}_{k=1,\dots,5; m=1,\dots,4}$, and $\boldsymbol{\theta}^{\text{struct}} = \{\boldsymbol{\beta}_j, \Sigma, \gamma, \sigma^2\}_{j=1,\dots,J}$, so $\boldsymbol{\theta} = \{\boldsymbol{\theta}^{\text{meas}}; \boldsymbol{\theta}^{\text{struct}}\}$. The conditional independence assumption is asserted so that,

$$P(\mathbf{Y}_{ij}|\eta_{ij}, \boldsymbol{\theta}^{\text{meas}}) = \prod_{m=1}^M P(Y_{ijm}|\eta_{ij}, \boldsymbol{\theta}^{\text{meas}}), \quad (4.10)$$

and Y_{ijm} is modelled as a categorical random variable,

$$P(Y_{ijm}|\eta_{ij}, \boldsymbol{\theta}^{\text{meas}}) = \prod_{k=1}^K P(Y_{ijm} = k|\eta_{ij}, \boldsymbol{\theta}^{\text{meas}})^{\mathbb{1}[Y_{ijm}=k]}. \quad (4.11)$$

The posterior distribution is given by Equation (4.9) multiplied by prior distributions for $\boldsymbol{\eta}$, $\boldsymbol{\beta}$, $\boldsymbol{\theta}$, $\boldsymbol{\gamma}$, and Σ . The structural equation model in Equation (4.8) is extended to allow for country-specific parameters (this was done previously in Section 3.3 in the case of hierarchical logistic regression). This hierarchical structural model is given by the equations,

$$P(\eta_{ij}|\mathbf{X}_{ij}, \boldsymbol{\theta}^{\text{struct}}) = \mathcal{N}(\mathbf{X}_{ij}\boldsymbol{\beta}_j, \sigma^2) \quad (4.12)$$

$$P(\boldsymbol{\beta}_j|\boldsymbol{\gamma}, \Sigma) = \mathcal{N}(\boldsymbol{\gamma}\mathbf{Z}_j, \Sigma), \quad (4.13)$$

where $\boldsymbol{\beta}_j = (\beta_{0j}, \beta_{1j}, \dots, \beta_{n_X-1,j})^T$ are the random-effect parameters for country j , and σ^2 is the individual-level variance [246]. At the country-level (Equation (4.13)), $\boldsymbol{\gamma} \in \mathbb{R}^{n_X \times n_Z}$ are the model fixed effects, $\mathbf{Z}_j \in \mathbb{R}^{n_Z \times 1}$ is a column vector of country-level explanatory data, and Σ is the covariance matrix, $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_{n_Z}^2)$. Individuals are labelled from $i = 1, \dots, n_j$ (where n_j is the number of individuals in country j) and countries range from $j = 1, \dots, J$ ($J = 72$ in this chapter). The number of individual- and country-level explanatory variables

are given by n_X and n_Z , respectively.

The posterior distribution is obtained by multiplying Equation (4.9) by prior distributions.

4.4.2 Prior distributions

In the graded response model, the difficulty parameters $b_{m,k}$ ($m = 1, \dots, M; k = 1, \dots, K$) satisfy the order constraints $b_{m,1} < b_{m,2} < \dots < b_{m,K-1}$ and the discrimination parameters $a_m > 0$ for all m . These conditions place restrictions on the choice of prior distributions for these parameters. A typical choice for the prior distribution for the discrimination and difficulty parameters is an exchangeable prior distribution (that is, a distribution with invariance under permutation of its parameters), so the order constraint can be readily satisfied [246]. An appropriate exchangeable distribution is the multivariate normal distribution,

$$\begin{pmatrix} \tilde{a}_m \\ \tilde{b}_{m,1} \\ \vdots \\ \tilde{b}_{m,K-1} \end{pmatrix} \sim \mathcal{N}(\boldsymbol{\mu}_{\text{item}}, \Sigma_{\text{item}}),$$

where $a_m = \exp(\tilde{a}_m)$ so that ordinality with respect to the latent variable is enforced and the difficulty parameters are reordered so that $b_{m,n} = \min_n \{\tilde{b}_{m,1}, \dots, \tilde{b}_{m,K-1}\}$ ($n = 1, \dots, K-1$), where $\min_n S$ denotes the n^{th} smallest element of set S . The parameters $\boldsymbol{\mu}_{\text{item}}$ and Σ_{item} are hyper-parameters for the mean and covariance of the item response parameters.

Choices of these hyper-parameters can be motivated through a variety of factors, such as past analyses, the data, and domain knowledge about the parameters [255]. In the case of the discrimination parameters, we find that large values of a_m can cause model identification issues if the hesitancy values are confined to a scale of order one (that is, if $\eta_i \sim \mathcal{O}(1)$ and $a_m \gg 1$, then $P(Y_{ijm} = k|a_m) \approx P(Y_{ijm} = k|a_m + \Delta)$ for a large range of values Δ), and so an uninformative prior over a_m can lead to convergence issues. In this study, hyper-priors are placed over $\boldsymbol{\mu}_{\text{item}}$ and Σ_{item} to represent prior uncertainty. Discrimination values are restricted to values of order 1 to prevent the model convergence issues discussed in Section 4.3.3.

A range of priors are also possible for the structural parameters, $\boldsymbol{\theta}^{\text{struct}}$. The choice of prior for the fixed effect parameters $\boldsymbol{\gamma}$ can be motivated by knowledge of the range of values for the latent hesitancy variables, since these fixed-effects are tied to hesitancy through the structural equation model. As discussed in Section 4.3.3, a constraint on the sum of the difficulty parameters and the product of discrimination parameters fixes the scale and location of the latent variable; in particular, these can be selected so that η_i is centred close

to zero with a variance of order one. Appropriate priors for the fixed-effects are therefore independent normal priors with a semi-informative variance component, $\gamma_{pq} \sim \mathcal{N}(0, 10)$ ($p = 1, \dots, n_X; q = 1, \dots, n_Z$). An inverse-gamma prior is chosen for the prior distribution $P(\sigma^2)$ (with hyper-priors motivated through simulation study – see Appendix C.1) and an inverse-Wishart prior is used for the covariance matrix Σ (the robustness of which is again shown in Appendix C.1).

We specify the following prior distributions for the Bayesian hierarchical IRT model specified in Section 4.4.1,

$$\begin{aligned} (\tilde{a}_m, \mathbf{b}_m)^T &\sim \mathcal{N}(\boldsymbol{\mu}_{\text{item}}, \Sigma_{\text{item}}) \\ \boldsymbol{\mu}_{\text{item}} | \Sigma_{\text{item}} &\sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{item}}) \\ \Sigma_{\text{item}}^{-1} &\sim \text{Wishart}(\mathbf{1}, n_{\text{item}} + 1), \end{aligned} \tag{4.14}$$

where $a_m = \exp(\tilde{a}_m)$. The prior distributions for the variance components are specified as,

$$\begin{aligned} \sigma_\eta^{-2} &\sim \text{Gamma}(1, 0.01) \\ \Sigma^{-1} &\sim \text{Wishart}(0.01 \times \mathbf{1}, n_x + 1), \end{aligned} \tag{4.15}$$

where n_x and n_{item} in the Wishart distributions contribute to prior non-informativeness [246].

A simulation and sensitivity study is performed in Appendix C.1 to demonstrate the robustness of parameter inference to the priors outlined in this section. For example, the sensitivity analysis of item response parameters shows that inferences under the specified priors are robust to a wide range of generated item response parameters, including those outside the range of values anticipated in this study (see Appendix C.1.3). Although there has been some criticism of the choice of the family of Gamma priors for hierarchical variance components [256], our sensitivity analysis shows that careful selection of the variance hyper-parameters for these distributions in this context – in which much is known about the range of parameter values – provides robust inferences over a range of possible parameter values.

4.4.3 Structural models for vaccine hesitancy

A number of structural models are considered to find an appropriate model for fitting explanatory data to inferred hesitancy. Model-fit criteria (see Section 4.4.4) are used to assess the best-fitting model. In the following models, $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$ and $\nu_{pj} \sim \mathcal{N}(0, \sigma_p^2)$ for

$p = 1, \dots, n_X$:

$$\begin{aligned}\mathcal{M}_1 : \eta_{ij} &= \beta_{0j} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_{00} + \nu_{0j},\end{aligned}$$

model $\mathcal{M}_1$ denotes an ‘empty’ model in which no explanatory data at either level are used to explain individual- or country-level variance;

$$\begin{aligned}\mathcal{M}_2 : \eta_{ij} &= \beta_{0j} + \beta_{1j} \cdot \mathbf{x}_{\text{SEX},ij} + \beta_{2j} \cdot \mathbf{x}_{\text{AGE},ij} + \beta_{3j} \cdot \mathbf{x}_{\text{EDU},ij} + \beta_{4j} \cdot \mathbf{x}_{\text{EMP},ij} + \epsilon_{ij} \\ \beta_{ij} &= \gamma_{i0} + \nu_{ij},\end{aligned}$$

$\mathcal{M}_2$ seeks to explain variance at the individual-level only by using an individual’s sex, age, education, and employment (SEX, AGE, EDU, EMP as described in Section 3.2.1, respectively)¹;

$$\begin{aligned}\mathcal{M}_3 : \eta_{ij} &= \beta_{0j} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_{00} + \gamma_{01}Z_{\text{GDP},j} + \gamma_{02}Z_{\text{SAN},j} + \gamma_{03}Z_{\text{WAT},j} + \gamma_{04}Z_{\text{BIR},j} + \gamma_{05}Z_{\text{U5M},j} + \gamma_{06}Z_{\text{HEX},j} + \nu_{0j},\end{aligned}$$

$\mathcal{M}_3$ uses the country-level variables GDP (GDP), access to sanitation (SAN), access to water (WAT), births attended by skilled health staff (BIR), and under-five mortality (U5M) to explain country-level variance (these variables were introduced in Section 3.2.2);

$$\begin{aligned}\mathcal{M}_4 : \eta_{ij} &= \beta_{0j} + \beta_{1j} \cdot \mathbf{x}_{\text{SEX},ij} + \beta_{2j} \cdot \mathbf{x}_{\text{AGE},ij} + \beta_{3j} \cdot \mathbf{x}_{\text{EDU},ij} + \beta_{4j} \cdot \mathbf{x}_{\text{EMP},ij} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_{00} + \gamma_{01}Z_{\text{GDP},j} + \gamma_{02}Z_{\text{SAN},j} + \gamma_{03}Z_{\text{WAT},j} + \gamma_{04}Z_{\text{BIR},j} + \gamma_{05}Z_{\text{U5M},j} + \gamma_{06}Z_{\text{HEX},j} + \nu_{0j},\end{aligned}$$

$\mathcal{M}_4$ is the *intercepts-as-outcomes* model used in the preceding chapter in Section 3.3, which includes individual-level explanatory data and country-level data to explain variation in the intercepts β_{0j} ; and,

$$\begin{aligned}\mathcal{M}_5 : \eta_{ij} &= \beta_{0j} + \beta_{1j} \cdot \mathbf{x}_{\text{SEX},ij} + \beta_{2j} \cdot \mathbf{x}_{\text{AGE},ij} + \beta_{3j} \cdot \mathbf{x}_{\text{EDU},ij} + \beta_{4j} \cdot \mathbf{x}_{\text{EMP},ij} + \epsilon_{ij} \\ \beta_{pj} &= \gamma_{00} + \gamma_{01}Z_{\text{GDP},j} + \gamma_{02}Z_{\text{SAN},j} + \gamma_{03}Z_{\text{WAT},j} + \gamma_{04}Z_{\text{BIR},j} + \gamma_{05}Z_{\text{U5M},j} + \gamma_{06}Z_{\text{HEX},j} + \nu_{pj},\end{aligned}$$

for $p = 1, \dots, n_X$; model $\mathcal{M}_5$ is the *intercepts-and-slopes-as-outcomes* model also introduced previously in Section 3.3 and, as the name suggests, uses explanatory data to explain variance in the intercept and slope random effects. The *one-hot* matrix notation used in Equation (3.1) is again used here. Unlike the *logistic* hierarchical models in Chapter 3, the models $\mathcal{M}_1$ to $\mathcal{M}_5$ here are *linear*: thus fixed-effect parameters are directly associated with an increase in hesitancy relative to the baseline group (rather than associated through odds ratios in

Section 3.4). The baseline group is an 18-44 year-old employed male with no education.

4.4.4 Model estimation

Analytical expressions for complicated probability densities – such as multidimensional integrals arising as posterior densities from item response theory in Equation (4.6) – can often be hard, if not impossible, to obtain [255]. A suite of methods exist to either solve or approximate these distributions, or to find desirable properties of them. Sampling techniques, such as the suite of Markov chain Monte Carlo (MCMC) methods (see, for example [257]), provide a sample of the posterior density (from which statistical properties can be obtained) through a Markov random walk through the parameter space. The posterior density (or an approximation to it) can also be obtained through methods seeking to make distributional simplifications, such as in variational Bayes and Laplace’s method [258]. Point estimates – rather than estimating the full posterior distribution – can also be obtained through numerical integration methods or iterative schemes such as expectation maximisation [255]. In this thesis, we focus on Gibbs sampling, an MCMC scheme that provides a sample from the posterior density through successive sampling of full conditional distributions [255].

Gibbs sampling

The goal of MCMC procedures is to construct a Markov chain of values whose stationary distribution converges to the posterior distribution of interest; that is, as we take an infinite number of samples the probability of visiting a set of parameters $\boldsymbol{\theta}$ is $P(\boldsymbol{\theta})$. Although there are many incarnations of MCMC, we focus on Gibbs sampling, formally introduced by Geman and Geman [259]. Gibbs sampling is an attractive choice of MCMC method when the full conditional distributions of each parameter of interest are known, which is the case with a choice of (conditionally conjugate) priors for hierarchical item response theory, which were introduced in Section 4.4.1.

Gibbs sampling involves generating an (ergodic) Markov chain of parameters $\{\boldsymbol{\theta}_s\}_{s=1}^S$, where S is the number of MCMC iterations, by sequentially sampling from the full conditional distributions of each parameter,

$$\begin{aligned}\theta_1^{(s+1)}|\cdot &\sim P(\theta_1|\theta_2^{(s)}, \theta_3^{(s)}, \dots, \theta_p^{(s)}) \\ \theta_2^{(s+1)}|\cdot &\sim P(\theta_2|\theta_1^{(s+1)}, \theta_3^{(s)}, \dots, \theta_p^{(s)}) \\ &\vdots \\ \theta_p^{(s+1)}|\cdot &\sim P(\theta_p|\theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{p-1}^{(s+1)}),\end{aligned}$$

where $|\cdot$ represents a conditioning on the most recent updates of the random vector $\boldsymbol{\theta}$ [255].

Each of the conditional distributions above satisfy detailed balance equations which, along with Markov chain ergodicity, result in sampling from a stationary distribution which asymptotically coincides with the posterior distribution. Ergodicity allows us to start the Markov chain from any initial position (as Gibbs sampling is asymptotically exact) but, in the finite-sample limit, choices far from the target distribution may take a long time to reach high-probability areas of the target distribution [190]. This may result in a large ‘burn-in’ period not representative of the posterior distribution (and thus needs to be discarded).

Convergence diagnostics

A number of convergence diagnostics exist for determining whether the MCMC sample has successfully converged to the target distribution (and thus also whether the sample has successfully burnt-in): the Gelman-Rubin and Raferty-Lewis [260, 261] diagnostics, and the Geweke statistic [262] are prevalently used. In this thesis we consider the latter of these diagnostic tools.

If parameter samples have converged to their stationary distribution then the means from two windows of this time-series should not differ and a two-sample t -test can be used to test this difference. For a stationary sequence (implying a constant mean), then the two-sample T -statistic is given by,

$$T = \frac{\hat{\mu}_\theta^A - \hat{\mu}_\theta^B}{n_A^{-1}\hat{S}^A(0) + n_B^{-1}\hat{S}^B(0)}, \quad (4.16)$$

where $\hat{\mu}_\theta^{A(B)}$ is an estimate of the sample mean across window A (B) of length $n_{A(B)}^{-1}$, and $\hat{S}^{A(B)}(0)$ is an estimate of the spectral density which provides an estimate for the time-series sample variances [262]. Under the assumption of stationarity T follows a standard normal distribution as $n_A, n_B \rightarrow \infty$, $T \sim \mathcal{N}(0, 1)$. The p -value associated with the test statistic T is known as the Geweke statistic and p -values with $p < \alpha$ provide evidence for non-converged chains (at significance level α) [262]. To simultaneously assess burn-in and convergence, we compute the Geweke statistic for the first tenth and final half of the MCMC chains. Successful burn-in and convergence are verified for our models in Appendix C.2.2.

Model selection

In a Bayesian formalism, the model selection procedure involves choosing a model (for example, from two competing models $\mathcal{M}_1$ and $\mathcal{M}_2$) if the posterior model odds favours model $\mathcal{M}_1$; that is, if

$$\text{BF} = \frac{P(\mathcal{M}_1|Y)}{P(\mathcal{M}_2|Y)} > 1,$$

where Y are data, and where the ratio of posterior model odds is known as the Bayes factor. A Bayes Factor greater than about three provides substantial evidence to favour model

$\mathcal{M}_1$ [263]. Application of Bayes' rule to this expression allows us to re-write the Bayes factor as,

$$\frac{P(\mathcal{M}_1|Y)}{P(\mathcal{M}_2|Y)} = \frac{P(Y|\mathcal{M}_1) P(\mathcal{M}_1)}{P(Y|\mathcal{M}_2) P(\mathcal{M}_2)}. \quad (4.17)$$

The marginal likelihoods which comprise the Bayes factor are integrals over the parameter vector $\boldsymbol{\theta}_i$ which parametrise model $\mathcal{M}_i$:

$$P(Y|\mathcal{M}_i) = \int P(Y|\mathcal{M}_i, \boldsymbol{\theta}_i) P(\boldsymbol{\theta}_i|\mathcal{M}_i) d\boldsymbol{\theta}_i, \quad (4.18)$$

for $i = 1, 2$. Unless a simple analytic form exists to this equation (which is seldom the case), this integral requires numerical integration techniques to evaluate.

A number of model-assessing information criteria exist. These criteria usually have the form $-2 \log \mathcal{L} + \Lambda$, where the model's log-likelihood $\log \mathcal{L}$ is penalised by an amount Λ , usually depending on the number of model parameters [255].

Bayesian information criterion The Bayesian information criterion (BIC) is a widely-used approximation to the Bayes factor [264] and is given by,

$$\text{BIC} = -2 \log \mathcal{L}_{\text{MAP}} + p \log N, \quad (4.19)$$

where $\log \mathcal{L}_{\text{MAP}}$ is the log-likelihood evaluated at the maximum a posteriori estimate (this is the mode of the posterior distribution), p is the number of model parameters, and N is the number of data points in the model. Thus, overfitting is guarded against as the likelihood is penalised by a term depending on the number of variables in the model.

Aikake information criterion Another commonly used measure of model fit is the Akaike information criterion (AIC) which quantifies the dissimilarity between a 'true' model and an observed model [265]. The AIC is given by,

$$\text{AIC} = -2 \log \mathcal{L}_{\text{ML}} + 2p, \quad (4.20)$$

where $\log \mathcal{L}_{\text{ML}}$ is the log-likelihood evaluated at the maximum likelihood estimate.

Deviance information criteria Deviance information criteria can be used for hierarchical models when the true number of parameters is masked by the hierarchical structure [266]. The deviance information criteria (DIC) penalises the log likelihood by a measure of model di-

mension proposed to combat ambiguity in parameter numbers in hierarchical models,

$$\text{DIC} = -2 \log \mathcal{L}(\bar{\boldsymbol{\theta}}) + 2p_D, \quad (4.21)$$

where $p_D = D(\bar{\boldsymbol{\theta}}) - \bar{D}(\boldsymbol{\theta})$ (where $D = -2 \log \mathcal{L}$ is the deviance), and where the likelihood is calculated at the posterior means $\bar{\boldsymbol{\theta}}$ [266]. An alternative to the DIC stated above proposed by Gelman replaces p_D with $p_V = \frac{1}{2} \widehat{\text{var}}(D(\boldsymbol{\theta}))$, as this quantity is an estimate of the number of free parameters in a non-hierarchical model and is trivial to calculate [255].

4.5 Results

Individuals' vaccine hesitancy is obtained via the Bayesian hierarchical IRT model introduced in Section 4.4.1. Global and national trends in this latent hesitancy index are investigated in Section 4.5.2. Socio-economic determinants of hesitancy are sought in Section 4.5.3 through the structural equation models introduced in Section 4.4.3. In Section 4.5.4 a vaccine hesitancy study is performed in Nigeria: we use the inferred hesitancy indices of individuals in Nigeria to map hesitancy and use the extended set of socio-economic determinants (SAHEL-NIGERIA introduced in Section 4.2.2) to explore variation in hesitancy in Nigeria. Results are fully discussed and interpreted in Section 4.6. We begin this chapter with a brief description of the model implementation procedure.

4.5.1 Model implementation and diagnostics

The Bayesian hierarchical IRT models outlined in Section 4.4.1 with the five structural equation models $\mathcal{M}_1$ to $\mathcal{M}_5$ are implemented using the JAGS programming language using the R package **rjags** [267].² Each model is run for a total of 51,000 iterations with the first 1,000 discarded as the burn-in phase and a thinning interval of five to avoid memory issues. Evidence for successful burn-in and model convergence for all models is provided by the Geweke statistic (Appendix C.2.2). Strong support is also found for conditional independence of observed vaccination response data given the inferred latent hesitancy (Appendix C.2.1).

Model assumptions for hierarchical models have been previously introduced in Section 3.3.2. In this chapter, linear regression (rather than the logistic regression) is used; additional tests for homoscedasticity and residual normality is thus required at the individual-level. We verify these assumptions in Appendix C.2.3.

Although the condition number of the individual- and country-level explanatory data matrix reveals a well-conditioned matrix (condition number of 13 which is below the value

²Our JAGS code for hierarchical Bayesian IRT is provided in Appendix C.3.

of 30 [229]), preliminary analysis of MCMC chains revealed poor convergence of country-level variables and an inflation of standard errors, particularly of GDP per capita (GDP) and health expenditure (HEX). Correlation between GDP and HEX was high (Pearson's $r = 0.92$); although z -scoring of country-level data appeared to have removed ill-conditioning, it is likely that the Gibbs sampler became 'trapped' along the highly correlated surface of these two variables and led to these issues. The country-level variable GDP is therefore removed from the analysis. This was not an issue for the hierarchical analysis in Chapter 3.4.2, a possible reason for this is the addition of the Sahel data which induces a higher correlation between GDP and HEX. The correlation between the country-level data in this chapter and the preceding chapter remained high (as expected) with Pearson correlations of: 0.77 (HEX), 1.00 (SAN), 0.96 (WAT), 0.96 (BIR), 1.00 (U5M).

4.5.2 Inference of vaccine hesitancy

The mean of the posterior distribution of an individual's inferred latent hesitancy,

$$\eta_{ij} = \frac{1}{S} \sum_{s=1}^S \eta_{ij}^s, \quad (4.22)$$

is termed the *vaccine hesitancy index* (VHI), where $S = 50,000$. Low values of η_{ij} represent individuals with high vaccine hesitancy. We note that throughout what follows, *high values of the vaccine hesitancy index are associated with low levels of individual vaccine hesitancy*.

In Figure 4.2 A-D, we plot the probability of responding with a particular Likert response at a given vaccine hesitancy value (that is, we plot $P(Y_{ijm} = k|\eta)$ against η for all individuals). This is done for each vaccination survey question: overall I think vaccines are important for children to have (A); overall I think vaccines are safe (B); vaccines are effective (C); and vaccines are compatible with my religious beliefs (D). The global distribution of the vaccine hesitancy index is shown as a grey smoothed histogram and these values are scaled such that an individual with hesitancy index $\eta_{ij} = 0$ has a 50% chance of replying **strongly agree** to Vaccinations are important for children to have. For a given η_{ij} , the probability of the Likert responses must sum to 1.

A number of interesting trends in vaccine hesitancy behaviours are revealed using this item response theory approach. The distribution of individuals' VHIs (grey smoothed histograms) reveals multiple peaks suggesting a clustering of vaccination behaviour. For instance, individuals within the right-most peak (high VHI η ; low hesitancy) almost certainly reply **strongly agree** to overall I think vaccines are important for children to have, overall I think vaccines are safe, overall I think vaccines are effective. As the vaccine hesitancy increases from this value

(that is, η decreases), the probability of replying **strongly agree** decreases. At values of $\eta \approx 0, 0.5, 0.4$, and -0.1 , individuals become more likely to reply **tend to agree** to importance, safety, effectiveness, and religious compatibility (respectively) than **strongly agree**. A second prominent peak at $\eta = 0.28$ is still associated with less hesitant vaccination beliefs for each item; however, there is now clear vaccine scepticism towards the safety of vaccinations and with religious compatibility concerns (Figure 4.2 B and D).

The trends in the link between Likert response probabilities and vaccine hesitancy are made clearer in Figure 4.2 E which shows item response curves for positive Likert responses (**tend to agree** and **strongly agree**). This plot reveals that although the discrimination parameters are similar across the questions on vaccine safety, importance, and effectiveness, there is a notable difference in difficulty parameters. This difference suggests that more respondents are hesitant towards vaccine safety than vaccine effectiveness and, similarly, more hesitant towards vaccine effectiveness than vaccine importance. We also find that the response curve for religious compatibility in Figure 4.2 E is less *discriminative* than the other three curves. This low relative discrimination reveals that responses to the survey item on religious compatibility inform less towards vaccine hesitancy than the other three items. These trends are in agreement with the findings in Section 3.4.1; however, they provide more detailed information into individuals' responses and the correlation between vaccine survey questions.

Global trends in hesitancy

In Figure 4.2 A-D, we observe the global distribution of the vaccine hesitancy index and its relation to the probability of responding in a particular way to the vaccination survey questions. Most individuals around the world have a low vaccine hesitancy (a high vaccine hesitancy index) as evidenced by high proportions of individuals responding positively to the vaccination survey items. In Figure 4.2 F, the global distribution of the vaccine hesitancy index (grey smoothed histogram) is compared to those in the Sahel (coloured histograms). The combined Sahel region has a median VHI of 1.627 [1.62, 1.63] compared to 0.44 [0.42, 0.46] for the rest of the world highlighting the relatively low hesitancy amongst individuals in the Sahel compared to the rest of the countries surveyed. These differences in median values between the Sahel and the rest of the world represent increases in probability of positive responses of 0.5%, 2.5%, 1%, and 14% for each survey question respectively, but even larger differences exist between **strongly agree** and **tend to agree**. (The median is a bootstrapped estimate from 10,000 resamples from the distributions of η_{ij} for all individuals either in the Sahel or globally and error bars denote a 99% confidence interval). Median hesitancy values within each country in the Sahel are also significantly higher than the global average: Burkina Faso 1.002 [0.927, 1.028], Mali 1.860 [1.858, 1.863], Mauritania 1.930 [1.928, 1.930],




Figure 4.2: Inference of vaccine hesitancy reveals hesitant behaviours and permits a worldwide comparison of hesitancy (A) The probability of responding with Likert response $k \in \{ \text{strongly disagree (SD), tend to disagree (TD), do not know (DK), tend to agree (TA), strongly agree (SA)} \}$ to vaccines are important for children to have at vaccine hesitancy level η . Higher levels of η represent *lower* levels of hesitancy; individuals' vaccine hesitancy index (VHI, Equation (4.22)) are shown in the grey smoothed histogram. (B-D) repeat (A) for safety, effectiveness, and religious compatibility (respectively). (E) shows item response curves for a positive response (TA or SA) to all questions. (F) shows low hesitancy levels in the Sahel compared to globally (GLO) [GLO: global; BFA: Burkina Faso; MLI: Mali; MRT: Mauritania; NER: Niger; NGA: Nigeria; and SEN: Senegal]. (G) median-ranked box-and-whisker plots of VHIs in each country.

Niger 1.806 [1.800, 1.809], Nigeria 1.645 [1.634, 1.652], and Senegal 1.574 [1.569, 1.579].

In Figure 4.2 G box-and-whisker plots show the median (vertical bar), interquartile range (rectangular box), and range (horizontal lines) of individuals' VHI values within each country (with Sahel countries at the top of G). Most countries have a median VHI above the mean hesitancy value; however, most countries have a heavy tail of low VHI values denoting large groups of vaccine hesitant individuals. Many countries, such as Mauritania, Argentina, Ethiopia, do not have a large number of vaccine hesitant individuals, and simply have a high median VHI with low variance about this value.

In the Sahel, Burkina Faso, Senegal and Nigeria, have relatively high median VHI values signifying populations with relatively low vaccine hesitancy. However, the Sahel countries also tend to have a very large range of VHI values, suggesting groups of vaccine hesitant individuals. This variability in VHI values between countries is interpreted in the discussion section, although further work is required to fully interpret international trends more carefully before we reach publication stage. The variability in hesitant individuals is now investigated through the hierarchical structural regression models $\mathcal{M}_1$ to $\mathcal{M}_2$ introduced in Section 4.4.3.

4.5.3 Socio-economic determinants of vaccine hesitancy

Global trends

Introduced in Section 4.4.3, models $\mathcal{M}_1$ to $\mathcal{M}_5$ aim to explain variance at differing levels of the hierarchical data. Model performance is evaluated using the four information criteria in Section 4.4.4. Models $\mathcal{M}_1$ (the 'empty' model with no explanatory variables at either level), $\mathcal{M}_3$ (explanatory variables included at the country-level only), and $\mathcal{M}_5$ (the *intercepts-and-slopes-as-outcomes* model) are found to have the highest model selection criteria across all model selection criteria.³ These high model selection criteria suggest that not enough individual- and country-level variance is explained in $\mathcal{M}_1$, which does not include any explanatory data, and $\mathcal{M}_3$ which ignores variance at the individual level. The high criteria values for $\mathcal{M}_5$, however, suggest that a high degree of overfit is present in this model, as country-level parameters are included to account for variance in every random-effects term. Models $\mathcal{M}_2$ and $\mathcal{M}_4$ are favoured by the information criteria: both use individual-level explanatory data to account for individual-level variance, but $\mathcal{M}_4$ also uses country-level variables to account for variability between the baseline groups across all countries. The selection of $\mathcal{M}_2$ and $\mathcal{M}_4$ reveals that it is crucial to monitor individual-level variance, and that there is large variability amongst baseline hesitancy levels within each country that should be explained (this is

³except for $\mathcal{M}_1$ which was found to have the lowest BIC value, however this result is discarded, as it is well-known that BIC can often provide unnecessarily harsh penalties to models, and because this information criteria is particularly sensitive to hierarchical models [266].

possibly suggested by the large variability in median hesitancy in Figure 4.2 G).

Fixed-effect parameters – representing global averages of trends across all countries – for these two models are displayed in Table 4.1 (with the ‘empty’ model $\mathcal{M}_1$ included for model fit comparison). The maximum a posteriori (MAP) estimates⁴ and 95% highest posterior density (HPD) intervals⁵ are shown, and the lowest information criteria for each model is highlighted. Positive (negative) fixed-effect parameters are associated with higher (lower) VHI values and thus less (more) hesitant vaccination belief. In model $\mathcal{M}_2$, for example, the baseline group (intercept) of an employed 18-44 year-old male with no education has a hesitancy index centred at $\eta = 0.42$; an employed 18-44 year-old female with no education would have $\eta = 0.46$ (note that 0.42 represents *higher* vaccine hesitancy than 0.46.)

Across models $\mathcal{M}_2$ and $\mathcal{M}_4$, it is observed that females, older age groups, and those with the highest level of education are more likely to have higher VHI values and are therefore less likely to be vaccine hesitant (than the respective baseline group). With regards to employment, housewives are more likely to have high VHI values (less vaccine-hesitant) than other groups, whereas the unemployed are more likely to have low VHI values (more vaccine-hesitant).

Model $\mathcal{M}_4$ includes country-level predictors to explain variability between countries’ baseline groups (the intercepts β_{0j} in $\mathcal{M}_4$ in Section 4.4.3). Countries with a higher proportion of births attended by skilled health staff (BIR) are likely to have more vaccine hesitant individuals, whilst countries with high infant mortality rates (U5M) have fewer. These two effects are relatively strong compared to individual variables, reflecting the large between-country variability: this is further evidenced through the individual- and intercept random-effect variance components shown in Table 4.1.

A comparison of trends between the Sahel and the rest of the world

The random-effect parameters – which represent a country’s deviation from the fixed-effect parameters – are displayed in Figure 4.3 A for model $\mathcal{M}_4$. (There is very strong correlation – $R^2 = 0.94$ – between the random-effect parameters in $\mathcal{M}_2$ and $\mathcal{M}_4$.) Each grey point represents a country’s random effect and coloured dots highlight the position of countries in the Sahel (a Sahel datum within a ring indicates that the 95% HPD interval does not contain zero).

There is large variability within the intercept random-effect terms, highlighting the large variability in VHI values across all countries.⁶ The high median VHI values in the Sahel countries relative to those from the rest of the world are shown in the intercept term in Figure

⁴The MAP estimate is the mode of the posterior distribution.

⁵The $(1 - \alpha)$ HPD interval is the narrowest interval around the MAP estimate that contains $100(1 - \alpha)\%$ of the probability mass.

⁶As a reminder, the baseline group is an 18-44 year-old unemployed male with no education.

	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_4$
ine			
Individual-level			
Intercept	0.52 (0.41,0.63)	0.42 (0.29, 0.55)	0.42 (0.30, 0.53)
Sex			
Male (baseline)			
Female		0.04 (0.01, 0.06)	0.04 (0.01, 0.06)
Age			
18-44 (baseline)			
45-64		0.04 (0.01, 0.07)	0.04 (0.01, 0.07)
65+		0.09 (0.03, 0.15)	0.09 (0.03, 0.14)
Education			
None (baseline)			
Low		0.02 (-0.05, 0.10)	0.03 (-0.05, 0.09)
Med		0.04 (-0.03, 0.12)	0.06 (-0.02, 0.12)
High		0.09 (-0.01, 0.17)	0.09 (0.00, 0.16)
Employment			
Employed (baseline)			
Unemployed		-0.07 (-0.11,-0.04)	-0.07 (-0.10,-0.03)
Student		-0.02 (-0.06, 0.03)	-0.02 (-0.06, 0.02)
Housewife		0.04 (0.01, 0.08)	0.05 (0.01, 0.08)
Retired/disabled		0.01 (-0.03, 0.05)	0.01 (-0.02, 0.06)
Country-level			
HEX (% of GDP)			0.05 (-0.08, 0.18)
SAN (% with access)			0.06 (-0.16, 0.31)
WAT (% with access)			0.05 (-0.10, 0.21)
BIR (% of total)			-0.18 (-0.34,-0.03)
U5M (per 1000 births)			0.21 (-0.01, 0.42)
Variance components			
σ^2	1	0.86	0.86
σ_0^2	0.48	0.21	0.16
ine			
Information criteria			
AIC	639,662	637,425	637,478
BIC	640,542	646,172	646,734
DIC (pV)	637,382	633,194	633,404
DIC (pD)	957,765	961,068	955,648

Table 4.1: **Fixed-effect structural equation model parameters for the Bayesian hierarchical response theory on augmented GLOBAL-67 and SAHEL-6 reveal global trends in determinants of hesitancy** The most parsimonious models as selected by a range of information criteria are structural models with individual-level predictors only ($\mathcal{M}_2$) and the intercepts-as-outcomes model with country-level predictors ($\mathcal{M}_4$). Fixed-effect parameters greater than zero represent a positive association between that variable and positive (that is, less hesitant) vaccination beliefs (relative to the baseline group) and vice versa. Females, older age groups, highly educated groups, and housewives are associated with lower vaccine hesitancy, whilst the unemployed are associated with higher hesitancy. Countries with a higher proportion of births attended by skilled health staff (BIR) hold a negative association with the baseline random effects, whilst countries with high infant mortality rates (U5M) are associated with lower hesitancy.

4.3 A. We also note that the large random-effect variability in the intercept relative to the other parameters provides explanation for parsimony of model $\mathcal{M}_4$ relative to the *intercepts-and-slopes-as-outcomes* model ($\mathcal{M}_5$), which uses country-level socio-economic variables to explain the variance in each of the other parameters.

The global fixed-effect trends in Table 4.1 indicate that males are more likely than females to have lower VHI values – Figure 4.3 A demonstrates this trend is true for all countries in the Sahel. The global trends in age – that older age groups have higher VHI values and are thus less vaccine hesitant – holds true in Niger and Nigeria, but there is no significance in this effect elsewhere in the Sahel.

There are striking trends in the effect of education on vaccine hesitancy in the Sahel countries. Whilst globally, high levels of education are associated with high VHI values (low vaccine hesitancy), this is emphatically not the case in Burkina Faso, Mali, Niger, and Senegal (Figure 4.3 A and B). In these four countries each additional level of education is associated with decreases in their VHI values. Higher levels of education appear, therefore, to be associated with vaccine hesitancy across much of the Sahel. This effect is emphasised in Figure 4.3 B where the posterior distributions of the random-effect education terms are shown across the Sahel. To provide an idea of the strength of decreasing VHI values with increasing education levels across the Sahel, we estimate the overlap regions between posterior distributions associated with the lowest and highest VHI values in Figure 4.3 B. A p -value, defined as the number of times a randomly generated sample from the low-education posterior is also in the high-education posterior, is reported. In Burkina Faso ($p < 0.01$), Mali ($p = 0.06$), Niger ($p < 0.01$), and Senegal ($p = 0.05$) there is strong evidence to suggest that those with the highest levels of education (university and masters/PhD degrees; see 4.2.1) have lower VHI values than those with low levels (primary educated). With regards to employment status, those unemployed in Burkina Faso and students in Senegal are likely to have higher VHI values and thus be more vaccine-hesitant than those employed. Housewives in Burkina Faso, Mali, Niger, and Senegal are associated with higher VHI values and are thus less vaccine-hesitant than the baseline group (those employed).

In Figure 4.3 C, the negative association between the random-effect intercepts and a country’s health expenditure (HEX) is shown. The Sahel countries are again highlighted, as are countries at the extreme end of the hesitancy scale (Argentina and Japan) and a selection of other countries of interest (India, China, Russia, UK, USA).

4.5.4 State-level vaccine hesitancy in Nigeria

The inferred VHI values for individuals in Nigeria ($\eta_{i,\text{Nigeria}}$ for individuals i in the SAHEL-NIGERIA dataset) obtained via Bayesian hierarchical IRT model $\mathcal{M}_4$ are used to explore state-level




Figure 4.3: **Random-effect parameters in the Sahel reveal high levels of education are associated with vaccine hesitancy in Burkina Faso, Mali, Niger, and Senegal** (A) Random-effects parameters from model $\mathcal{M}_4$ are shown for the Sahel (coloured) and the rest of the world (grey). Random-effect parameters greater than zero are associated with more positive (less hesitant) vaccination views relative to the baseline category and vice versa. (A Sahel datum within a ring indicates that the 95% HPD interval does not contain zero.) The Sahel shows broad consistency with the global trends in 4.2; however, high levels of education are associated with vaccine hesitant views in the Burkina Faso, Mali, Niger, and Senegal. (B) Posterior distributions for the education random effect parameters in the Sahel further highlighting the education trend in (A). (C) Random-effect intercept terms β_{0j} against health expenditure HEX reveals countries with high national spends on healthcare tend to have lower levels of hesitancy. (Health expenditure has high correlation with GDP per capita as discussed in Section 4.5.1.)

hesitancy.⁷ As we introduce a state-level label for the VHI of individuals in Nigeria, the variable $\eta_{i,\text{Nigeria}}$ is relabelled $\eta_{is}^{\text{Nigeria}}$ in this section, where the subscript s denotes Nigerian state.

State-level variation in Nigerian VHI values is mapped to explore trends across Nigeria's 37 states; variation in these trends is explored via a linear regression model using the extended set of explanatory variables in Nigeria outlined in Section 4.2.2.⁸ In this regression model, we include interaction terms between variables to seek more nuanced effects between factors and hesitancy. The choice of interaction terms was motivated by domain knowledge of likely vaccine hesitancy-influencing factors: to probe the effect that religious leaders may have on vaccine hesitancy, we explore the interaction between ethnicity and religious schooling, and we also consider the interaction between ethnicity and sex, and sex and education (this latter interaction is motivated by the findings in Chapter 2). The following hierarchical linear regression model is therefore used:

$$\begin{aligned}\eta_{is}^{\text{Nigeria}} &= \beta_{0s} + \beta_{1s} \cdot \mathbf{x}_{\text{SEX},is} + \beta_{2s} \cdot \mathbf{x}_{\text{AGE},is} + \beta_{3j} \cdot \mathbf{x}_{\text{EDU},is} + \beta_{4j} \cdot \mathbf{x}_{\text{EMP},is} \\ &+ \beta_{5j} \cdot \mathbf{x}_{\text{ETH},is} + \beta_{6j} \cdot \mathbf{x}_{\text{REL},is} + \beta_{7j} \cdot \mathbf{x}_{\text{DEC},is} + \beta_{8j} \cdot \mathbf{x}_{\text{EC},is} \\ &+ \beta_{9j} \cdot \mathbf{x}_{\text{PEC},is} + \beta_{9j} \cdot \mathbf{x}_{\text{ETH:REL},is} + \beta_{9j} \cdot \mathbf{x}_{\text{SEX:ETH},is} + \beta_{10j} \mathbf{x}_{\text{SEX:EDU},is} + \epsilon_{is} \\ \beta_{ps} &= \gamma_{p0} + \nu_{ps},\end{aligned}\tag{4.23}$$

for $p = 1, \dots, 33$, $s = 1, \dots, 37$, ‘ $\cdot$ ’ denotes an interaction between variables, where the random-effect parameters $\beta_{[.]s}$ are again vectors with length equal to the number of levels less one, and $\nu_{ps} \sim \mathcal{N}(0, \sigma_p^2)$. The number of individuals in each state (over which the index i runs) ranges from 100 in Ebonyi to 130 in Benue. The baseline category, β_{0j} , is an 18-44 year-old employed male with no education; the extended variable baseline groups are given in Section 4.2.2, though it is noted that the baseline ethnicity is the Hausa group.

A state's overall VHI value is calculated by taking the median average across individuals' VHI in each state. Clear spatial trends are observed (Figure 4.4). North-western states (with the exception of Zamfara) have lower VHI values (higher vaccine hesitancy) than the rest of the country (except Ekiti in the south west). VHI values in east Nigeria are very high signifying low concerns with vaccines; although VHI values are high across south Nigeria, there are low state-level VHI values in Rivers, Ondo, and Ekiti. These trends mirror historic adverse vaccine events in Nigeria: these have already been discussed in Section 1.2.1, and are

⁷VHI values from $\mathcal{M}_2$ could equally be used as the correlation between VHI values across these two models is high, $R^2 = 0.99$.

⁸This regression method is the ‘two-step’ approach discussed in Section 4.3.4. Although method can lead to biased parameter estimates as VHI values are assumed to be point estimates, this approach – rather than creating a three-level hierarchical model – is undertaken for its simplicity.




Figure 4.4: **Spatial vaccine hesitancy trends across Nigeria reveal high hesitancy in the north-Western states and in Ekiti** Median vaccine hesitancy (obtained by taking the median average of all individuals' inferred VHI within each state) is shown for each of Nigeria's 37 states. Vaccine hesitancy is high in the north west: Kano, in particular, has strikingly high hesitancy. Ekiti, in south-western Nigeria, has high hesitancy in contrast to its neighbouring states.

further commented on in the discussion section of this chapter.

The results of the hierarchical regression in Equation (4.24) are displayed in Figure 4.5 A. The fixed-effect terms are shown as black dots with 95% HPD intervals; the variation in random-effects across states are shown behind in grey.

Males in Nigeria are more likely than females to have lower VHI values (that is, they are more hesitant towards vaccines) and there is a strong interaction effect between ethnicity and sex. The precise nature of this relationship is somewhat obscured in Figure 4.5 A so in Figure 4.5 B, the predicted distributions of males and females within each ethnicity are displayed (all the other explanatory variables are held in their baseline groups). Within each ethnicity, females are more likely to have higher VHI values than males and this interaction effect is pronounced for the Hausa and Fulani/Peulh/Haalpulaar (FPH) ethnicities. Amongst the FPH ethnicities, VHI values are low and there is a striking difference between males and females. After controlling for a large range of factors, there is no strong association between




Figure 4.5: **A hierarchical regression with an extended set of explanatory variables uncovers a strong relationship between vaccine hesitancy and ethnicity and sex** (A) Fixed- (black dots) and random-effect (grey dots) parameters of the hierarchical regression in Equation (4.24) with 95% HPD intervals for the fixed effects shown. Individuals belonging to the FPH ethnicities who are religiously schooled are likely to be more vaccine hesitant than those who are not religiously schooled. (B) The predicted distributions of hesitancy for different ethnicities and the two sexes (whilst holding other variables constant) highlights the strong interaction effect between ethnicity and sex, with large differences in vaccine hesitancy between the sexes in the amongst the FPH and Hausa ethnicities.

age and VHI, though there is some evidence to suggest that the 65+ age group has higher VHI values than the baseline group. The strength of the age effect observed here is not as strong as the models in the preceding section, which could be due to the larger number of controlling factors. There is no evidence to support an interaction term between sex and education, although larger samples within each state may yet reveal a sex-education trend.

Although religious schooling is not generally associated with low VHI values (after controlling for other variables), those with religious upbringing in the FPH ethnicity higher VHI values (lower vaccine hesitancy) than those without. This result is striking: the Fulani ethnicity is a large ethnic group in northern Nigeria and this result may suggest that religious schooling may be a key contributing factor in individuals' development of hesitancy towards vaccines.

4.6 Discussion

In this chapter, we derive a quantitative index of vaccine hesitancy according to SAGE's '3Cs' definition [58]. Expressing hesitancy in this way allows a rigorous understanding of many properties of vaccine hesitancy, for example: spatial trends in hesitancy can be established and – when vaccine attitude data is routinely collected – we can quantitatively assess the spatio-temporal dynamics of hesitancy, including how and why hesitancy spreads and predicting future levels. These analyses will better equip us to anticipate where and why gaps in vaccination coverage will occur.

The vaccine hesitancy index is derived via Bayesian hierarchical item response theory: low index values highlight hesitancy concerns, whilst high index values suggest strong trust in vaccines and thus low vaccine hesitancy. This method simultaneously infers the VHI (our marker for hesitancy) and allows for its regression against socio-economic determinants. We use this index to investigate hesitancy globally, in the Sahel, and at the regional level in Nigeria. Vaccine hesitancy is low overall in the Sahel (compared to the rest of the world), with Mauritania the least hesitant country and Burkina Faso the most hesitant country in the Sahel. Despite the high relative hesitancy of Burkina Faso to the rest of Sahel, hesitancy is still low here compared to most other countries around the world (Figure 4.2). Similarly, Nigeria performs well globally, but has a wide variability in sentiment. Northern Nigerian states and Ekiti have very high hesitancy.

An investigation of the contribution to each survey question to vaccine hesitancy reveals religious compatibility has a smaller effect (that is, a smaller *discrimination*) in determining the rate of change of positive survey response with hesitancy than importance, safety, or effectiveness (Figure 4.2). This difference implies that religious compatibility concerns were

not as important as importance, safety, or effectiveness concerns in determining the latent hesitancy variable. The *difficulty* parameters – which inform the degree of hesitancy expressed for each survey question – reveal that vaccine safety is answered less positively by more individuals than effectiveness (in turn, effectiveness is answered less positively than vaccine importance). A vaccine hesitancy value in excess of 0.5 is associated with **tend to agree** or **strongly agree** responses for vaccine importance, safety, and effectiveness; however, there are moderate religious compatibility concerns around this value.

Variation in VHI values, given by Equation (4.22), is explained via a number of competing structural equation models to seek a parsimonious way of explaining variance at the individual and country levels. Models $\mathcal{M}_2$ and $\mathcal{M}_4$ are favoured. These two models explain the large variability at the individual-level; model $\mathcal{M}_4$ also captures the large unexplained variance amongst intercept terms (that is, model $\mathcal{M}_4$ models the difference amongst baseline groups between each country). Being female, a housewife, belonging to an older age group, and having a high level of education are associated with lower vaccine hesitancy values. Although age was not found to be significant across each individual survey question in Section 3.4, it is likely that income dominated the effect of age.

An investigation of country-level trends reveals a striking relationship between education and vaccine hesitancy in Burkina Faso, Mali, Niger, and Senegal (Figure 4.3 A-B). In these four countries, individuals with the highest education levels (undergraduate and postgraduate degrees) are significantly more likely to be hesitant than those with low education levels (primary education). However, globally, individuals with higher education levels are less likely to be vaccine hesitant (Table 4.1). In addition to the association between education and hesitancy, we also find that housewives are less likely to be vaccine-hesitant in Burkina Faso, Mali, Niger, and Senegal (Figure 4.3 A).

Individuals' VHI values in Nigeria are summarised and mapped in Figure 4.4. A high degree of vaccine hesitancy (that is, low VHI values) is found in north west Nigeria with the lowest rates in Kano and neighbouring states; hesitancy is low in the rest of Nigeria. These hesitancy values coincide with vaccine-acceptance issues in Kano, Zamfara, and Katsina due to drug-trial deaths and anti-vaccination messages from religious leaders [268,269]. The polio vaccine was suspended in Kano, and similar bans were implemented in Bauchi, Kaduna, and Zamfara [269]. Our inferred levels of hesitancy indeed indicate that Kano and Katsina have amongst the highest hesitancy levels in Nigeria, but also that individuals in Zamfara are less hesitant. Despite the ban in Bauchi, individuals do not seem to have acquired the same high levels of hesitancy as in Kano and Katsina. Our results reveal the extent of vaccine hesitancy in Nigeria and quantify the impact of *Boko Haram*, whose presence has led to the murder of nine polio workers in 2013 and has raised fears over the safety of health-seeking

behaviours [270].

A regression of individuals' VHI values against an extended set of socio-economic variables (Section 4.5.4) reveals the importance of the interactions between sex and ethnicity, and ethnicity and religion in determining levels of vaccine hesitancy. Males are more likely to be hesitant than females, and this effect is particularly strong amongst the Hausa and Fulani/Peulh/Haalpulaar (FPH) ethnicities (Figure 4.5 B). The Igbo, Yoruba, and other ethnic groups are less likely to be hesitant than the FPH ethnicities. Of those in the FPH grouping, religious schooling effects vaccine hesitancy: those without such schooling are more likely to be vaccine hesitant. Hausa and Fulani ethnicities – who are the main ethnic groups in northern Nigeria – have previously shown high levels of vaccine scepticism towards the polio vaccine [240].

A key limitation of the data used in this chapter is that only four dimensions on attitudes are surveyed; precluding, for example, an investigation of the extent to which intent to vaccinate affects hesitancy. Additional survey questions could yield further insights into the quantitative determinants of hesitancy and can be naturally incorporated into the item response theory modelling framework. Indeed, an item response theory analysis on a large set of variables putatively linked to hesitancy can reveal the most informative survey questions: this could have cost-saving consequences for vaccine hesitancy monitors.

A key assumption in hierarchical item response theory is that of *measurement invariance*. Measurement invariance assumes that the measurement instrument (in our case the four-question vaccine attitude survey) exhibits cross-national equivalence [246]. In the context of vaccine hesitancy, this assumption is required so that two individuals from different countries with the same overall hesitancy reply in the same way to the survey questions. The absence of standardising questions (for example on known behaviours of hesitancy such as previous vaccine refusal) complicates our ability to verify this assumption, though there exist methods to validate this assumption which should be applied [271].

In Chapter 5, in-depth descriptions of future analyses methods are provided and discussed in terms of their relevance for public health research. We particularly stress the impact of time-varying hesitancy models.

5 | Thesis summary and specific future research avenues

Talk about events (7)

— Philistine, *Guardian*
crossword No. 26,815

5.1 Summary

Over the past few decades, there has been a shift in the primary barrier to vaccination uptake from access to vaccines to vaccine hesitancy [238]. Although access is still the primary barrier for millions of individuals, there have been significant increases in the number of parents who are delaying or rejecting vaccines for their children in areas where vaccine are freely available [272, 273]. Vaccine hesitancy is curtailing progress in immunisation programmes and addressing hesitancy-based coverage gaps is a public health priority [153, 154].

In this thesis a number of tools in applied mathematics, statistics, and machine learning are used to enhance our understanding of the current global state of vaccine coverage and hesitancy. In Chapter 2, we introduce the first national-level forecasts of vaccine coverage via Gaussian process regression. From forecasts of coverage time-series we derive a vaccine performance index: a measure of the ability of national immunisation programmes to maintain and sustain high coverage. These forecasts and index are used – and can henceforth be used – to evaluate the progress of immunisation goals such as those from WHO’s Global Vaccine Action Plan [142]. We perform a large-scale correlative study into the strength of association between coverage and socio-economic determinants. This is the first study to perform such a broad analysis of the factors putatively linked to coverage. We identify several factors which possibly limit vaccine coverage rates and find a declining strength of association between coverage and socio-economic statistics in many regions. We speculate this decline could be indicative of this shift away from access as the primary barrier.

In Chapter 3 we use data collected from the largest survey on vaccine attitudes to date to establish global perceptions on the importance, safety, and effectiveness and identify individuals who have religious compatibility issues. Countries are ranked according to their overall level of sentiment towards each survey question and global attitudes to the safety of vaccines are mapped to reveal the countries with safety-based concerns. A regression study into the factors associated with these four attitudes uncovers individual- and country-level determinants of vaccine attitudes.

In Chapter 4 we derive a measure of individual vaccine hesitancy. This index summarises an individual’s overall vaccine views, and we propose that this measure can be interpreted as vaccine hesitancy. This is the first attempt to quantitatively infer vaccine hesitancy according to SAGE’s ‘3C’ definition [58]. We note that further questions relating to vaccination confidence, complacency, and convenience – in particular those which relate to vaccination intent – could reveal a more nuanced view of hesitancy. Nonetheless we have laid the groundwork for further studies into the spatio-temporal properties of this important measure with a sensible approach that can naturally incorporate further survey items.

The analyses in this thesis reveal valuable public health insights and build a platform upon which research in a number of key areas can be developed. For example, Gaussian process regression can be readily applied to sub-national immunisation rates. In the USA and UK, for instance, the Centers for Disease Control [274] and Public Health England [275] collect sub-national level immunisation data and WHO have recently published sub-national immunisation data collected from other national health bodies [27]. In these settings, we can begin to investigate regional goals, such as GVAP’s target of 80% coverage at every sub-national district [142].

Models with higher predictive power can be sought by, for example, including SEFs and auxiliary variables such as media sentiment. Social media has already been used to predict flu epidemics [276], localise vaccination scares [277], and correlate sentiment with vaccine coverage [278]. A further pressing analysis is to update our forecasts of vaccine coverage using the latest coverage data and to compare the resulting VPIs with the national summaries of vaccine hesitancy obtained in Chapter 4. This comparison will establish the utility of Gaussian process forecasts in determining hesitancy, and can be used as a proxy for hesitancy when survey data is not available.

In addition to analyses such as those outlined above, there are a number of novel mathematical modelling approaches that can yield greater insight into vaccine hesitancy and its dynamics. These methods can provide answers to challenges 4 and 5 outlined in the Section 1.4.1. These proposed methods are:

- **Hesitancy inference using non-survey data sources;**

- **Modelling the hesitancy-coverage link;** and
- **Identifying like-minded behavioural groups.**

The first two of these proposed analyses are naturally linked: inference of historic hesitancy levels enables an examination of the quantitative link between hesitancy and coverage. The third proposed point involves modelling hesitancy as a discrete – rather than continuous (as in Chapter 4) – latent variable. This method has the virtue that individuals clustered into discrete groups might yield information on shared vaccination beliefs.

5.2 Specific future research

In the following two sections (5.2.1 and 5.2.2) I outline specific research directions which follow from my thesis work and that I plan to pursue.

5.2.1 Hesitancy inference using non-survey data sources

In the preceding chapters, we sought to gauge global, national, and – in the case of Nigeria in Chapter 4 – sub-national levels of vaccine hesitancy. Although the latent variable approach in Chapter 4 offers the most direct way of obtaining vaccine hesitancy, historic levels are unobtainable via this method as relevant survey data does not exist. In this section, I outline how we can infer historic (and thus predict future) levels of vaccine hesitancy in the total absence of survey data probing attitudes to vaccines. If successful, this method has profound consequences for the quantification of past hesitancy, enabling it to be linked to socio-economic factors and other auxiliary data such as anti-vaccination media campaigns and health policies in a diverse range of settings (we estimate that we can achieve this in dozens of LMICs).

We propose decomposing vaccine coverage into attitude and hesitancy which – as they are unobserved – are to be modelled as latent variables. By doing this, we are able to combine a range of factors linked to healthcare access to gauge an individual’s overall access to immunisations. We suggest, then, that any residual variance in immunisation data after controlling for access is vaccine hesitancy (see model below). We will use data from the Demographic and Health Surveys (DHS) [279] to develop novel item response theory models to achieve the goals set out above. DHS is a program collecting childhood immunisation data and socioeconomic data (such as wealth, access to healthcare, socioeconomic status, religion, to name a few) for individuals in about 90 LMICs. These data are collected at the district levels within these countries.

Consider the following DHS variables:

- $\mathbf{Y}_i$: a vector of individual i 's immunisation records for the vaccinations BCG, DTP1, DTP2, DTP3, MCV, POL1, POL2, and POL3, where $Y_{im} \in \{0, 1\}$ for $m = 1, \dots, 8$ and where 1 (0) denotes presence (absence) of a given vaccination;¹
- $\mathbf{A}_i$: a vector of discrete-valued data representing the amount of healthcare individual i has received towards a certain healthcare-access-related variable. In the DHS dataset there are multiple variables informative of an individual's access to healthcare: for example, birth location of child, number of pre/post-natal check-ups, and distance to healthcare facility, to name a few (we estimate there to be about 20 healthcare-access-related variables in total). $A_{ip} \in \{0, K_p\}$ for $p = 1, \dots, N_p$, where N_p is the number of access-related variables, each with differing number of possible levels K_p ;
- ξ_i : healthcare *access* (a proxy for access to immunisations) for individual i (a latent variable);
- η_i : vaccine *hesitancy* for individual i (a latent variable); and
- $\mathbf{X}_i$: a vector of hesitancy-explaining variables.

We therefore seek to develop a model that simultaneously: infers vaccine access from an individual's healthcare variables; infers hesitancy given an individual's access; and – like our model in Chapter 4 – explains variance in hesitancy with variables putatively linked to attitude. The primary motivating assumption in this model (as stated above) is that vaccine uptake is determined *solely* as a combination of an individual's *access* and *hesitancy*. This assumption is based on SAGE's definition of vaccine hesitancy as non-access barriers to vaccination uptake [58]. A further important model assumption is that healthcare access is used as proxy to immunisation access: while likely a good approximation, we will need to substantiate this claim.

We therefore look to use the range of demographic factors to disentangle the effects of access and hesitancy in determining an individual's vaccination uptake behaviour. We here provide an outline of this novel approach for a group of individuals, but we note that it can easily be generalised to account for nested data structures, such as individuals within districts or states.

We begin by noting that the joint distribution of the variables enumerated above is proportional to the posterior distribution over the parameters of interest,

$$P(\eta_i, \xi_i, \boldsymbol{\theta} | \mathbf{Y}_i, \mathbf{A}_i) \propto P(\mathbf{Y}_i, \mathbf{A}_i, \eta_i, \xi_i, \boldsymbol{\theta}),$$

¹DTP2, POL1, and POL2 are, respectively, the second dose of the DTP programme, and the first and second doses of the polio programme.

where $\boldsymbol{\theta}$ is some parameter set. The joint distribution can be expressed as,

$$P(\mathbf{Y}_i, \mathbf{A}_i, \eta_i, \xi_i, \boldsymbol{\theta}) = P(\mathbf{Y}_i|\eta_i, \xi_i, \boldsymbol{\theta}_1)P(\mathbf{A}_i|\xi_i, \boldsymbol{\theta}_2)P(\xi_i)P(\eta_i)p(\boldsymbol{\theta}_1)p(\boldsymbol{\theta}_2),$$

where $\boldsymbol{\theta}$ is now decomposed into separate parameter sets $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$. $\boldsymbol{\theta}_1$ will contain item response parameters linking hesitancy and access to immunisation data, whilst $\boldsymbol{\theta}_2$ will contain item response parameters linking access to healthcare variables. The first probability distribution on the right-hand side models the relationship between immunisation data and vaccine hesitancy and vaccine access (both to be treated as latent variables), whilst the second uses healthcare to access variables to infer latent vaccine access. Thus, item response theory models can be used in both of these distributions and our parameter set $\boldsymbol{\theta}$ contains item response parameters relating latent variables to observed data.

Individual-level access latent variables can be treated a nuisance parameter (as our primary focus is η_i rather than ξ_i). Access can therefore be marginalised out of the joint distribution², so that the final model can be written as,

$$P(\eta_i, \boldsymbol{\theta}|\mathbf{Y}_i, \mathbf{A}_i) \propto \int P(\mathbf{Y}_i|\eta_i, \xi_i, \boldsymbol{\theta}_1)P(\mathbf{A}_i|\xi_i, \boldsymbol{\theta}_2)P(\xi_i)d\xi_i \times P(\eta_i)p(\boldsymbol{\theta}_1)p(\boldsymbol{\theta}_2). \quad (5.1)$$

It is left to specify models for $P(\mathbf{Y}_i|\eta_i, \xi_i, \boldsymbol{\theta}_1)$ and $P(\mathbf{A}_i|\xi_i, \boldsymbol{\theta}_2)$ (and prior distributions – but we do not consider these further). We propose that $P(\mathbf{Y}_i|\eta_i, \xi_i, \boldsymbol{\theta}_1)$ could be modelled as,

$$P(\mathbf{Y}_i|\eta_i, \xi_i, \boldsymbol{\theta}_1) \sim \text{Bern}(P(\mathbf{Y}_i = 1|\eta_i, \xi_i, \boldsymbol{\theta}_1)),$$

where

$$\text{logit}(P(\mathbf{Y}_i = 1|\eta_i, \xi_i, \boldsymbol{\theta}_1)) = \alpha\eta_i + \beta\xi_i,$$

that is, analogously to Equation (4.1), but with immunisation records depending on a linear combination of access and hesitancy. Similarly, an item response model can be used for $P(\mathbf{A}_i|\xi_i, \boldsymbol{\theta}_2)$ so that access can be both derived and item response parameters can be used to establish the health variables strongly linked to access.

Covariates of hesitancy can be introduced by simply replacing $\eta_i \rightarrow \eta_i|\mathbf{X}_i$ in the prior distribution for hesitancy in Equation (5.1) and using a structural equation model to capture this dependency (as we did in Section 4.3.4).

There are three main modelling challenges in this analysis. Firstly, immunisation data is often collected for children who are not yet eligible to receive further doses in a vaccination programme (which constitutes a missing data hurdle); secondly (and related to the first

²We could of course return these latent variables, but information about the amount access variable p relates to ξ_i is represented through item response parameters given in the parameter set $\boldsymbol{\theta}$.




Figure 5.1: **Violation of the conditional independence assumption for immunisation records** Conditional independence typically assumed in item response theory models (see Section 4.3.3) will be violated for individuals’ immunisation records where completion of later doses is conditional on receiving earlier doses such as the DTP and polio immunisation programmes.

modelling challenge), conditional independence will be violated for immunisation response data since requiring latter doses is dependent on former completion (see Figure 5.1); and, thirdly, time-variation in the model proposed above is required to provide temporal trends and forecasts.

Modelling the hesitancy-coverage link

The goal of the research outlined above is to obtain a time-varying, state-level hesitancy quantity $\eta_j(t)$ (say, for state j). Once this quantity has been derived, it can be linked to state-level vaccine coverage $y_j(t)$ (which is known) and we can attempt to model the interaction between these two variables. We could, for example, develop models to infer local hesitancy dynamics. That is, we could investigate the relationship $\eta_j(t) \rightarrow \eta_{j'}(t')$ and the impact on coverage rates $y_j(t) \rightarrow y_j(t')$ (for states j, j' and at times t and t'). This analysis could lead to valuable lessons on the spread of hesitancy from sceptical clusters of individuals, for instance, by investigating optimal vaccination interventions based on inferences made about hesitancy spread and future coverage levels or the introduction of local measures of herd immunity, based on local coverage-hesitancy dynamics.

We note that there have been many attempts to model how individual behaviour modulates vaccine uptake (see [280, 281] for overviews), including recent attempts to use social media to capture vaccine sentiment and model measles uptake [277]. We hope that these models could be furthered with our historic hesitancy inference or through survey-based hesitancy measures introduced in Chapter 4.

5.2.2 Identifying like-minded behavioural groups

In Chapter 2, item response theory is used to model vaccine hesitancy (as a continuous variable) by which individuals (and countries) can be ranked. Although providing useful insights into the relative level of hesitancy of each individual, modelling hesitancy as a continuous trait may not be the optimal way to capture hesitancy. A complementary analysis, which could provide an insight into how similar beliefs cluster, is to treat hesitancy as a *discrete* latent variable. Modelling hesitancy in this way would allow individuals to be grouped according to their ‘like-minded’ perceptions and behaviours. *Latent class analysis* [282] is the discrete analogue to item response theory: we briefly outline this method here.

Latent class analysis

Latent class analysis (LCA) concerns the inference of discrete latent variables (latent classes) from discrete multivariate data within a mixture model framework [283]. The goal of this analysis is to classify individuals into discrete classes according to their multivariate vaccination response data. Methods exist for inference of both ordinal [284] and nominal latent classes [282]: we focus on the latter, which breaks the ordinality of hesitancy enforced by the item response theory in Chapter 2.

Arranging the multivariate response data in the same way as in Section 4.3.3 (that is, with $\mathbf{Y} = [\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_N]^T$ as a matrix of individual items across all items), the likelihood $P(\mathbf{Y}|\boldsymbol{\theta})$ is again sought. In this modelling framework, the likelihood is given by,

$$P(\mathbf{Y}|\boldsymbol{\theta}, \boldsymbol{\pi}) = \prod_{i=1}^N \sum_{c=1}^{N_c} \prod_{m=1}^M \theta_{mc, Y_{im}} \pi_c, \quad (5.2)$$

where the parameters of interest are $\boldsymbol{\pi}$ – which give the probability that any randomly chosen individual belongs to a latent class $c = 1, \dots, C$ – and $\boldsymbol{\theta}_{mck}$, which is the probability of a specific item response given individual i ’s membership to a latent class C_i ,

$$\theta_{mck} = P(Y_{im} = k | C_i = c), \quad (5.3)$$

for $k = 1, \dots, K$ (the number of polytomous responses), $m = 1, \dots, M$ (the number of items), and $c = 1, \dots, N_C$, where N_C is the number of latent classes [285]. The posterior distribution may be again be specified through the inclusion of priors for $\boldsymbol{\theta}$ and $\boldsymbol{\pi}$, which we do not consider in detail here, but note that uninformative Dirichlet distributions would likely be an appropriate prior choice. Inferences can then be made about the class assignment probabilities and class membership for each individual. Dependency on covariates can be introduced [285] and hierarchical generalisations can be made in a similar way to IRT.

If fundamental characteristics of latent classes are determined, then opinion-modulating intervention strategies could be designed and developed for each type of hesitant group. This would provide novel insight into the nature of hesitancy and mechanisms to address it.

5.3 Final Thoughts

In this thesis we have considered global analyses of vaccine hesitancy and vaccine coverage using a variety of large datasets. Although many of the key challenges outlined in [Chapter 2](#) have been successfully met, we have noted several clearly delineated new research methods. Included in these methods is the development of novel models to infer historic levels of vaccine hesitancy. This research avenue is of particular importance as it permits a historical assessment of hesitancy with respect to vaccine coverage, media sentiment, and national (and sub-national) policy programmes.

Bibliography

- [1] World Health Organization *et. al.* Measles vaccination has saved an estimated 17.1 million lives since 2000 (news release). 2015. [Online; accessed 27-Dec-2017].
- [2] Joel G Breman, Isao Arita, Smallpox Eradication Unit, and World Health Organization *et. al.* The confirmation and maintenance of smallpox eradication. 1980.
- [3] Walter A Orenstein and Rafi Ahmed. Simply put: Vaccination saves lives. *Proceedings of the National Academy of Sciences*, 114(16):4031–4033, 2017.
- [4] Jeffrey P Baker and Samuel L Katz. Childhood vaccine development: an overview. *Pediatric research*, 55(2):347–356, 2004.
- [5] Peter Andersen and T Mark Doherty. The success and failure of BCG – implications for a novel tuberculosis vaccine. *Nature Reviews Microbiology*, 3(8):656–662, 2005.
- [6] Gregory J Fox, Marianna Orlova, and Erwin Schurr. Tuberculosis in newborns: The lessons of the “Lübeck disaster” (1929–1933). *PLoS pathogens*, 12(1):e1005271, 2016.
- [7] TM Daniel. Leon Charles Albert Calmette and BCG vaccine [founders of our knowledge]. *The International Journal of Tuberculosis and Lung Disease*, 9(9):944–945, 2005.
- [8] Stanley A Plotkin. Vaccines: past, present and future. *Nature medicine*, 11:S5–S11, 2005.
- [9] Jennifer Hamborsky, Andrew Kroger, and Charles Wolfe. *Epidemiology and Prevention of Vaccine-preventable Diseases: The Pink Book: Course Textbook*. Public Health Foundation, 2015.
- [10] World Health Organization *et. al.* WHO vaccine-preventable diseases: monitoring system: 2015 global summary. 2015. 6, 2015.
- [11] Rotem Lapidot and Christopher J Gill. The pertussis resurgence: putting together the pieces of the puzzle. *Tropical Diseases, Travel Medicine and Vaccines*, 2(1):26, 2016.
- [12] James D Cherry. Pertussis: challenges today and for the future. *PLoS pathogens*, 9(7):e1003418, 2013.
- [13] James D Cherry. Epidemic pertussis in 2012—the resurgence of a vaccine-preventable disease. *New England Journal of Medicine*, 367(9):785–787, 2012.
- [14] Victoria Romanus, Ragnar Jonsell, and Seth-olof Bergquist. Pertussis in sweden after the cessation of general immunization in 1979. *The Pediatric infectious disease journal*, 6(4):364–371, 1987.
- [15] Stanley A Plotkin. The pertussis problem. *Clinical infectious diseases*, 58(6):830–833, 2013.
- [16] Stuart Blume and Ingrid Geesink. A brief history of polio vaccines. *Science*, 288(5471):1593–1594, 2000.

- [17] Lorraine Ni, Jane F Seward, Tammy A Santibanez, Mark A Pallansch, Olen M Kew, D Rebecca Prevots, Peter M Strebel, Joanne Cono, Melinda Wharton, and Walter A *et. al* Orenstein. Vaccine policy changes and epidemiology of poliomyelitis in the United States. *The Journal of the American Medical Association*, 292(14):1696–1701, 2004.
- [18] Ciro A de Quadros, Jon K Andrus, Jean-Marc Olive, Carlyle Guerra de Macedo, and Donald A Henderson. Polio eradication from the western hemisphere. *Annual review of public health*, 13(1):239–252, 1992.
- [19] Benjamin M Nkowane, Steven GF Wassilak, Walter A Orenstein, Kenneth J Bart, Lawrence B Schonberger, Alan R Hinman, and Olen M Kew. Vaccine-associated paralytic poliomyelitis: United States: 1973 through 1984. *The Journal of the American Medical Association*, 257(10):1335–1340, 1987.
- [20] T Jacob John. Vaccine-associated paralytic polio in India. *Bulletin of the World Health Organization*, 80(11):917–917, 2002.
- [21] Adam Finn and Frank Bell. Polio vaccine: is it time for a change? *Archives of disease in childhood*, 78(6):571–574, 1998.
- [22] William L Atkinson, Walter A Orenstein, and Saul Krugman. The resurgence of measles in the United States, 1989-1990. *Annual review of medicine*, 43(1):451–463, 1992.
- [23] Jan Hendriks and Stuart Blume. Measles vaccination before the measles-mumps-rubella vaccine. *American journal of public health*, 103(8):1393–1401, 2013.
- [24] Heikki Peltola, Annamari Patja, Pauli Leinikki, Martti Valle, Irja Davidkin, and Mikko Paunio. No evidence for measles, mumps, and rubella vaccine-associated inflammatory bowel disease or autism in a 14-year prospective study. *The Lancet*, 351(9112):1327–1328, 1998.
- [25] Andrew J Wakefield, Simon H Murch, Andrew Anthony, John Linnell, DM Casson, Mohsin Malik, Mark Berelowitz, Amar P Dhillon, Michael A Thomson, and Peter *et. al* Harvey. Retracted: Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children, 1998.
- [26] Justin Lewis and Tammy Speers. Misleading media reporting? The MMR story. *Nature Reviews Immunology*, 3(11):913–918, 2003.
- [27] WHO-UNICEF. WHO-UNICEF coverage estimates. http://apps.who.int/immunization_monitoring/globalsummary/timeseries/tscoveragemcv1.html, 2017. [Online; accessed 21-Dec-2017].
- [28] WHO-UNICEF. WHO-UNICEF disease incidence. http://apps.who.int/immunization_monitoring/globalsummary/timeseries/tsincidencediphtheria.html, 2017. [Online; accessed 21-Dec-2017].
- [29] Fiona Godlee. The fraud behind the MMR scare. *British Medical Journal*, 342(d.22), 2011.
- [30] Fiona Godlee, Jane Smith, and Harvey Marcovitch. Wakefield’s article linking MMR vaccine and autism was fraudulent. *British Medical Journal*, 342, 2011.
- [31] Liam Smeeth, Claire Cook, Eric Fombonne, Lisa Heavey, Laura C Rodrigues, Peter G Smith, and Andrew J Hall. MMR vaccination and pervasive developmental disorders: a case-control study. *The Lancet*, 364(9438):963–969, 2004.

- [32] Anjali Jain, Jaclyn Marshall, Ami Buikema, Tim Bancroft, Jonathan P Kelly, and Craig J Newschaffer. Autism occurrence by MMR vaccine status among US children with older siblings with and without autism. *The Journal of the American Medical Association*, 313(15):1534–1540, 2015.
- [33] Luke E Taylor, Amy L Swerdfeger, and Guy D Eslick. Vaccines are not associated with autism: an evidence-based meta-analysis of case-control and cohort studies. *Vaccine*, 32(29):3623–3629, 2014.
- [34] Nadja Durbach. ‘They might as well brand us’: Working-class resistance to compulsory vaccination in Victorian England. *Social History of Medicine*, 13(1):45–63, 2000.
- [35] Dorothy Porter and Roy Porter. The politics of prevention: anti-vaccinationism and public health in nineteenth-century England. *Medical history*, 32(3):231, 1988.
- [36] Heidi J Larson, Caitlin Jarrett, Elisabeth Eckersberger, David MD Smith, and Pauline Paterson. Understanding vaccine hesitancy around vaccines and vaccination from a global perspective: a systematic review of published literature, 2007–2012. *Vaccine*, 32(19):2150–2159, 2014.
- [37] Martin Fichman and Jennifer E Keelan. Resister’s logic: the anti-vaccination arguments of Alfred Russel Wallace and their role in the debates over compulsory vaccination in England, 1870–1907. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 38(3):585–607, 2007.
- [38] Robert M Wolfe and Lisa K Sharp. Anti-vaccinationists past and present. *BMJ: British Medical Journal*, 325(7361):430, 2002.
- [39] E Peter Hennock. Vaccination policy against smallpox, 1835–1914: a comparison of England with Prussia and Imperial Germany. *Social History of Medicine*, 11(1):49–71, 1998.
- [40] Small-pox and vaccination in germany. *BMJ*, 2(1957):25–28, 1998.
- [41] Frank Fenner, Donald Ainslie Henderson, Isao Arita, Zdenek Jezek, and Ivan D *et. al* Ladnyi. *Smallpox and its eradication*, volume 6. World Health Organization Geneva, 1988.
- [42] Historical Medical Library of the College of Physicians of Philadelphia. Death of the Vaccinator. Obtained via <https://www.historyofvaccines.org/content/death-vaccinator>. [Online; accessed 21-Dec-2017].
- [43] Stanley A Plotkin. Lessons learned concerning vaccine safety. *Vaccine*, 20:S16–S19, 2001.
- [44] Khabir Ahmad. Nigerian government investigates pfizer drug trial allegations. *Lancet*, 357(9250):129, 2001.
- [45] Ayodele Samuel Jegede. What led to the Nigerian boycott of the polio vaccination campaign? *PLoS medicine*, 4(3):e73, 2007.
- [46] World Health Organization. Who removes nigeria from polio-endemic list. <http://www.who.int/mediacentre/news/releases/2015/nigeria-polio/en/>, 2015. [Online; accessed 27-Dec-2017].
- [47] Sebastian Taylor, Mahmud Khan, Ado Muhammad, Okey Akpala, Marit van Strien, Chris Morry, Warren Feek, and Ellyn Ogden. Understanding vaccine hesitancy in polio eradication in northern Nigeria. *Vaccine*, 35(47):6438–6443, 2017.
- [48] Haider J Warraich. Religious opposition to polio vaccination. *Emerging infectious diseases*, 15(6):978, 2009.

- [49] Eugene J Gangarosa, AM Galazka, CR Wolfe, LM Phillips, E Miller, RT Chen, and RE Gangarosa. Impact of anti-vaccine movements on pertussis control: the untold story. *The Lancet*, 351(9099):356–361, 1998.
- [50] Julie Leask. Target the fence-sitters. *Nature*, 473(7348):443–445, 2011.
- [51] Stanley Plotkin, Jeffrey S Gerber, and Paul A Offit. Vaccines and autism: a tale of shifting hypotheses. *Clinical Infectious Diseases*, 48(4):456–461, 2009.
- [52] Michael Eisenstein. Public health: An injection of trust. *Nature*, 507(7490):S17–S19, 2014.
- [53] Ohid Yaqub, Sophie Castle-Clarke, Nick Sevdalis, and Joanna Chataway. Attitudes to vaccination: a critical review. *Social Science & Medicine*, 112:1–11, 2014.
- [54] R Henderson, K Oates, H Macdonald, and WC Smith. General practitioners’ concerns about childhood immunisation and suggestions for improving professional support and vaccine uptake. *Commun Dis Public Health*, 7(4):260–6, 2004.
- [55] Jacqui Wise. Largest group of children affected by measles outbreak in Wales is 10-18 year olds. *BMJ: British Medical Journal*, 346, 2013.
- [56] Jessica E Atwell, Josh Van Otterloo, Jennifer Zipprich, Kathleen Winter, Kathleen Harriman, Daniel A Salmon, Neal A Halsey, and Saad B Omer. Nonmedical vaccine exemptions and pertussis in california, 2010. *Pediatrics*, 132(4):624–630, 2013.
- [57] Ayodele Samuel Jegede. What led to the Nigerian boycott of the polio vaccination campaign? *PLoS medicine*, 4(3):e73, 2007.
- [58] World Health Organization. Report of the SAGE working group on vaccine hesitancy. http://www.who.int/immunization/sage/meetings/2014/october/1_Report_WORKING_GROUP_vaccine_hesitancy_final.pdf, 2014. [Online; accessed 06-Jun-2016].
- [59] David H Peters, Anu Garg, Gerry Bloom, Damian G Walker, William R Brieger, and M Hafizur Rahman. Poverty and access to health care in developing countries. *Annals of the New York Academy of Sciences*, 1136(1):161–171, 2008.
- [60] Bruce Y Lee and Leila A Haidari. The importance of vaccine supply chains to everyone in the vaccine world. *Vaccine*, 2017.
- [61] David Sarley, Mustafa Mahmud, Jide Idris, Modele Osunkiyesi, Onome Dibosa-Osador, Peter Okebukola, and Owens Wiwa. Transforming vaccines supply chains in Nigeria. *Vaccine*, 35(17):2167–2174, 2017.
- [62] Yarlini Balarajan, Selvaraj Selvaraj, and SV Subramanian. Health care and equity in India. *The Lancet*, 377(9764):505–515, 2011.
- [63] Fowad Murtaza, Tajammal Mustafa, and Rabia Awan. Determinants of nonimmunization of children under 5 years of age in Pakistan. *Journal of family & community medicine*, 23(1):32, 2016.
- [64] Babar T Shaikh and Juanita Hatcher. Health seeking behaviour and health service utilization in Pakistan: challenging the policy makers. *Journal of public health*, 27(1):49–54, 2004.

- [65] Ruhul Amin, Nirali M Shah, and Stan Becker. Socioeconomic factors differentiating maternal and child health-seeking behavior in rural Bangladesh: A cross-sectional analysis. *International journal for equity in health*, 9(1):9, 2010.
- [66] Lori K Handy, Stefania Maroudi, Maura Powell, Bakanuki Nfila, Charlotte Moser, Ingrid Japa, Ndibo Monyatsi, Elena Tzortzi, Ismini Kouzeli, and Anthony *et. al* Luberti. The impact of access to immunization information on vaccine acceptance in three countries. *PloS one*, 12(8):e0180759, 2017.
- [67] Zulfiqar A Bhutta, Johannes Sommerfeld, Zohra S Lassi, Rehana A Salam, and Jai K Das. Global burden, distribution, and interventions for infectious diseases of poverty. *Infectious diseases of poverty*, 3(1):21, 2014.
- [68] John Grange and Alimuddin Zumla. Tuberculosis and the poverty-disease cycle, 1999.
- [69] Jorge Alvar, Sergio Yactayo, and Caryn Bern. Leishmaniasis and poverty. *Trends in parasitology*, 22(12):552–557, 2006.
- [70] Roshan Bastani, Beth A Glenn, Jennifer Tsui, L Cindy Chang, Erica Marchand, Victoria M Taylor, and Rita Singhal. Understanding sub-optimal HPV vaccine uptake among ethnic minority girls. *Cancer Epidemiology and Prevention Biomarkers*, pages cebp–0267, 2011.
- [71] R Bruce Aylward, Daniela Porta, Lucia Fiore, Barbara Ridolfi, Patrizia Chierchini, and Francesco Forastiere. Unimmunized gypsy populations and implications for the eradication of poliomyelitis in Europe. *The Journal of infectious diseases*, 175(Supplement_1):S86–S88, 1997.
- [72] Allison T Walker, Philip J Smith, and Maureen *et. al* Kolasa. Reduction of racial/ethnic disparities in vaccination coverage, 1995–2011. *MMWR Surveill Summ*, 63(Suppl 1):7–12, 2014.
- [73] Kevin Fiscella, Peter Franks, Marthe R Gold, and Carolyn M Clancy. Inequality in quality: addressing socioeconomic, racial, and ethnic disparities in health care. *The Journal of the American Medical Association*, 283(19):2579–2584, 2000.
- [74] World Health Organization *et. al*. Global Vaccine Action Plan 2011-2020. *Global Vaccine Action Plan 2011-2020.*, 2013.
- [75] Jeffrey D Sachs. From millennium development goals to sustainable development goals. *The Lancet*, 379(9832):2206–2211, 2012.
- [76] Access to medicine foundation. Access to vaccines index 2017. <https://accesstovaccinesindex.org/media/atvi/2017-Access-to-Vaccines-Index.pdf>, 2018. [Online; accessed 10-Jan-2018].
- [77] Marty Makinen, Miloud Kaddar, Vivikka Mollidrem, and Lara Wilson. New vaccine adoption in lower-middle-income countries. *Health policy and planning*, 27(suppl_2):ii39–ii49, 2012.
- [78] BA Ushie, OA Fayehun, and DB Ugal. Trends and patterns of under-5 vaccination in Nigeria, 1990-2008: what manner of progress? *Child: Care, Health and Development*, 40(2):267–274, 2013.
- [79] Lisa M Calhoun, Anna M van Eijk, Kim A Lindblade, Frank O Odhiambo, Mark L Wilson, Elizabeth Winterbauer, Laurence Slutsker, and Mary J Hamel. Determinants and coverage of vaccination in children in western Kenya from a 2003 cross-sectional survey. *Am. J. Trop. Med. Hyg.*, 90(2):234–241, 2003.

- [80] Yu Hu, Yaping Chen, Jing Guo, Xuewen Tang, and Lingzhi Shen. Completeness and timeliness of vaccination and determinants for low and late uptake among young children in eastern China. *Human Vaccines & Immunotherapeutics*, 10(5):1408–1415, 2014.
- [81] Michiko Ueda, Naoki Kondo, Misato Takada, and Hideki Hashimoto. Maternal work conditions, socioeconomic and educational status, and vaccination of children: A community-based household survey in Japan. *Preventable Medicine*, 66:17–21, 2014.
- [82] Sonalde Desai and Soumya Alva. Maternal education and child health: Is there a strong causal relationship? *Demography*, 35(1):71–78, 1998.
- [83] M.C. Sullivan, A. Tegegn, F. Tessema, et. al. Minding the immunization gap: family characteristics associated with completion rates in rural ethiopia. *Journal of community health*, 35(1):53–59, 2010.
- [84] M. K. Akmatov, M. Kretzschmar, and A. Krämer, R. Mikolajczyk. Determinants of childhood vaccination coverage in kazakhstan in a period of societal change: implications for vaccination policies. *Vaccine*, 25(10):1756–1763, 2007.
- [85] I. A. Semali. Trends in immunization completion and disparities in the context of health reforms: the case study of Tanzania. *BMC health services research*, 10(1):299, 2010.
- [86] Daouda Sissoko, Helen Trottier, Denis Malvy, and Mira Johri. The influence of compositional and contextual factors on non-receipt of basic vaccines among children of 12-23-month old in India: A multilevel analysis. page e106528, 2014.
- [87] Suresh Sharma. Socioeconomic factors of full immunisation coverage in India. *World Journal of Vaccines*, 3(03):102, 2013.
- [88] Charles S Wiysonge, Olalekan A Uthman, Peter M Ndumbe, and Gregory D Hussey. Individual and contextual factors associated with low childhood immunisation coverage in sub-Saharan Africa: a multilevel analysis. *PLoS One*, 7.5:e37905, 2012.
- [89] Diddy Antai. Inequitable childhood immunization uptake in Nigeria: a multilevel analysis of individual and contextual determinants. *BMC infectious diseases*, 9(1):181, 2009.
- [90] Oyelola A Adegboye, Danelle Kotze, and Olasunkanmi A Adegboye. Multi-year trend analysis of childhood immunization uptake and coverage in Nigeria. *Journal of biosocial science*, 46(02):225–239, 2014.
- [91] Rohini P Pande and Abdo S Yazbeck. What’s in a country average? Wealth, gender and regional inequalities in immunization in India. *Social Science & Medicine*, 57(11):2075–2088, 2003.
- [92] Maureen E Canavan, Heather L Sipsma, Getnet M Kassie, and Elizabeth H Bradley. Correlates of complete childhood vaccination in East African countries. *PLoS One*, 9(4):e95709, 2014.
- [93] Cheryl A Moyer, Lia Tadesse, and Senait Fisseha. The relationship between facility delivery and infant immunization in Ethiopia. *International Journal of Gynecology & Obstetrics*, 123(3):217–220, 2013.
- [94] Anu Rammohan, Niyi Awofeso, and Renae C Fernandez. Paternal education status significantly influences infants’ measles vaccination uptake, independent of maternal education status. *BMC public health*, 12(1):336, 2012.

- [95] K Danis, T Georgakopoulou, T Stavrou, D Laggas, and T Panagiotopoulos. Socioeconomic factors play a more important role in childhood vaccination coverage than parental perceptions: a cross-sectional study in Greece. *Vaccine*, 28(7):1861–1869, 2010.
- [96] Karen S Wagner, Johan CJ van Wijgerden, Nick Andrews, Khushbu Goulden, and Joanne M White. Childhood vaccination coverage by ethnicity within London between 2006/2007 and 2010/2011. *Arch. Dis. Child*, 99:348–353, 2014.
- [97] Anna Pearce, Catherine Law, David Elliman, Tim J Cole, and Helen Bedford. Factors associated with uptake of measles, mumps, and rubella vaccine (MMR) and use of single antigen vaccines in a contemporary uk cohort: prospective cohort study. *Bmj*, 336(7647):754–757, 2008.
- [98] Rohini P Pande and Abdo S Yazbeck. What’s in a country average? wealth, gender, and regional inequalities in immunization in India. *Social Science & Medicine*, 57(11):2075–2088, 2003.
- [99] Daniel J Corsi, Diego G Bassani, Rajesh Kumar, Shally Awasthi, Raju Jotkar, Navkiran Kaur, and Prabhat Jha. Gender inequity and age-appropriate immunization coverage in India from 1992 to 2006. *BMC international health and human rights*, 9(1):S3, 2009.
- [100] Ayaz Muhammad. Relationship between child immunization and household socio-demographic characteristic in Pakistan. *Research on Humanities and Social Sciences*, 4(7), 2014.
- [101] Maureen E Canavan, Heather L Sipsma, Getnet M Kassie, and Elizabeth H Bradley. Correlates of complete childhood vaccination in East African countries. *PLoS One*, 9(4):e95709, 2014.
- [102] Karen S Wagner, Johan CJ van Wijgerden, Nick Andrews, Khushbu Goulden, and Joanne M White. Childhood vaccination coverage by ethnicity within London between 2006/2007 and 2010/2011. *Archives of disease in childhood*, 99(4):348–353, 2014.
- [103] Jeremy I Hawker, Babatunde Olowokure, Annette L Wood, Richard C Wilson, and Richard Johnson. Widening inequalities in MMR vaccine uptake rates among ethnic groups in an urban area of the UK during a period of vaccine controversy (1994–2000). *Vaccine*, 25(43):7516–7519, 2007.
- [104] BA Ushie, OA Fayehun, and DB Ugal. Trends and patterns of under-5 vaccination in nigeria, 1990–2008: what manner of progress? *Child: care, health and development*, 40(2):267–274, 2014.
- [105] Anu Rammohan and Niyi Awofeso. District-level variations in childhood immunizations in India: the role of socio-economic factors and health infrastructure. *Social Science & Medicine*, 145:163–172, 2015.
- [106] Jessica E Atwell and Daniel A Salmon. Pertussis resurgence and vaccine uptake: implications for reducing vaccine hesitancy. *Pediatrics*, 134(3):602–604, 2014.
- [107] Laura AV Marlow, Jo Waller, and Jane Wardle. Trust and experience as predictors of HPV vaccine acceptance. *Human vaccines*, 3(5):171–175, 2007.
- [108] Robert McMurray, Francine M Cheater, Anna Weighall, Carolyn Nelson, Martin Schweiger, and Suzanne Mukherjee. Managing controversy through consultation: a qualitative study of communication and trust around MMR vaccination decisions. *Br J Gen Pract*, 54(504):520–525, 2004.
- [109] Andrea L Benin, Daryl J Wisler-Scher, Eve Colson, Eugene D Shapiro, and Eric S Holmboe. Qualitative analysis of mothers’ decision-making about vaccines for infants: The importance of trust. *Pediatrics*, 117(5):1532–1541, 2006.

- [110] Koji Wada and Derek R Smith. Mistrust surrounding vaccination recommendations by the Japanese government: results from a national survey of working-age individuals. *BMC public health*, 15(1):426, 2015.
- [111] Willemien van der Weerd, Daniëlle RM Timmermans, Desirée JMA Beaujean, Jurriaan Oudhoff, and Jim E van Steenbergen. Monitoring the level of government trust, risk perception and intention of the general public to adopt protective measures during the influenza a (h1n1) pandemic in the netherlands. *BMC public health*, 11(1):575, 2011.
- [112] A. Benin, D.J. Wisler-Scher, E. Colson, E.D. Shapiro, and E.S. Holmboe. Qualitative analysis of mothers’ decision-making about vaccines for infants: The importance of trust. *Pediatrics*, 117(5):1532–1541, 2006.
- [113] Noni E MacDonald, Jennifer Smith, and Mary Appleton. Risk perception, risk management and safety assessment: what can governments do to increase public confidence in their vaccine system? *Biologicals*, 40(5):384–388, 2012.
- [114] E Hak, Y Schönbeck, H De Melker, GA Van Essen, and EAM Sanders. Negative attitude of highly educated parents and health care workers towards future vaccinations in the Dutch childhood vaccination program. *Vaccine*, 23(24):3103–3107, 2005.
- [115] Paolo Bonanni and Mauro Bergamini. Factors influencing vaccine uptake in Italy. *Vaccine*, 20:S8–S12, 2001.
- [116] P. Bonanni and M Bergamini. Factors influencing vaccine uptake in Italy. *Vaccine*, 20(1):S8–12, 2001.
- [117] World Health Organization. Immunization coverage factsheet. <http://www.who.int/mediacentre/factsheets/fs378/en/>, 2017. [Online; accessed 27-Dec-2017].
- [118] Mary E Ramsay. Measles: the legacy of low vaccine coverage. *Archives of disease in childhood*, 98(10):752–754, 2013.
- [119] Alan Smith, Joanne Yarwood, and David M Salisbury. Tracking mothers’ attitudes to MMR immunisation 1996–2006. *Vaccine*, 25(20):3996–4002, 2007.
- [120] Brendan W Mason and Peter D Donnelly. Impact of a local newspaper campaign on the uptake of the measles mumps and rubella vaccine. *Journal of Epidemiology & Community Health*, 54(6):473–474, 2000.
- [121] Michael McCarthy. Measles outbreak linked to disney theme parks reaches five states and Mexico. *BMJ: British Medical Journal (Online)*, 350, 2015.
- [122] Denise Antona, Daniel Lévy-Bruhl, Claire Baudon, François Freymuth, Mathieu Lamy, Catherine Maine, Daniel Floret, and Isabelle Parent Du Chatelet. Measles elimination efforts and 2008–2011 outbreak, France. *Emerging infectious diseases*, 19(3):357, 2013.
- [123] Anita WM Suijkerbuijk, Tom Woudenberg, Susan JM Hahné, Laura Nic Lochlainn, Hester E de Melker, Wilhelmina LM Ruijs, and Anna K Lugné. Economic costs of measles outbreak in the netherlands, 2013–2014. *Emerging infectious diseases*, 21(11):2067, 2015.
- [124] Mark Muscat, Henrik Bang, Jan Wohlfahrt, Steffen Glismann, and Kåre *et. al* Mølbak. Measles in Europe: an epidemiological assessment. *The Lancet*, 373(9661):383–389, 2009.

- [125] Public Health England. Measles notifications and deaths in England and Wales: 1940 to 2016. <https://www.gov.uk/government/publications/measles-deaths-by-age-group-from-1980-to-2013-ons-data/measles-notifications-and-deaths-in-england-and-wales-1940-to-2013>, 2016. [Online; accessed 21-Dec-2017].
- [126] Washington State Department of Health. News release. <https://www.doh.wa.gov/Newsroom/2015NewsReleases/15119WAMeaslesRelatedDeath>, 2015. [Online; accessed 21-Dec-2017].
- [127] Muhammad Umair Khan and Akram Ahmad. Polio vaccination in Pakistan. *The Lancet*, 386(9991):337, 2015.
- [128] Seye Abimbola, Asmat Ullah Malik, and Ghulam Farooq Mansoor. The final push for polio eradication: addressing the challenge of violence in Afghanistan, Pakistan, and Nigeria. *PLoS medicine*, 10(10):e1001529, 2013.
- [129] Jeffrey P Baker. The pertussis vaccine controversy in Great Britain, 1974–1986. *Vaccine*, 21(25):4003–4010, 2003.
- [130] Svetlana S Markina, Nina M Maksimova, Charles R Vitek, Erika Y Bogatyreva, and Anatoly A Monisov. Diphtheria in the Russian Federation in the 1990s. *Journal of Infectious Diseases*, 181(Supplement_1):S27–S34, 2000.
- [131] The Public Health Agency of Sweden. Barriers and motivating factors to MMR vaccination in communities with low coverage in Sweden. <https://www.folkhalsomyndigheten.se/contentassets/5db4b41a40f94e98b0e1d0d4a596bae8/barriers-motivating-factors-{M}-{R}-vaccination-communities-low-coverage-sweden-15027.pdf>, 2015. [Online; accessed 20-Dec-2017].
- [132] World Health Organization *et. al.* News release. <http://www.who.int/mediacentre/news/releases/2004/pr32/en/>, 2004. [Online; accessed 27-Dec-2017].
- [133] Qanta A Ahmed, Sania Nishtar, and Ziad A Memish. Poliomyelitis in Pakistan: time for the muslim world to step in. *The Lancet*, 381(9877):1521–1523, 2013.
- [134] Sara Reardon. Decrying CIA vaccination sham, health workers brace for backlash. *Science*, 333(6041):395–395, 2011.
- [135] Naoki Masuda and Petter Holme. Temporal network epidemiology, 2017.
- [136] Emilia Vynnycky and Richard White. *An introduction to infectious disease modelling*. Oxford University Press, 2010.
- [137] CDC state vaccination requirements. <https://www.cdc.gov/vaccines/imz-managers/laws/state-reqs.html>. [Online; accessed 10-January-2018].
- [138] Frank H Beard, Brynley P Hull, Julie Leask, Aditi Dey, and Peter B McIntyre. Trends and patterns in vaccination objection, Australia, 2002–2013. *The Medical Journal of Australia*, 204(7):275, 2016.
- [139] Editorial. Our responsibility to protect the Rohingya. *The Lancet*, 390(10114):2740, 2017.
- [140] World Health Organization. Disease outbreak news: Diphtheria – Cox’s bazar in Bangladesh. <http://www.who.int/csr/don/13-december-2017-diphtheria-bangladesh/en/>, 2017. [Online; accessed 27-Dec-2017].

- [141] World Health Organization. Epi fact sheet: Myanmar. <http://www.searo.who.int/immunization/data/myanmar.pdf>, 2017. [Online; accessed 20-Dec-2017].
- [142] Global Vaccine Action Plan 2011-2020. Annex 1., <http://www.unicef.org/immunization/files/GVAP.pdf>.
- [143] Gavi 2011-2015 indicators. <http://www.gavi.org/results/goal-level-indicators/health-systems-goal-indicators/>. [Online; accessed 10-Jan-2018].
- [144] The guide to tailoring immunization programmes (TIP). http://www.euro.who.int/__data/assets/pdf_file/0003/187347/The-Guide-to-Tailoring-Immunization-Programmes-TIP.pdf?ua=1. [Online; accessed 10-Jan-2018].
- [145] Public Health England. Tailoring immunisation programmes: Charedi community, north London. <https://www.gov.uk/government/publications/tailoring-immunisation-programmes-charedi-community-north-london>, 2016. [Online; accessed 27-Dec-2017].
- [146] V. Nankabirwa, T. Tylleskär, J.K. Tumwine, and H. Sommerfelt. Maternal education is associated with vaccination status of infants less than 6 months in eastern Uganda: a cohort study. *BMC Pediatrics*, 10:92–101, 2010.
- [147] Mira Johri, SV Subramanian, Marie-Pierre Sylvestre, Sakshi Dudeja, Dinesh Chandra, Georges K Koné, Jitendar K Sharma, and Smriti Pahwa. Association between maternal health literacy and child vaccination in India: a cross-sectional study. *J Epidemiol Community Health*, pages jech–2014, 2015.
- [148] Steven Black and Rino Rappuoli. A crisis of public confidence in vaccines. *Sci. Trans. Med.*, 2(61):61mr1, 2010.
- [149] E Hak, Y Schönbeck, H De Melker, GA Van Essen, and EAM Sanders. Negative attitude of highly educated parents and health care workers towards future vaccinations in the Dutch childhood vaccination program. *Vaccine*, 23(24):3103–3107, 2005.
- [150] Bola Murele, Rui Vaz, Alex Gasasira, Pascal Mkanda, Tesfaye Erbetto, and Joseph Okeibunor. Vaccine perception among acceptors and non-acceptors in Sokoto state, Nigeria. *Vaccine*, 32(26):3323–3327, 2014.
- [151] Brendan Nyhan, Jason Reifler, Sean Richey, and Gary L Freed. Effective messages in vaccine promotion: a randomized trial. *Pediatrics*, 133(4):e835–e842, 2014.
- [152] Elizabeth Wolf, Ali Rowhani-Rahbar, Jeffrey Duchin, M Patricia DeHart, and Douglas Opel. The challenges in measuring local immunization coverage: a statewide case study. *Pediatrics*, 137(5):e20153755, 2016.
- [153] S. Ozawa and M. L. Stack. Public trust and vaccine acceptance-international perspectives. *Human Vaccines and Immunotherapeutics*, 9(8):1774–1778, 2013.
- [154] H. Larson and W. Schulz. The state of vaccine confidence 2015. <http://www.vaccineconfidence.org/The-State-of-Vaccine-Confidence-2015.pdf>, 2016. [Online; accessed 29-Apr-2016].
- [155] P Carrillo-Santistevé and PL Lopalco. Measles still spreads in Europe: who is responsible for the failure to vaccinate? *Clinical Microbiology and Infection*, 18(s5):50–56, 2012.

- [156] U.S. Department of Health and NVAC Vaccine Confidence Working Group Human Services. Assessing the state of vaccine confidence in the United States: recommendations from the national vaccine advisory committee. <http://www.hhs.gov/sites/default/files/nvpo/nvac/reports/nvac-vaccine-confidence-public-health-report-2015.pdf>, 2016. [Online; accessed 23-Jun-2016].
- [157] Alexandre de Figueiredo, Iain G Johnston, David MD Smith, Sumeet Agarwal, Heidi J Larson, and Nick S Jones. Forecasted trends in vaccination coverage and correlations with socioeconomic factors: a global time-series analysis over 30 years. *The Lancet Global Health*, 4(10):e726–e735, 2016.
- [158] Michiel van Boven and Alies van Lier. Global trends in vaccination coverage. *The Lancet Global Health*, 4(10):e670–e671, 2016.
- [159] Heidi J Larson, Alexandre de Figueiredo, Zhao Xiahong, William S Schulz, Pierre Verger, Iain G Johnston, Alex R Cook, and Nick S Jones. The state of vaccine confidence 2016: global insights through a 67-country survey. *EBioMedicine*, 12:295–301, 2016.
- [160] Steven Black. Recognizing the importance of vaccine confidence. *EBioMedicine*, 12:28–29, 2016.
- [161] Anthony Burton, Roeland Monasch, Barbara Lautenbach, Marta Gacic-Dobo, Maryanne Neill, Rouslan Karimov, Lara Wolfson, Gareth Jones, and Maureen Birmingham. WHO and UNICEF estimates of national infant immunization coverage: methods and processes. *Bulletin of the World Health Organization*, 87(7):535–541, 2009.
- [162] Ingrid Torjesen. Uk adds hepatitis b to infant vaccination schedule, 2017.
- [163] GAVI. Mission indicators. Obtained via <http://www.gavi.org/results/measuring/2011-2015-indicators/>, 2011. [Online; accessed 29-Dec-2017].
- [164] Gapminder. www.gapminder.org/data/. [Data obtained from various sources via Gapminder].
- [165] World Bank Group. <http://www.worldbank.org>.
- [166] International Labor Organization. <http://www.ilo.org>.
- [167] World Health Organization. <http://www.who.int>.
- [168] Rebecca M Warner. *Applied statistics: From bivariate through multivariate techniques*. Sage, 2008.
- [169] World Health Organization. WHO regional offices. www.who.int/about/regions/en. [Online; accessed 08-Aug-2013].
- [170] A Rogier T Donders, Geert JMG van der Heijden, Theo Stijnen, and Karel GM Moons. A gentle introduction to imputation of missing values. *Journal of clinical epidemiology*, 59(10):1087–1091, 2006.
- [171] Craig K Enders. *Applied missing data analysis*. Guilford Press, 2010.
- [172] Andrew Gelman and Jennifer Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press New York, NY, USA, 2007.
- [173] Heikki Junninen, Harri Niska, Kari Tuppurainen, Juhani Ruuskanen, and Mikko Kolehmainen. Methods for imputation of missing values in air quality data sets. *Atmospheric Environment*, 38(18):2895–2907, 2004.

- [174] David Banks, Leanna House, Frederick R McMorris, Phipps Arabie, and Wolfgang A Gaul. *Classification, Clustering, and Data Mining Applications: Proceedings of the Meeting of the International Federation of Classification Societies (IFCS), Illinois Institute of Technology, Chicago, 15–18 July 2004*. Springer Science & Business Media, 2011.
- [175] Nicholas L Crookston and Andrew O Finley. yaimpute: an R package for kNN imputation. 2008.
- [176] Danh V Nguyen, Naisyin Wang, and Raymond J Carroll. Evaluation of missing value estimation for microarray data. *Journal of Data Science*, 2(4):347–370, 2004.
- [177] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [178] J Edward Jackson. *A user’s guide to principal components*, volume 587. John Wiley & Sons, 2005.
- [179] Trevor F Cox and Michael AA Cox. *Multidimensional scaling*. CRC press, 2000.
- [180] John W Sammon. A nonlinear mapping for data structure analysis. *IEEE Transactions on computers*, 100(5):401–409, 1969.
- [181] Geoffrey E Hinton and Sam T Roweis. Stochastic neighbor embedding. In *Advances in neural information processing systems*, pages 857–864, 2003.
- [182] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [183] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [184] J Bernardo, J Berger, A Dawid, and A *et. al* Smith. Regression and classification using Gaussian process priors. *Bayesian statistics*, 6:475, 1998.
- [185] Stephen Roberts, Michael Osborne, Mark Ebden, Steven Reece, Neale Gibson, and Suzanne Aigrain. Gaussian processes for time-series modelling. *Phil. Trans. R. Soc. A*, 371(1984):20110550, 2013.
- [186] Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (gpml) toolbox. *Journal of Machine Learning Research*, 11(Nov):3011–3015, 2010.
- [187] Gregory W Corder and Dale I Foreman. *Nonparametric statistics: A step-by-step approach*. John Wiley & Sons, 2014.
- [188] Ronald K Pearson. *Mining imperfect data: Dealing with contamination and incomplete records*. SIAM, 2005.
- [189] Jerome L Myers, Arnold Well, and Robert Frederick Lorch. *Research design and statistical analysis*. Routledge, 2010.
- [190] Larry Wasserman. *All of statistics: a concise course in statistical inference*. Springer Science & Business Media, 2013.
- [191] John H McDonald. *Handbook of biological statistics*, volume 2. Sparky House Publishing Baltimore, MD, 2009.
- [192] Robert Delongchamp, Taewon Lee, and Cruz Velasco. A method for computing the overall statistical significance of a treatment effect among a group of genes. *BMC bioinformatics*, 7(2):S11, 2006.

- [193] World Bank Data. Births attended by skilled health staff. <https://data.worldbank.org/indicator/SH.STA.BRTC.ZS>, 2018. [Online; accessed 10-Jan-2018].
- [194] World Bank Data. Improved sanitation facilities. <https://data.worldbank.org/indicator/SH.STA.ACSN>, 2018. [Online; accessed 10-Jan-2018].
- [195] World Bank Data. Improved water source. <https://data.worldbank.org/indicator/SH.H2O.SAFE.ZS?>, 2018. [Online; accessed 10-Jan-2018].
- [196] World Bank Data. Primary completion rate. <https://data.worldbank.org/indicator/SE.PRM.CMPT.ZS>, 2018. [Online; accessed 10-Jan-2018].
- [197] World Bank Data. Education enrolment ratio. <http://data.worldbank.org/indicator/SE.ENR.PRSC.FM.ZS>, 2018. [Online; accessed 10-Jan-2018].
- [198] WHO Global Health Expenditure Database. Central government health expenditure. <http://apps.who.int/nha/database/PreDataExplorer.aspx?d=1>, 2018. [Online; accessed 10-Jan-2018].
- [199] World Bank Data. Gross domestic product per capita. <http://data.worldbank.org/indicator/NY.GDP.PCAP.PP.KD>, 2018. [Online; accessed 10-Jan-2018].
- [200] World Bank Data. Gross national income per capita. <http://data.worldbank.org/indicator/NY.GNP.PCAP.PP.CD/countries>, 2018. [Online; accessed 10-Jan-2018].
- [201] World Bank Data. Phone subscribers. <http://go.worldbank.org/UOFSM7AQ40>, 2018. [Online; accessed 10-Jan-2018].
- [202] Carbon Dioxide Information Analysis Center. Carbon dioxide usage per person. http://cdiac.ornl.gov/trends/emis/meth_reg.html, 2018. [Online; accessed 10-Jan-2018].
- [203] World Bank data. Improved sanitation facilities access (rural). <http://data.worldbank.org/indicator/SH.STA.ACSN.RU>, 2018. [Online; accessed 10-Jan-2018].
- [204] WHO. World immunization week 2016: immunization game-changers should be the norm world-wide. <http://www.who.int/mediacentre/news/releases/2016/world-immunization-week/en/>, 2016. [Online; accessed 29-Apr-2016].
- [205] Márcia de Cantuária Tauil, Ana Paula Sayuri Sato, and Eliseu Alves Waldman. Factors associated with incomplete or delayed vaccination across countries: a systematic review. *Vaccine*, 34(24):2635–2643, 2016.
- [206] N.A. Otieno, B.O. Nyawanda, A. Audi, et al. Demographic, socio-economic and geographic determinants of seasonal influenza vaccine uptake in rural western Kenya, 2011. *Vaccine*, 32(49):6699–6704, 2014.
- [207] E.M. Galarce, S. Minsky, and K. Viswanath. Socioeconomic status, demographics, beliefs and A(H1N1) vaccine uptake in the United States. *Vaccine*, 29(32):5284–5289, 2011.
- [208] V. Friederichs, J.C. Cameron, and C. Robertson. Impact of adverse publicity on MMR vaccine uptake: a population based analysis of vaccine uptake records for one million children, born 1987-2004. *Arch. Dis. Child*, 91:465–468, 2006.
- [209] R. Reading, A. Colver, S. Openshaw, and S. Jarvis. Do interventions that improve immunisation uptake also reduce social inequalities in uptake? *BMJ*, 308(6379):1142–1144, 1994.

- [210] S. Mitchell, N. Andersson, N.M. Ansari, K. Omer, J.L. Soberanis, and A. Cockcroft. Equity and vaccine uptake: a cross-sectional study of measles vaccination in Lasbela district, Pakistan. *BMC Int. Health Hum. Rights*, 9(suppl. 1):S7, 2009.
- [211] O.T. Le Polain de Waroux, J.R. Armstrong Schellenberg, F. Manzi, et al. Timeliness and completeness of vaccination and risk factors for low and late vaccine uptake in young children living in rural southern Tanzania. *int. Health*, 5(2):139–147, 2013.
- [212] Public Health England. Parental attitudes to vaccination in young children. <https://www.gov.uk/government/collections/parental-attitudes-to-vaccination-in-young-children>, 2017. [Online collection; accessed 28-November-2017].
- [213] Eve Dubé, Dominique Gagnon, Emily Nickels, Stanley Jeram, and Melanie Schuster. Mapping vaccine hesitancy—country-specific characteristics of a global phenomenon. *Vaccine*, 32(49):6649–6654, 2014.
- [214] Communicable diseases network Australia. Human papillomavirus (hpv) surveillance plan – an integrated approach to monitoring the impact of hpv vaccine in Australia. <http://www.health.gov.au/hpv-surveillance-plan.pdf>, 2013. [Online; accessed 28-November-2017].
- [215] Eve Dubé, Dominique Gagnon, and Noni E *et. al* MacDonald. Strategies intended to address vaccine hesitancy: Review of published reviews. *Vaccine*, 33(34):4191–4203, 2015.
- [216] Heidi J Larson, William S Schulz, Joseph D Tucker, and David MD Smith. Measuring vaccine confidence: introducing a global vaccine confidence index. *PLoS currents*, 7, 2015.
- [217] WIN/Gallup International Association. End of Year Survey. http://www.wingia.com/en/survey/end_of_year_survey/, 2017. [Online; accessed 28-November-2017].
- [218] World Bank. Socio-economic data. [http://data.worldbank.org/indicator/X, X = NY.GDP.PCAP.KD, SH.STA.BRTC.ZS, SH.DYN.MORT, SH.STA.ACSN, or SH.H20.SAFE.RU.ZS](http://data.worldbank.org/indicator/X,X=NY.GDP.PCAP.KD,SH.STA.BRTC.ZS,SH.DYN.MORT,SH.STA.ACSN,orSH.H20.SAFE.RU.ZS), 2016. [Online; accessed 02-May-2016].
- [219] World Health Organization. Socio-economic data. http://apps.who.int/gho/indicatorregistry/App_Main/view_indicator.aspx?iid=122, 2016. [Online; accessed 02-May-2016].
- [220] United Nations Development Programme. Socio-economic data. <http://hdr.undp.org/en/content/mean-years-schooling-adults-years>, 2016. [Online; accessed 02-May-2016].
- [221] Tom Snijders and Roel Bosker. *Multilevel analysis: An introduction to basic and applied multilevel analysis (2nd Edition)*. London: Sage, 1999.
- [222] Andrew Gelman. Multilevel (hierarchical) modeling: what it can and cannot do. *Technometrics*, 48(3):432–435, 2006.
- [223] Clifford H Wagner. Simpson’s paradox in real life. *The American Statistician*, 36(1):46–48, 1982.
- [224] Stephen W Raudenbush and Anthony S Bryk. *Hierarchical linear models: Applications and data analysis methods*, volume 1. Sage, 2002.
- [225] WIN/GALLUP International Association. End of Year, survey methodology. http://www.wingia.com/web/files/richeditor/filemanager/Methosheet_Global.pdf, 2016. [Online; accessed 01-March-2016].

- [226] John Neter, Michael H Kutner, Christopher J Nachtsheim, and William Wasserman. *Applied linear statistical models*, volume 4. Irwin Chicago, 1996.
- [227] George AF Seber and Alan J Lee. *Linear regression analysis*, volume 936. John Wiley & Sons, 2012.
- [228] James E Gentle. *Numerical linear algebra for applications in statistics*. Springer Science & Business Media, 2012.
- [229] Carsten F Dormann, Jane Elith, Sven Bacher, Carsten Buchmann, Gudrun Carl, Gabriel Carré, Jaime R García Marquéz, Bernd Gruber, Bruno Lafourcade, and Pedro J *et. al* Leitão. Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. *Ecography*, 36(1):27–46, 2013.
- [230] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015.
- [231] Pierre Verger, Lisa Fressard, Fanny Collange, Arnaud Gautier, Christine Jestin, Odile Launay, Jocelyn Raude, Céline Pulcini, and Patrick Peretti-Watel. Vaccine hesitancy among general practitioners and its determinants during controversies: a national cross-sectional survey in france. *EBioMedicine*, 2(8):891–897, 2015.
- [232] Eliot Marshall. A shadow falls on hepatitis B vaccination effort, 1998.
- [233] Giulia Piccirilli, Tiziana Lazzarotto, Angela Chiereghin, Laura Serra, Liliana Gabrielli, and Marcello Lanari. Spotlight on measles in italy: why outbreaks of a vaccine-preventable infection continue in the 21st century. *Expert review of anti-infective therapy*, 13(3):355–362, 2015.
- [234] Francesco Aquino, Gabriele Donzelli, Emanuela De Franco, Gaetano Privitera, Pier Luigi Lopalco, and Annalaura Carducci. The web and public confidence in MMR vaccination in italy. *Vaccine*, 35(35):4494–4498, 2017.
- [235] J. Lewis and T. Speers. Misleading media reporting? The MMR stor. *Nature Reviews Immunology*, 3(11):913–918, 2003.
- [236] Marcel Salathé and Sebastian Bonhoeffer. The effect of opinion clustering on disease outbreaks. *Journal of The Royal Society Interface*, 5(29):1505–1508, 2008.
- [237] Saki Takahashi, C Jessica E Metcalf, Matthew J Ferrari, Andrew Tatem, and Justin Lessler. The geography of measles vaccination in the african great lakes region. *Nature Communications*, 8, 2017.
- [238] Centers for Disease Control and Prevention. Telebriefing, 2013. National immunization survey, vaccine for children program, and recent measles outbreaks in the U.S. press briefing transcript. http://www.cdc.gov/media/releases/2013/t0912_measles-outbreaks-data.html, 2014. [Online; accessed 04-June-2016].
- [239] Fumiko Samejima. Graded response model. In *Handbook of modern item response theory*, pages 85–100. Springer, 1997.
- [240] Charles A Michael, NSTOP OPV Refusal Study Team, Ikechukwu U Ogbuanu, NSTOP OPV Refusal Study Team, Aaron D Storms, NSTOP OPV Refusal Study Team, Chima J Ohuabunwo, NSTOP OPV Refusal Study Team, Melissa Corkum, and NSTOP OPV Refusal Study Team *et. al*. An assessment of the reasons for oral poliovirus vaccine refusals in northern Nigeria. *The Journal of infectious diseases*, 210(suppl_1):S125–S130, 2014.

- [241] C Bishop. Pattern recognition and machine learning (information science and statistics), 1st edn. 2006. corr. 2nd printing edn. *Springer, New York*, 2007.
- [242] Zoubin Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, 2015.
- [243] Jake Lever, Martin Krzywinski, and Naomi Altman. Points of significance: Principal component analysis, 2017.
- [244] Frederic M Lord and Melvin R Novick. *Statistical theories of mental test scores*. IAP, 2008.
- [245] Frederic M Lord. *Applications of item response theory to practical testing problems*. Routledge, 1980.
- [246] Jean-Paul Fox. *Bayesian item response modeling: Theory and applications*. Springer Science & Business Media, 2010.
- [247] Jing Cao and S Lynne Stokes. Bayesian IRT guessing models for partial guessing behaviors. *Psychometrika*, 73(2):209, 2008.
- [248] Georg Rasch. *Probabilistic models for some intelligence and attainment tests*. ERIC, 1993.
- [249] Jacob Cohen. The cost of dichotomization. *Applied psychological measurement*, 7(3):249–253, 1983.
- [250] Joseph Bafumi, Andrew Gelman, David K Park, and Noah Kaplan. Practical issues in implementing and understanding bayesian ideal point estimation. *Political Analysis*, 13(2):171–187, 2005.
- [251] Ronald K Hambleton and Hariharan Swaminathan. *Item response theory: Principles and applications*. Springer Science & Business Media, 2013.
- [252] Kimberly S Maier. A rasch hierarchical measurement model. *Journal of Educational and Behavioral Statistics*, 26(3):307–330, 2001.
- [253] Matthew S Johnson, Sandip Sinharay, and Eric T Bradlow. 17 hierarchical item response theory models. *Handbook of statistics*, 26:587–606, 2006.
- [254] Joop Hox and J Kyle Roberts. *Handbook of advanced multilevel analysis*. Psychology Press, 2011.
- [255] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*, volume 2. CRC press Boca Raton, FL, 2014.
- [256] Andrew *et. al* Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian analysis*, 1(3):515–534, 2006.
- [257] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of markov chain Monte Carlo*. CRC press, 2011.
- [258] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, (arXiv preprint arXiv:1601.00670), 2017.
- [259] Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, (6):721–741, 1984.
- [260] Andrew Gelman and Donald B Rubin. Inference from iterative simulation using multiple sequences. *Statistical science*, pages 457–472, 1992.
- [261] Adrian E Raftery and Steven Lewis. How many iterations in the Gibbs sampler? Technical report, 1991.

- [262] John *et. al* Geweke. *Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments*, volume 196. Federal Reserve Bank of Minneapolis, Research Department Minneapolis, MN, USA, 1991.
- [263] Harold Jeffreys. *The theory of probability*. OUP Oxford, 1998.
- [264] Gideon *et. al* Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
- [265] Hirotugu Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.
- [266] David J Spiegelhalter, Nicola G Best, Bradley P Carlin, and Angelika Van Der Linde. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639, 2002.
- [267] Martyn Plummer. rjags: Bayesian graphical models using MCMc. Rpackage version 2.1. 0-10, 2010.
- [268] Victoria M Gammimo, Adamu Nuhu, Sue Gerber, Alex Gasasira, David E Sugerman, Fadinding Manneh, Paul Chenoweth, Molly R Kurnit, and Emmanuel A Abanida. An evaluation of polio supplemental immunization activities in Kano, Katsina, and Zamfara states, Nigeria: Lessons in progress. *The Journal of infectious diseases*, 210(suppl_1):S91–S97, 2014.
- [269] Jennifer G Cooke and Farha Tahir. Polio in Nigeria: the race to eradication. 2012.
- [270] Sophie Arie. Murder of nine polio workers in Nigeria threatens eradication efforts. *BMJ: British Medical Journal (Online)*, 346, 2013.
- [271] Robert J Vandenberg and Charles E Lance. A review and synthesis of the measurement invariance literature: Suggestions, practices, and recommendations for organizational research. *Organizational research methods*, 3(1):4–70, 2000.
- [272] Catherine Hough-Telford, David W Kimberlin, Inmaculada Aban, William P Hitchcock, Jon Almquist, Richard Kratz, and Karen G O’Connor. Vaccine delays, refusals, and patient dismissals: a survey of pediatricians. *Pediatrics*, page e20162127, 2016.
- [273] Saad B Omer, Daniel A Salmon, Walter A Orenstein, M Patricia Dehart, and Neal Halsey. Vaccine refusal, mandatory immunization, and the risks of vaccine-preventable diseases. *New England Journal of Medicine*, 360(19):1981–1988, 2009.
- [274] Centers for Disease Control. www.cdc.gov/.
- [275] Public Health England. www.gov.uk/government/organisations/public-health-england.
- [276] Aron Culotta. Towards detecting influenza epidemics by analyzing twitter messages. In *Proceedings of the first workshop on social media analytics*, pages 115–122. ACM, 2010.
- [277] A Demetri Pananos, Thomas M Bury, Clara Wang, Justin Schonfeld, Sharada P Mohanty, Brendan Nyhan, Marcel Salathé, and Chris T Bauch. Critical dynamics in population vaccinating behavior. *Proceedings of the National Academy of Sciences*, page 201704093, 2017.
- [278] Marcel Salathé and Shashank Khandelwal. Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control. *PLoS computational biology*, 7(10):e1002199, 2011.

- [279] Demographic and Health Surveys. <https://dhsprogram.com/>.
- [280] Zhen Wang, Michael A Andrews, Zhi-Xi Wu, Lin Wang, and Chris T Bauch. Coupled disease–behavior dynamics on complex networks: A review. *Physics of life reviews*, 15:1–29, 2015.
- [281] Frederik Verelst, Lander Willem, and Philippe Beutels. Behavioural change models for infectious disease transmission: a systematic review (2010–2015). *Journal of The Royal Society Interface*, 13(125):20160820, 2016.
- [282] Allan L McCutcheon. *Latent class analysis*. Number 64. Sage, 1987.
- [283] Arthur White and Thomas Brendan Murphy. Bayeslca: an R package for bayesian latent class analysis. 2014.
- [284] Marcel Croon. Latent class analysis with ordered latent classe. *British Journal of Mathematical and Statistical Psychology*, 43(2):171–192, 1990.
- [285] Drew A Linzer and Jeffrey B *et. al* Lewis. polca: An R package for polytomous variable latent class analysis. *Journal of statistical software*, 42(10):1–29, 2011.
- [286] Jun Shao and Bob Zhong. Last observation carry-forward and last observation analysis. *Statistics in medicine*, 22(15):2429–2441, 2003.
- [287] Per Jonsson and Claes Wohlin. An evaluation of k-nearest neighbour imputation using Likert data. In *Software Metrics, 2004. Proceedings. 10th International Symposium on*, pages 108–118. IEEE, 2004.
- [288] Roderick P McDonald and Magdalena M-C Mok. Goodness of fit in item response models. *Multivariate Behavioral Research*, 30(1):23–40, 1995.
- [289] Konstantinos Sechidis and Gavin Brown. Markov blanket discovery in positive-unlabelled and semi-supervised data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 351–366. Springer, 2015.

A | Appendix A

A.1 GAPMINDER and COVERAGE Data

A.1.1 Socio-economic factors

The 190 socio-economic factors used in the correlative study in Section 2.2.2 are listed in Table A.1 with missing data statistics. The reduced set of socio-economic factors are shown in the left-hand column with accompanying statistics – for each SEF – on: the number of countries with greater than three data points ($N_{>3}$), the number of countries with no data (N_0), and the average number of data points per country (A). For each SEF, we then display – for each world region – the overall missing data percentage. (ALL = all regions; AFR = Africa region; AMR = Americas region; EMR = Eastern Mediterranean region; EUR = European region; SEAR = South-East Asia region; WPR = Western Pacific region [169], and see Appendix A.1.2).

These missing data statistics provide information on the nature of each time series. SEFs such as **child mortality** (3) and **mean years in school** (91-96) have on average 28.4 and 24.9 time-series data points per country, respectively. However, many SEFs have high missing data fractions, suggesting measurement at regularly-spaced intervals; for example, **dollar billionaires** (20), **gasoline price** (101), and **population structures** (339-366). (Numbers in parentheses correspond to numbered SEFs in the table.) This variation in missing data structure requires a sensible treatment of missing data.

A.1.2 WHO regional classifications

Countries are classified according to WHO regional groupings [169]. These groupings are as follows:

Africa region (AFR)

Algeria, Angola, Benin, Botswana, Burkina Faso, Burundi, Cameroon, Cabo Verde, Central African Republic, Chad, Comoros, Congo, Côte d’Ivoire, Democratic Republic of the

Socio-economic factor (SEF)	missing data (%)									
	$N_{>3}$	N_0	A	ALL	AFR	AMR	EMR	EUR	SEAR	WPR
0 Children per woman (total fertility)	168	28	26.5	14.4	8.5	11.8	4.7	18.5	9.1	29.6
1 CO2 emissions (tonnes per person)	175	21	25.3	18.5	10.9	10.7	4.5	31.8	19.6	25.1
2 Income per person (GDP capita PPPD inflation adjusted)	181	14	28.0	9.8	6.5	2.9	6.0	18.9	9.1	8.7
3 Child mortality (0-5 year olds dying per 1000 born)	182	14	28.4	8.5	6.5	4.4	5.0	15.1	9.1	6.2
4 Life expectancy (years)	168	28	26.6	14.3	8.5	11.8	4.5	18.2	9.1	29.6
18 Aid received per person (current USD)	142	54	20.3	34.4	5.6	16.1	10.7	80.5	18.2	39.7
19 Aid received (% of GNI)	140	56	18.7	39.7	10.6	19.3	28.4	81.2	28.7	45.4
20 Dollar billionaires per million people	183	13	3.7	88.0	87.9	87.5	87.7	88.5	88.3	87.6
21 Average age of dollar billionaires (years)	183	13	3.7	88.0	87.9	87.5	87.7	88.5	88.3	87.6
22 Total number of dollar billionaires	183	13	3.7	88.0	87.9	87.5	87.7	88.5	88.3	87.6
26 GDP capita (USD inflation adjusted)	172	23	25.0	19.3	9.0	10.8	25.5	23.4	31.7	29.3
29 GDP capita growth over next 10 years	151	39	10.4	66.3	59.4	62.0	68.2	74.8	62.5	66.8
30 GDP capita growth (% per year)	178	18	25.6	17.6	11.5	4.3	26.5	23.2	23.5	23.8
33 GNI capita (Atlas method current USD)	175	21	23.6	23.9	11.5	12.0	29.3	32.3	33.1	35.2
34 GNI per capita (PPP current international D)	167	29	23.5	24.3	12.2	15.2	28.6	30.9	32.0	37.0
43 Inflation (annual p)	176	20	25.0	19.4	9.2	10.1	22.1	26.2	25.2	30.6
45 Investments (% of GDP)	167	28	23.6	23.9	19.0	5.5	28.3	28.9	24.6	41.7
47 Foreign direct investment net inflows (% of GDP)	174	22	23.9	22.9	14.3	4.8	20.8	37.4	35.5	27.4
56 Agriculture (% of GDP)	163	32	22.3	27.9	15.4	22.1	29.9	35.7	24.9	41.0
57 Industry (% of GDP)	169	27	22.9	26.0	16.7	16.9	37.1	31.4	23.2	34.9
58 Services (% of GDP)	168	28	22.8	26.5	16.7	16.5	36.8	33.3	23.2	35.2
59 Exports (% of GDP)	169	26	24.0	22.6	11.0	10.8	24.8	31.3	19.9	39.4
60 Imports (% of GDP)	170	25	24.2	22.1	15.2	5.6	28.2	31.1	18.8	32.7
62 Arms imports (USD inflation adjusted)	155	26	15.5	50.0	60.2	52.5	24.5	48.8	34.6	58.3
64 High technology exports (% of manufactured exports)	151	38	12.8	58.6	64.4	45.3	61.6	54.0	65.4	69.2
65 Merchandise trade (% of GDP)	173	23	24.6	20.7	10.9	4.2	19.9	36.2	27.6	24.6
67 Trade balance (% of GDP)	170	25	24.2	22.0	15.2	5.6	28.2	30.8	18.8	32.7
71 Military expenditure (% of GDP)	150	43	15.1	51.3	45.2	53.6	41.5	49.5	55.4	68.7
74 Ratio of girls to boys in primary and secondary education (%)	169	23	16.8	45.7	44.9	44.7	35.3	41.2	56.3	61.8
82 Children out of school primary	156	31	13.3	57.2	52.9	52.8	50.1	55.8	74.8	71.3
83 Children out of school primary female	150	33	11.0	64.6	57.4	61.9	58.7	63.5	78.9	82.0
84 Children out of school primary male	150	33	11.0	64.6	57.4	61.9	58.7	63.5	78.9	82.0
85 Primary completion rate total (% of relevant age group)	158	30	14.2	54.3	41.8	48.7	47.5	58.3	71.0	73.5
86 Primary school completion (% of boys)	154	33	12.7	59.0	44.2	56.7	50.4	63.9	78.0	76.6
87 Primary school completion (% of girls)	154	33	12.7	59.0	44.2	56.7	50.4	63.9	78.0	76.6
91 Mean years in school (women 25 years and older)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
92 Mean years in school (men 25 years and older)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
93 Mean years in school (women 25 to 34 years)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
94 Mean years in school (men 25 to 34 years)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
95 Mean years in school (women of reproductive age 15 to 44)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
96 Mean years in school (women p men 25 to 34 years)	163	33	24.9	19.5	9.4	20.3	12.0	24.3	29.6	28.3
98 Energy use per person	153	42	19.2	37.9	55.5	26.1	14.8	29.0	33.4	61.2
101 Pump price for gasoline (USD per liter)	149	41	4.6	85.3	81.1	85.1	83.9	86.3	86.2	91.4
128 Flood deaths annual number	155	41	22.9	26.0	22.4	17.5	27.7	30.3	15.0	37.6
129 Storm deaths annual number	143	53	21.2	31.7	44.3	14.7	44.7	33.7	15.0	23.8
136 Flood affected annual number	155	41	22.9	26.0	22.4	17.5	27.7	30.3	15.0	37.6
137 Storm affected annual number	143	53	21.2	31.7	44.3	14.7	44.7	33.7	15.0	23.8
140 CO2 int of econ output (kg CO2 per 2005 PPP D of GDP)	169	27	22.3	28.0	21.4	13.1	31.2	38.2	28.2	35.0
142 Sulfur emissions per person (kg)	153	43	16.4	47.1	45.2	46.2	35.3	45.8	44.6	64.9
143 Yearly CO2 emissions (1000 tonnes)	171	25	24.7	20.3	13.0	13.6	8.9	33.0	19.6	25.1
144 Total sulfur emission (kilotonnes)	153	43	16.4	47.1	45.2	46.2	35.3	45.8	44.6	64.9
149 Forest coverage (%)	173	23	2.6	91.5	90.9	90.6	91.6	92.4	91.2	91.4
155 Agricultural land (% of land area)	178	18	25.4	18.2	11.2	6.1	9.4	34.0	12.0	22.9
158 Renewable water (cu meters per person)	161	35	4.7	84.8	82.7	83.5	81.8	86.8	82.4	89.2
159 Internal renewable water (cu meters per person)	160	36	4.7	84.8	82.7	83.5	82.7	86.6	82.4	89.2
167 Infant mortality (rate per 1000 births)	183	13	24.3	21.7	22.3	16.7	13.3	29.3	20.8	18.6
253 Total health spending (% of GDP)	144	52	11.7	62.3	59.6	59.3	63.6	63.4	67.2	65.7
254 Government share of total health spending (%)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
255 Private share of total health spending (%)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
256 Out of pocket share of total health spending (%)	140	56	11.3	63.6	52.5	92.4	50.9	63.4	57.8	60.0
257 Government health spending of total gov. spending (%)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
258 Total health spending per person (international D)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
259 Government health spending per person (international D)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
260 Total health spending per person (USD)	176	20	14.2	54.2	51.4	53.2	50.9	55.0	67.2	56.0
261 Government health spending per person (USD)	178	18	14.4	53.7	51.4	53.2	50.9	55.0	57.8	56.0
262 Medical Doctors (per 1000 people)	174	20	11.5	62.9	76.3	69.4	64.5	41.1	69.2	71.8
264 Adults with HIV (% age 15 to 49)	138	52	15.0	51.6	41.2	45.6	62.8	51.3	44.0	71.8
273 Births attended by skilled health staff (% of total)	146	24	5.7	81.6	87.9	76.0	85.0	75.8	82.4	86.7
275 Maternal mortality ratio (per 100000 live births)	166	24	2.6	91.6	90.8	91.4	90.3	92.0	92.1	93.4
277 Maternal deaths lifetime risk (per 1000)	158	30	2.5	92.0	91.1	91.9	90.3	92.2	92.1	94.9
279 Suicide (per 100000 people)	179	15	8.9	71.2	89.4	63.0	87.0	48.1	89.1	77.1
281 Sugar per person (g per day)	159	37	19.1	38.3	27.1	21.7	46.8	48.9	34.0	52.2
282 Food supply (kilocalories person and day)	160	34	21.6	30.3	18.3	12.3	40.2	40.8	26.1	46.5
287 Body Mass Index (BMI) men Kg m2	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
288 Body Mass Index (BMI) women Kg m2	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
289 Blood pressure (SBP) men mmHg	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
290 Blood pressure (SBP) women mmHg	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
291 Cholesterol (fat) in blood men mmol L	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
292 Cholesterol (fat) in blood women mmol L	178	18	26.3	15.0	12.4	9.2	10.7	21.8	15.0	16.8
293 Infectious TB new cases per 100000 estimated	182	14	16.7	46.1	44.4	43.6	47.2	48.3	47.2	46.2
294 Infectious TB new cases per 100000 reported	182	14	12.7	59.0	60.1	55.4	59.2	60.9	57.2	58.1
297 Infectious TB detection rate (%)	177	19	12.4	59.9	60.8	57.3	59.2	62.2	57.2	58.3
298 Infectious TB detection rate (%) DOTS only	168	23	8.8	71.8	70.1	71.3	67.3	77.8	63.9	69.7
299 Infectious TB treatment (DOTS) completed (%)	165	27	8.2	73.6	70.7	73.2	68.9	80.6	64.5	72.2
300 TB programme (DOTS) population coverage (%)	182	14	10.8	65.3	66.4	64.0	64.7	65.8	63.3	65.0
301 All forms of TB new cases per 100000 estimated	182	13	16.7	46.1	44.4	43.6	47.2	48.2	47.2	46.2
302 All forms of TB existing cases per 100000 estimated	182	14	16.7	46.1	44.4	43.6	47.2	48.3	47.2	46.2
303 All forms of TB deaths per 100000 estimated	182	14	16.7	46.1	44.4	43.6	47.2	48.3	47.2	46.2
304 All forms of TB new cases per 100000 reported	182	14	24.5	21.1	22.4	14.6	24.2	22.6	23.2	20.2

Socio-economic factor (SEF)	missing data (%)									
	$N_{>3}$	N_0	A	ALL	AFR	AMR	EMR	EUR	SEAR	WPR
309 All forms of TB detection rate (%)	177	19	11.1	64.1	65.0	63.5	63.0	65.6	60.7	62.7
310 All forms of TB detection rate (%) DOTS only	168	23	8.8	71.5	70.1	70.7	67.3	77.8	63.9	68.6
311 TB with HIV+ new cases per 100000 estimated	182	14	15.5	50.1	45.6	47.1	55.1	52.5	47.2	54.1
312 TB with HIV+ existing cases per 100000 estimated	182	14	15.7	49.4	45.6	47.1	51.5	51.3	47.2	54.1
313 TB with HIV+ deaths per 100000 estimated	150	46	13.7	55.9	48.1	53.9	59.2	57.2	47.2	69.9
324 Broadband subscribers (per 100 people)	178	18	9.4	69.6	70.6	66.6	69.1	70.1	67.7	71.7
326 Cell phones (per 100 people)	176	20	27.1	12.6	4.8	10.1	3.2	17.0	18.2	25.6
327 Fixed line and mobile phone subscribers (per 100 people)	176	20	22.6	27.1	21.6	19.5	22.1	33.3	24.0	38.6
329 Internet users (per 100 people)	176	20	15.5	50.1	48.9	47.4	47.2	48.5	53.4	59.6
331 Personal computers (per 100 people)	167	24	10.2	67.0	69.9	64.9	64.5	66.4	65.4	68.5
332 Improved sanitation overall access (%)	170	26	17.3	44.2	39.2	40.8	40.8	47.4	46.9	52.4
333 Improved sanitation urban access (%)	173	23	18.0	42.0	37.9	40.2	33.0	44.6	46.9	51.3
334 Improved sanitation rural access (%)	169	27	17.1	44.7	38.8	40.8	40.8	48.6	46.9	54.1
335 Roads paved (p of total roads)	156	26	9.8	68.3	71.7	65.7	63.3	66.8	69.2	72.8
336 Improved water source overall access (%)	174	22	17.9	42.3	38.5	42.3	39.3	44.8	46.6	44.2
337 Improved water source urban access (%)	177	19	18.6	40.0	37.2	39.1	34.2	42.3	44.6	44.2
338 Improved water source rural access (%)	171	25	16.8	45.8	41.1	44.8	41.9	48.4	49.0	51.6
339 Population aged 0 to 4 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
340 Population aged 5 to 9 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
341 Population aged 10 to 14 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
342 Population aged 15 to 19 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
343 Population aged 20 to 39 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
344 Population aged 40 to 59 years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
345 Population aged 60+ years both sexes (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
346 Population aged 0 to 4 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
347 Population aged 5 to 9 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
348 Population aged 10 to 14 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
349 Population aged 15 to 19 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
350 Population aged 20 to 39 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
351 Population aged 40 to 59 years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
352 Population aged 60+ years (total number)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
353 Population aged 0 to 4 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
354 Population aged 5 to 9 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
355 Population aged 10 to 14 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
356 Population aged 15 to 19 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
357 Population aged 20 to 39 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
358 Population aged 40 to 59 years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
359 Population aged 60+ years female (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
360 Population aged 0 to 4 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
361 Population aged 5 to 9 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
362 Population aged 10 to 14 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
363 Population aged 15 to 19 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
364 Population aged 20 to 39 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
365 Population aged 40 to 59 years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
366 Population aged 60+ years male (%)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
367 Population growth (annual %)	176	20	27.8	10.3	4.3	8.8	0.6	12.7	18.2	22.3
368 Crude birth rate (births per 1000 population)	168	28	5.1	83.4	82.3	82.9	81.5	84.2	82.4	86.4
369 Crude death rate (deaths per 1000 population)	168	28	5.1	83.4	82.3	82.9	81.5	84.2	82.4	86.4
370 Teen fertility rate (births per 1000 women ages 15 19)	180	15	17.1	44.8	54.8	43.1	49.4	30.7	54.3	50.5
372 Sex ratio (all age groups)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
373 Sex ratio (0 to 14 years)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
374 Sex ratio (15 to 24 years)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
375 Sex ratio (15 to 49 years)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
376 Sex ratio (above 50 years)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
379 Urban population (% of total)	176	20	27.8	10.2	4.3	8.8	0.0	12.7	18.2	22.2
380 Urban population growth (annual %)	176	20	27.8	10.4	4.3	8.8	0.6	13.3	18.2	22.2
381 Children and elderly (per 100 adults)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
382 Median age (years)	168	28	6.0	80.6	79.3	80.1	78.4	81.5	79.5	84.1
383 Population density (per square km)	178	18	28.2	9.2	4.3	8.8	0.0	12.7	18.2	14.8
386 Murder (per 100000 people)	180	15	9.5	69.5	88.9	57.6	86.4	46.9	85.3	76.5
388 Democracy score (use as color)	154	42	24.0	22.5	12.2	29.4	6.0	23.8	18.2	44.4
389 HDI (Human Development Index)	163	33	6.0	80.7	78.7	78.7	78.2	81.9	81.8	85.7
408 Aged 15 to 24 employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
409 Aged 15+ employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
410 Females aged 15 24 employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
411 Females aged 15+ employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
412 Males aged 15 24 employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
413 Males aged 15+ employment rate (%)	157	39	13.6	56.1	51.0	56.5	50.1	57.1	50.1	69.5
414 Aged 15 64 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
415 Aged 25 54 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
416 Aged 15+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
417 Aged 65+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
418 Females aged 15 64 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
419 Females aged 25 54 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
420 Females aged 15+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
421 Females aged 65+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
422 Males aged 15 64 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
423 Males aged 25 54 labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
424 Males aged 15+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
425 Males aged 65+ labour force participation rate (%)	164	32	23.4	24.4	17.4	23.0	13.8	29.4	17.9	39.8
441 Under five mortality from CME (per 1000 born)	182	14	28.3	8.6	6.5	4.4	5.0	15.1	12.0	6.2
442 Under five mortality from IHME (per 1000 born)	159	37	25.1	18.9	8.5	14.7	9.1	21.8	18.2	44.4
443 Old version of Income per person (version 3)	174	21	23.1	25.4	16.3	23.0	27.6	32.7	17.9	30.9
444 Old version of Income per person (version 8)	180	15	26.9	13.2	9.5	6.2	9.2	21.6	12.3	14.8
445 Alternative GDP capita (PPP inflation adjusted) from PWT	160	32	18.6	40.0	30.4	35.3	43.3	49.0	36.1	42.9
446 Subsistence incomes per person	181	14	27.1	12.7	9.5	6.2	9.2	21.6	12.0	11.2

Table A.1: **Reduced set of SEFs** Reduced SEFs as described in Section 2.2.2; SEFs highlighted grey are found to have the strongest correlations with DTP3 coverage.

Congo, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Gambia, Ghana, Guinea, Guinea-Bissau, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritania, Mauritius, Mozambique, Namibia, Niger, Nigeria, Rwanda, Sao Tome and Principe, Senegal, Seychelles, Sierra Leone, South Africa, South Sudan, Swaziland, Togo, Uganda, United, Republic of Tanzania, Zambia, Zimbabwe.

Americas region (AMR)

Antigua and Barbuda, Argentina, Bahamas, Barbados, Belize, Bolivia (Plurinational State of), Brazil, Canada, Chile, Colombia, Costa Rica, Cuba, Dominica, Dominican Republic, Ecuador, El Salvador, Grenada, Guatemala, Guyana, Haiti, Honduras, Jamaica, Mexico, Nicaragua, Panama, Paraguay, Peru, Saint Kitts and Nevis, Saint Lucia, Saint Vincent and the Grenadines, Suriname, Trinidad and Tobago, United States of America, Uruguay, Venezuela (Bolivarian Republic of).

Eastern Mediterranean region (EMR)

Afghanistan, Bahrain, Djibouti, Egypt, Iran (Islamic Republic of), Iraq, Jordan, Kuwait, Lebanon, Libya, Morocco, Oman, Pakistan, Qatar, Saudi Arabia, Somalia, Sudan, Syrian Arab Republic, Tunisia, United Arab Emirates, Yemen.

EURpean region (EUR)

Albania, Andorra, Armenia, Austria, Azerbaijan, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Greece, Hungary, Iceland, Ireland, Israel, Italy, Kazakhstan, Kyrgyzstan, Latvia, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, Norway, Poland, Portugal, Republic of Moldova, Romania, Russian Federation, San Marino, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Tajikistan, The former Yugoslav Republic of Macedonia, Turkey, Turkmenistan, Ukraine, United Kingdom, Uzbekistan.

South-East Asia region (SEAR)

Bangladesh, Bhutan, Democratic People's Republic of Korea, India, Indonesia, Maldives, Myanmar, Nepal, Sri Lanka, Thailand, Timor-Leste.

Western Pacific region (WPR)

Australia, Brunei Darussalam, Cambodia, China, Cook Islands, Fiji, Japan, Kiribati, Lao People's Democratic Republic, Malaysia, Marshall Islands, Micronesia (Federated States of), Mongolia, Nauru, New Zealand, Niue, Palau, Papua New Guinea, Philippines, Republic of Korea, Samoa, Singapore, Solomon Islands, Tonga, Tuvalu, Vanuatu, Viet Nam.

A.1.3 Socio-economic factor imputation

The GAPMINDER dataset introduced in Section 2.2.2 comprises 36,100 socio-economic time-series data (190 SEFs and 190 countries). The primary challenge in imputing these time-series data is that the missing data fraction varies across time-series and that some time-series have no values whatsoever (see Table A.1). Without imputation, pairwise deletion [171]¹ would be required when calculating Spearman’s rank between two vectors and this approach hinders our ability to provide unbiased estimates [172]. Therefore, instead of deleting missing data, we show how a combination of missing data methods can provide robust estimates of the time-varying Spearman’s rank correlation (the quantity of interest; Equation (2.13)). In the rest of this section, it is instructive to define two matrix representations of the same socio-economic data to more easily explain the different imputation methods considered. The ‘country-year’ formulation is a matrix of data $X_{ct,s}$, where rows represent country-years and columns represent socio-economic variables, $c = 1, \dots, 190$, $t = 1, \dots, 31$, and $s = 1, \dots, 190$. The ‘time-series’ formulation contains the same information but in a matrix $X'_{cs,t}$ so that each row of the matrix contains the full time-series for socio-economic factor s and country c . The following imputation methods are considered.

Mean replacement (MR) Mean replacement replaces all missing points in a time-series with the mean value across all points that are present [171].

Linear interpolation and 2-point/6-point/all-point extrapolation (L2, L6, LA)

The linear interpolant is used between any two points which contain missing data between them. Extrapolation is based on a linear fit to the most recent two (2-point), six (6-point), or all (all-point) data points. Interpolation is performed first such that the extrapolation may use interpolated points [173].

KNN imputation KNN imputation differs from MR and linear interpolation methods in that it does not impute a missing data point using data from that same time-series. Instead, KNN imputation averages over the K ‘nearest’ rows in the matrix X and replaces missing values in a given row (country-year) with an average of K corresponding values across the K nearest-neighbour rows [174]. K is a parameter to be learned. A common implementation of KNN imputation is to only allow missing values to be replaced from rows with no missing values [174, 175]; however, only a small fraction of rows are complete cases in the GAPMINDER dataset, and so we allow any country-year (row) to be utilised in the distance calculation

¹In the case of seeking correlations between two variables, pairwise deletion is the deletion of an entire data point if at least one variable has a missing value.

between two rows, so long as it contains a corresponding value for the target. We account for distances calculated between rows with large missing data fractions by adding a small penalisation term to the (Euclidean) distance between rows by multiplying the Euclidean distance by $\sqrt{(p/p_*)}$, where p is the total number of socio-economic factors (rows) and p_* is the number of non-missing instances in each row. Data across each SEF is linearly scaled to the $[0,1]$ interval before applying K NN imputation.

To test imputation methods, we generate a data matrix C with no missing values from the complete matrix X' . Missing values are generated in this matrix according to the structure of real missing data in Table A.1. We model the missing data structure in the full time-series matrix X' by forming the 190 by 190 matrix P whose entries are given by

$$P_{cs} = \frac{1}{31} \sum_{t=1}^{31} M_{cs,t}^{X'}, \quad (\text{A.1})$$

where $M^{X'}$ is a missing data indicator matrix for X' (that is, $M_{cs,t}^{X'} = 1$ if $X'_{cs,t}$ is missing and 0 otherwise). P corresponds to the fraction of missing data for country c and factor s . (We assume here that each data point in a socio-economic factor time-series is equally likely to be missing). A random sub matrix of P (of dimension equal to the dimension of C) is generated, and missing values are inserted into C according to,

$$M_{cs,t}^C = \begin{cases} 1 & \text{if } U_{cs,t} < P_{cs} \\ 0 & \text{otherwise} \end{cases}$$

where M^C is a missing data indicator matrix for C , and $U_{cs,t} \sim \mathcal{U}(0,1)$, and where U has the same dimension as C . This method allows us to recover an instance of the missing data structure in X' as the sampling will induce a correlation structure between missing data values between SEFs and country-years.

Imputation performance of each method is evaluated (on each datapoint) by introducing a normalised absolute error (NAE),

$$\epsilon_{cs,t} = \left| \frac{C_{cs,t} - I_{cs,t}}{\frac{1}{n_s} \sum_{c,t} C_{cs,t}} \right|, \quad (\text{A.2})$$

where I is the matrix of imputed values and where we only consider errors ϵ on imputed values. The denominator is the mean value of SEF s (n_s is the number of non-missing data points for SEF s). This error metric allows us to compare errors both between socio-economic factors (which may have different scales) and within scales (which might also vary between

several orders of magnitude). For KNN imputation the matrix C is first z -scored and then values are rescaled to their original values. The procedure outlined above is repeated 20 times to ensure we obtain fair estimates of errors, since the mechanism above may result in a high fraction of missing data for a given SEF which will increase the error estimates for this factor.

A.1.4 Imputation results

Linear interpolation with 2-point extrapolation is found to be the optimal imputation strategy when the number of missing data points is at least two and mean imputation is preferred to KNN when there is only one data point. The median across all NAE errors for these methods are displayed as a function of the fraction f of non-missing time-series data in Figure A.1 A. Errors increase as the number of missing data points increase. For KNN imputation, we train over the optimal value of K for when there are zero and a non-zero number of missing data points per time-series. We find that $K = 1$ is optimal when there is at least one missing data per time series (which is most likely a model implementing ‘last observation carry-forward’ imputation see, for example, [286]), and $K = 40$ optimal when all time-series data are missing. This larger value suggests that when there are no missing data, it cannot




Figure A.1: **Linear interpolation and extrapolation outperforms KNN imputation and mean imputation for the GAPMINDER dataset** (A) Linear interpolation with extrapolation using the two most recent data points outperforms all other methods when the fraction of non-missing time-series data is greater than $1/31$. Mean replacement outperforms the optimal KNN method when there is one time-series datum ($f = 1/31$) and KNN is the only method considered when $f = 0$. (B) Error distributions for the optimal method are highly skewed, but suggest that most imputed values have a low error, but there are also some very large errors.

simply ‘carry-one-forward’ from time-series values in that given time-series, and relies on an average over many country-years. The median error is used since error distributions are highly skewed across imputation methods (Figure A.1 B).

There often exist many orders of magnitude between the observed and imputed values in the test matrix C (Figure A.2 A). In KNN imputation the imputation error increases with the number of missing values in a country-year row (Figure A.2 B). By choosing to consider country-years with above a threshold value of present data values we can limit the KNN imputation error. This error increase is because there are fewer values used in the calculation of the Euclidean distance between country-years (Figure A.2 C). Large imputation errors are also associated with SEFs whose time-series are prone to sudden, large jumps; for example, factors relating to natural disasters (127,128,136,137), inflation (43), and GDP growth (30) (Figure A.2 C). Factors which vary more smoothly over time such blood pressure or education levels have much lower errors (Figure A.2 C).

Although KNN imputation is associated with large imputation error in some cases, the combined method overall compares favourably to a pairwise deletion approach (see Section A.1.5). Examples of KNN imputed time-series are shown in Figure A.3 A. In Figure A.3 B, example time-series are shown for **food supply** (282) which use L2 imputation.

Many SEFs with large prediction error were identified in the imputation testing process. The raw dataset X contains more SEFs, and many of these we believe will also be associated with large error. These SEFs are not removed from our analysis, but rather we make a note of these SEFs in the case that they correlate strongly (and significantly) with vaccine coverage. However, this was not found to be the case. Since KNN error increases with the fraction of missing data in a given country-year, data is not imputed (by KNN) when there exists more than 50% missing data values in a given country-year. Example KNN imputations for **income per person** (3) for a random selection of countries are shown in Figure A.3 A demonstrating the utility of KNN imputation.

The optimal value of K ($K = 40$) when $f = 0$ is in keeping with the recommendation of using $K = \sqrt{n_{obs}}$ [287], where $n_{obs} = 1669$ is the number of cases (rows) used in the imputation method C . We thus increase this value to $K = 80$ for KNN imputation of the matrix X as we have 6076 rows.

As outlined in the main text (Section 2.2.2), the **GAPMINDER** data has a missing data fraction of 48.4%, which is reduced to 9.5% after we apply the combination imputation method. Of the data that is imputed, most of it (87.9%) is by linear interpolation (and of which a very large proportion – 91.0% – have greater than five data points per time series, so we can expect a very small imputation error); a very small amount of time-series have one data point and are thus imputed by mean-replacement (0.02% of the total imputed); and

the remainder (12.1%) is *KNN* imputed. After imputation, there are a total of 23 countries which have no time-series data for at least one socio-economic factor (the mean number of socio-economic factors missing for these countries is 153), including six which have no data for any SEF (denoted R in the following list): Andorra, Brunei Darussalam, Central African Republic, DR Congo (R), Cook Islands, Czech Republic (and Czechoslovakia), Dominican Republic, Ivory Coast (R), Marshall Islands, Monaco, Nauru, Niue, occupied Palestinian territory (R), Republic of Moldova (R), Russia (and former USSR), Saint Kitts and Nevis, San Marino, Slovakia, Timor Leste (R), Tuvalu, Yemen. Countries with no data for any SEF are removed from the analysis: as these countries are from a range of WHO regions, we do not expect any noticeable bias to be introduced from the omission of these countries.




Figure A.2: **Overall imputation error varies by SEF and KNN-imputation error increases with fraction of missing data in a country-year.** (A) Scatter plot of imputed values versus real values for our clean matrix *C*. (B) Median error increases with the number of missing data values in a country-year for *KNN* imputation. (C) Box and whisker plots for all errors across all SEFs in our test matrix *C*: factors prone to sudden jumps in value have larger median errors. (The *x*-axis has been cut off at 10^{-10} for clarity.) Numbers refer to the SEFs referenced in Table A.1.




Figure A.3: **Example imputations using KNN and L2 imputation.** Time-series from the clean data matrix C have missing values generated according to the process as described in the text. (A) KNN imputed values are shown for the SEF income per person (3). (B) Imputed values are shown from the L2 interpolation method for the SEF food supply (282). The green time-series corresponds to the clean data from C ; red-squares denote imputed values; and blue-squares are the data that remain after generating missing values in C . The mean NAE error for each time-series is shown.

A.1.5 Test correlation study

A test correlation study is implemented to compare our L2+MR+KNN imputation method to pairwise deletion. In calculating the correlation between two variables, pairwise deletion will remove a data point if either one of the data is missing [171]. The correlation study replicates the time-varying correlation study introduced in Section 2.3.3. Specifically, we calculate the time-averaged Spearman’s correlation given by Equation (2.13) for our imputation method compared to pairwise deletion. The correlation is calculated over the years 1980-2010 (to observe historic trends) and 2001-2010 (to observe recent ones). As in the main text, the statistic is only calculated when there exist at least eight data points (in our test data set C , this means that only three WHO regions have data available).




Figure A.4: **Time-averaged Spearman’s rank values are robust to imputation of large numbers of missing data values** (A) For each WHO region present in our imputed data matrix we display the true and predicted Spearman’s rank values calculated on the clean data C . (B) The true value of the statistic is also plotted against values obtained from pairwise deletion of data in C .

We calculate the difference between the true time-averaged correlation and that calculated using imputed data. We report absolute mean differences (and standard deviations) in statistic values of $.03 \pm .06$ for AFR, $.05 \pm .04$ for AMR, and $.06 \pm .06$ for EUR for data between 1980 and 2010 and $.05 \pm .09$ for AFR, $.08 \pm .07$ for AMR, and $.06 \pm .06$ for EUR for data between 2001 and 2010. In comparing the true value of $\bar{\rho}$ to the pairwise deletion-predicted value, we report absolute mean differences (and standard deviations) in statistic values of $.06 \pm .07$ for AFR, $.08 \pm .07$ for AMR, and $.12 \pm .11$ for EUR for data between 1980 and 2010

and $.06 \pm .09$ for AFR, $.11 \pm .10$ for AMR, and $.09 \pm .11$ for EUR for data between 2001 and 2010. These results indicate the improved accuracy of our imputation method versus pairwise deletion.

To observe differences between the accuracy of pairwise deletion versus our combined approach, we display scatterplots between the true statistic value and the value obtained from an imputed data set (Figure A.4 A) and between the true statistic value and the value obtained from pairwise deletion (Figure A.4 B). These results highlight the problematic approach of pairwise deletion and underlines the robustness of our statistic of interest to our imputation method.

A.2 Coverage correlations with socio-economic factors supporting material

A.2.1 Multiple hypothesis correction using meta p -values

A total of 5,890 ($190 \text{ SEFs} \times 31 \text{ time points}$) individual correlations are performed for each region and so a multiple hypothesis testing approach needs to be applied to control for the simultaneous correlations. A procedure such as Bonferroni correction will apply a large correction to p -values since each p -value is multiplied by the number of tests [190]. This correction will be unnecessarily harsh since we know that data from successive years is correlated: we are thus overstating the number of independent hypotheses by using conventional Bonferroni correction to correct for correlations.

To address this concern, we consider a meta p -value approach for combining dependent p -values which we adapt from Delongchamp *et. al.* [192] and which accounts for the correlation structure amongst p -values which is present when considering averages of correlations of variables dependent throughout time. The meta p -value – described below – may be interpreted as an overall significance level for a set of p -values from correlated observations. After these meta p -values have been calculated for each correlation of coverage against socio-economic (within each region and over time), a Bonferroni method can then be used to control for multiple hypothesis amongst this smaller set of p -values.

Meta p -values

Under the null hypothesis, the p -value for a test statistic follows a $U(0, 1)$ random [192], and thus it can be transformed to a standard normal random variable using the inverse cumulative normal distribution function $z_i = \Phi^{-1}(1 - p_i)$. When there exists a correlation structure between p -values, we can assume that $\mathbf{z} \sim N(\mathbf{0}, R)$ under the null hypothesis, where R is the correlation matrix between the z -transformed p -values [192]. We can then obtain an overall

significance level for the set of p -values by forming the test statistic $T = \sum_{i=1}^m z_i / \sqrt{\mathbf{1}^T R \mathbf{1}}$, where $\mathbf{1} = (1, 1, \dots, 1)^T$. R is unknown and rather than estimate R from the specific type of data as suggested in [192], we instead estimate R by bootstrapping a number of p -values from the raw data and then estimate the covariance of the z -transformed values. The overall p -value is then obtained from $P = 1 - \Phi(T)$ and multiplied by 190 to Bonferroni correct for the number of socio-economic factors. These meta p -values are displayed Figure 2.5 in the main text.

A.2.2 Selection of SEFs for Figure 2.5

To restrict the number of socio-economic factors we consider in our comparative analysis between socio-economic factors across time and between WHO regions, we choose the top-ten correlating SEFs in each region and for each time-period over which our Spearman statistic is calculated (as described in Section 2.3.3). We provide more explicit details of this selection method here.

A top-ten correlating list of SEFs with Bonferroni-corrected meta p -values of less than 0.01 is obtained. During this process, many factors appear in the agglomerated list more than once; furthermore, these lists contain factors which are almost identical to each other, for example, income per person (2) and GD% per capita (26) and, when this is the case, we only include one of the two. Moreover, for utility of comparison, we include some factors which have a natural counterpart, and are thus useful for comparison; for instance, **water access: urban** (337) is added for contrast with the other water and sanitation variables. In total for DTP3 we remove factors which are replicated for GDP and GNI (26, 34, 444, 445, 456), primary education (74), child death and maternal mortality (277, 370, 441, 495, 508), cell phones (326), and population age distributions (108-111, 113, 122, 128) which can be summarised by median age (382). We *include* the following factors as they provide useful comparison to included factors with high correlation and significance: **primary completion male** (86), **improved sanitation access overall/rural** (332/334), and **improved water access urban** (337). We do not go through the same process as described above for each of the other vaccines (whose correlations are displayed in Supplementary Figures A.5 and A.6). However, we list here correlates within the agglomerated top-ten list which differ from DTP3. **BCG**: Dollar billionaires per million (20), Total health spending per person (260), Blood pressure women (290), All forms TB death per 100000 (303), Children and elderly per 100 adults (381), human development index (HDI) (389); **DTP1 no differences**; **MCV** Mean years in school (women to men) (96), energy use per person (98), CO2 intensity of economic output (140), total health spending (253), Out of pocket share of total health spending (256), HDI (389); **POL3** All forms of TB deaths per 100000 (303), 15-24 employment rate (male) (412). (Numbers refer to variable definitions in Table A.1).

A.2.3 Time-varying correlations of BCG, DTP1, MCV, POL3

The correlative study performed in Section 2.3.3 is repeated here for the vaccination programmes DTP1, MCV, POL3, and BCG. Following the same process as outlined in this Section, Figure 2.5 is repeated for the BCG and DTP1 vaccines in Figure A.5 and for the MCV and POL vaccines in Figure A.6.

The following differences between these figures and the main-text figure are now highlighted. The effect of public health spending and access to sanitation in South-East Asia, which was moderate for DTP3, was very high and often increasing for other vaccines. Another observation was that the Eastern Mediterranean, where we recorded stable or decreasing correlations with several socioeconomic factors for DTP3, instead showed increasing links with all factors (except investments) for BCG. The Americas, which had decreasing correlations for most factors with DTP3, showed decreases to almost EURpe-like absences of coverage correlation in DTP1 and MCV, although all other regions and factors had similar patterns for DTP3 and MCV coverage. EURpe is broadly the same across all vaccines except BCG, for which correlations seem to be negative between factors such as government health spending, sanitation access, income, and BCG coverage.





Figure A.5: **Historic and recent best-performing correlates with BCG and DTP1 coverage** Time-averaged Spearman's rank correlation values between socio-economic factors and BCG (A) and DTP coverage (B).

A



B




Figure A.6: **Historic and recent best-performing correlates with MCV and POL3 coverage** Time-averaged Spearman's rank correlation values between socio-economic factors and MCV (A) and POL3 coverage (B).

A.3 Vaccine performance index supporting material

A.3.1 Gaussian process forecast performance

The accuracy of Gaussian process regression forecasts of coverage data is compared to linear regression. The best linear regression model for coverage time-series data was found to use 5 prior coverage points to fit to a future point. The Gaussian process regression (GPR) forecast method was found to have an RMSE of 5.5 compared to 6.7 for the benchmark linear regression model.

To visualise the predictive accuracy of GPR, we trained the GPR model on time-series of length 20 across all countries and predicted the 21st data point for the years 1980-1999 through to 1993-2012. These predictions are shown in Figure A.7 below.




Figure A.7: **Forecast accuracy of Gaussian process regression** Observed versus predicted values of the Gaussian process method as described in Section 2.3.1 shows good prediction accuracy across all DTP3 values.

A.3.2 Alternative forms of the Vaccine Performance Index

A strength of the index structure that we propose in the main text is its flexibility. We have chosen a weighting of terms congruent with our picture of desirable vaccination characteristics; readers and policymakers focussed on different aspects of vaccine coverage can re-weight these terms to, for example, focus more strongly on avoiding sudden drops in coverage. To illustrate this re-weighting potential, we consider the alternative indices obtained when setting $d = 5$ and $d = 0$ (with $\nu = 95$ and $\tau = 1$, as before) in Equation (2.11) in the manuscript

for DTP3 coverage. In addition, we also consider only the second term of Equation (2.11), setting $d = 0$, to investigate the probability of a forecasted drop to coverage.

These indices are compared in Figure A.8. Each alternative index is plotted against the original VPI suggested in the manuscript (Equation (2.11) with parameters $d = 2$, $\tau = 1$, and $\nu = 95$) to forecast (and thus obtain a VPI value) for coverage in 2013. In Figure A.8 A and B we display the two alternatives to the VPI in 2013 with $d = 5$ (A) and $d = 0$ (B), referring to the former as $P(> 95) - P(\text{drop} > 5)$ and the latter as $P(> 95) - P(\text{drop})$. Each country's data point is coloured by the WHO region and we label the two countries in each region which see the biggest shift between original VPI – denoted ‘VPI’ in Figure A.8 – and the alternative. (We plot DTP3 coverage time-series of these twelve countries for the first two alternatives in Figure A.9). In Figure A.8 A, we see that using an increased value of a drop serves to increase the VPI values of all countries, with notable increases for countries with lower VPI values, resulting in a tighter grouping of these countries. In Figure A.8 B, however, we observe the opposite effect: rather than an increased similarity of VPI values for many countries, we see that countries with VPI values around +1 now have a range of values under the alternate scheme. Although this may seem beneficial – that there are now a range of values for a group of countries with similarly high VPI values – this alternative formulation causes harsh punishment of countries with high (e.g. with coverage greater than 95%) but steady (e.g. with high probability of obtaining a value at least as big as the previous year) coverage. This punishment arises as, under the Gaussian process model, countries with stable coverage will have a roughly 50% chance of any drop. This is readily evidenced for Finland and Sweden, two countries with exemplary DTP3 coverage rates (Figure A.9 B). This drop in value is – in the case of Sweden and Finland – due to stable coverage levels which result in a symmetric Gaussian probability density over the forecasted point, resulting in a roughly 50% change of drop (despite a low variance). This is motivation for the choice of drop parameter $d = 2$, which does not unfairly punish countries with optimal time-series properties, nor bunches together countries with different time-series properties, as is the case if we increase the value of d above 2.

In Figure A.8 C, we display the original VPI against the probability of a drop between 2012 and 2013 under our Gaussian Process fit. Interestingly, we find that most of the correlation structure is broken, since countries are now no longer rewarded for having high coverage (as they were under the previous schemes), and countries are simply identified by the probability of a drop in coverage. We see clearly that there are a number of EURpean countries which have a large probability of a drop, suggesting either high but volatile coverage, or high but decreasing coverage (or, as previously mentioned, stable coverage rates with an equal chance of increase and decrease).




Figure A.8: **Possible alternatives for the Vaccine Performance Index reveal sensitivity to drop parameter values** VPI values are sensitive to values of the drop parameter d in Equation (2.11). (A) Increasing the value of the drop parameter to $d = 5$ results in higher index values for all countries compared to the original VPI (VPI). (B) Allowing $d = 0$ in Equation (2.11) – resulting in a term representing the probability of a drop of any size – has the effect of decreasing the index values for all countries, including large increases for many countries with with an original VPI value of $\sim +1$. This is an undesirable effect since these countries have stable coverage rates. (C) The probability of a drop $P(\text{drop})$ against our original VPI.

In Figure A.9, we display those countries which have the biggest difference between alternate and original VPI. In comparing the original index with the alternative when $d = 5$, we see that the biggest movers are countries with high recent variability in DTP3 coverage, and thus have a large probability of a drop in DTP3 coverage from the most recent training point (2012) to 2013 (Figure A.9 A). This set of countries may be interpreted as the most volatile that are also most in danger of not reaching the $\nu = 95\%$ threshold. In Figure A.9 B, we display the biggest movers for the second alternative VPI when $d = 0$, we see that most countries would be unfairly punished by this alternative form of VPI, since a time-series with zero gradient would result in a predictive distribution symmetric about the most recent coverage point. This, in turn, results in a roughly evens chance of a drop in coverage. Our original form of the VPI with $d = 2$ mitigates this.

A.3.3 Vaccine performance index values

Vaccine performance index values for the parameterisation $\nu = 95\%$, $d = 2$, and $\tau = 2$ as shown in Figure 2.3 are tabulated here. VPI values for 2001, 2005, 2009, and 2013 are provided.




Figure A.9: **Time-series of countries with the biggest difference between alternative and VPI values** (A) Those countries with the biggest difference (two from each WHO region) for our first alternative (with $d = 5$) and (B) for our second alternative (with $d = 0$).

AFR	2001	2005	2009	2013	EMR	2001	2005	2009	2013	SEAR	2001	2005	2009	2013
Algeria	-0.01	-0.09	-0.38	0.46	Afghanistan	-0.29	-0.52	-0.22	-0.10	Bangladesh	-0.28	0.96	0.45	0.71
Angola	-0.15	-0.19	-0.73	-0.22	Armenia	-0.38	0.15	-0.16	-0.01	Bhutan	0.29	0.14	0.18	0.57
Benin	-0.19	-0.34	-0.19	-0.30	Bahrain	0.94	0.94	0.92	1.00	DPR Korea	-0.26	-0.45	-0.12	0.99
Botswana	1.00	0.94	0.96	0.92	Djibouti	-0.14	-0.56	-0.34	-0.29	India	-0.34	-0.15	-0.64	-0.26
Burkina Faso	-0.15	-0.62	-0.31	-0.11	Egypt	0.93	0.87	0.88	0.17	Indonesia	-0.33	-0.25	-0.03	-0.03
Burundi	-0.32	-0.34	0.28	0.76	Iran	1.00	0.98	0.94	0.97	Maldives	0.77	0.69	0.96	0.90
Cameroon	-0.21	-0.18	-0.17	-0.03	Iraq	-0.30	-0.47	-0.05	-0.40	Burma	-0.20	-0.21	-0.36	-0.35
Cape Verde	-0.03	0.06	0.98	0.19	Jordan	0.71	0.70	0.70	0.76	Nepal	-0.01	-0.16	-0.05	-0.13
CAR	-0.20	-0.36	-0.54	-0.22	Kuwait	0.95	0.82	0.97	0.96	Sri Lanka	1.00	1.00	0.97	0.99
Chad	-0.30	-0.19	-0.21	-0.38	Lebanon	-0.62	-0.21	-0.03	-0.12	Thailand	1.00	1.00	1.00	1.00
Comoros	-0.45	-0.39	-0.63	-0.53	Libya	0.64	0.97	1.00	1.00	Timor Leste	NaN	NaN	-0.02	-0.53
Congo	-0.26	-0.11	-0.46	-0.53	Morocco	0.84	0.69	0.88	1.00	WPR				
Cote d'Ivoire	-0.22	-0.16	-0.06	-0.39	Oman	0.99	0.99	1.00	1.00	Australia	0.04	-0.12	-0.17	-0.07
DRC	-0.44	-0.58	-0.19	-0.21	Pakistan	-0.21	-0.16	-0.02	-0.07	Brunei Dar.	0.54	0.01	0.81	0.09
Equat. Guinea	-0.84	-0.42	-0.54	-0.52	Qatar	-0.78	0.51	0.72	0.44	Cambodia	-0.11	-0.09	-0.38	0.68
Eritrea	-0.31	0.81	0.58	0.62	Saudi Arabia	0.70	0.77	0.90	1.00	China	-0.28	-0.36	0.97	1.00
Ethiopia	-0.45	-0.28	-0.45	-0.32	Somalia	-0.14	-0.08	-0.09	-0.25	Cook Is.	0.23	0.26	0.91	0.48
Gabon	-0.17	-0.24	-0.19	-0.65	Sudan	-0.04	-0.12	-0.14	0.03	Fiji	0.12	0.30	0.99	1.00
Gambia	-0.21	-0.26	0.42	0.96	Syria	-0.15	-0.21	-0.48	-1.00	Japan	-0.31	0.19	0.65	0.86
Ghana	-0.01	-0.06	-0.05	0.02	Tunisia	1.00	0.89	0.98	0.96	Kiribati	-0.12	-0.17	-0.41	-0.04
Guinea	-0.32	-0.21	-0.22	-0.26	UAE	0.36	0.33	-0.09	0.04	Laos	-0.37	-0.33	-0.29	-0.23
Guinea-Bissau	-0.39	-0.38	-0.43	-0.32	Yemen	-0.26	-0.39	-0.38	-0.18	Malaysia	0.99	0.35	0.93	0.78
Kenya	-0.15	-0.21	-0.43	-0.54	EUR					Marshall Is.	-0.37	-0.64	-0.30	-0.26
Lesotho	-0.56	-0.47	-0.02	0.34	Albania	0.70	0.75	0.90	0.99	Micronesia	-0.62	-0.23	-0.62	-0.17
Liberia	NaN	-0.97	-0.45	-0.65	Andorra	-0.23	0.97	0.95	0.98	Mongolia	0.04	1.00	0.46	1.00
Madagascar	-0.14	-0.00	-0.73	-0.24	Austria	-0.00	-0.47	-0.44	-0.40	Nauru	-0.35	-0.44	0.48	0.44
Malawi	-0.32	-0.23	-0.11	0.18	Azerbaijan	-0.07	-0.05	-0.04	-0.08	New Zeal.	-0.27	-0.08	-0.11	-0.05
Mali	-0.40	-0.24	-0.23	-0.18	Belarus	1.00	0.32	0.63	0.57	Niue	0.97	0.93	0.62	0.76
Mauritania	-0.47	-0.50	-0.31	-0.35	Belgium	-0.36	0.03	0.99	1.00	Palau	0.13	0.56	0.64	-0.12
Mauritius	-0.26	0.11	0.95	0.98	Bosnia & Herz.	-0.26	-0.19	0.26	0.03	Papua N.G.	-0.20	-0.20	-0.11	-0.50
Mozambique	-0.17	-0.01	-0.07	-0.09	Bulgaria	0.08	0.07	-0.07	0.11	Philippines	-0.38	-0.06	-0.16	-0.03
Namibia	-0.46	-0.28	-0.14	-0.11	Croatia	-0.00	0.99	1.00	0.99	Korea, Rep.	0.20	0.20	0.33	0.66
Niger	-0.12	-0.07	-0.22	-0.03	Cyprus	0.85	0.99	1.00	1.00	Samoa	0.92	-0.10	-0.26	0.00
Nigeria	-0.26	-0.36	-0.56	-0.12	Czech Republic	0.99	0.99	0.99	1.00	Singapore	0.94	0.73	0.86	0.82
Rwanda	-0.25	-0.22	0.48	0.85	Czechoslovakia	NaN	NaN	NaN	NaN	Solomon Is.	-0.16	-0.20	-0.16	-0.50
ST and Princ.	-0.37	0.20	0.85	0.81	Denmark	0.14	0.73	0.16	0.25	Tonga	0.62	0.32	0.89	0.89
Senegal	-0.23	-0.58	-0.13	-0.02	Estonia	0.12	0.39	0.74	-0.05	Tuvalu	0.18	0.21	0.30	0.25
Seychelles	0.82	0.99	0.99	1.00	Finland	0.99	0.93	0.99	1.00	Vanuatu	-0.20	-0.32	-0.38	-0.48
Sierra Leone	NaN	-0.51	-0.44	-0.13	France	0.88	0.89	1.00	1.00	Vietnam	0.37	0.74	0.40	0.15
South Africa	-0.21	-0.20	-0.43	-0.29	Georgia	-0.27	-0.26	0.25	0.22					
Suriname	-0.26	-0.38	-0.39	-0.26	Germany	-0.05	0.43	0.78	0.81					
Swaziland	-0.30	-0.30	-0.52	-0.40	Greece	-0.33	0.41	1.00	1.00					
Togo	-0.24	-0.26	-0.27	-0.26	Hungary	1.00	1.00	1.00	1.00					
Uganda	-0.01	-0.06	-0.06	-0.19	Iceland	0.99	1.00	0.59	0.02					
Tanzania	-0.17	0.55	-0.20	-0.15	Ireland	-0.11	-0.04	0.01	0.83					
Zambia	-0.21	-0.20	-0.76	-0.17	Israel	0.92	0.93	0.87	0.38					
Zimbabwe	-0.28	-0.44	-0.29	0.95	Italy	0.10	0.35	0.72	0.79					
AMR					Kazakhstan	0.34	0.39	0.68	0.91					
Antigua and B.	0.95	0.89	0.93	0.96	Kyrgyzstan	0.90	0.99	0.18	0.72					
Argentina	-0.07	-0.07	0.11	0.18	Latvia	0.25	0.86	0.88	0.11					
Bahamas	0.09	0.27	0.21	0.74	Lithuania	0.51	0.36	0.98	0.51					
Barbados	0.20	-0.27	-0.03	-0.18	Luxembourg	1.00	1.00	1.00	1.00					
Belize	-0.20	0.23	0.78	0.86	Malta	0.25	0.51	-0.62	1.00					
Bolivia	-0.25	-0.23	-0.11	0.30	Monaco	1.00	1.00	1.00	1.00					
Brazil	0.99	0.29	0.99	0.94	Montenegro	NaN	NaN	NaN	0.47					
Canada	-0.15	-0.12	0.22	0.65	Netherlands	1.00	1.00	1.00	1.00					
Chile	-0.43	0.21	0.29	0.08	Norway	-0.15	0.19	0.08	0.19					
Colombia	-0.47	-0.15	-0.14	-0.42	Palestinian T.	NaN	NaN	NaN	NaN					
Costa Rica	-0.28	-0.32	-0.32	-0.50	Poland	1.00	1.00	1.00	1.00					
Cuba	0.65	0.35	-0.23	-0.15	Portugal	0.95	0.99	0.85	0.90					
Dominica	0.89	0.98	0.66	0.69	Moldova	-0.24	0.95	0.07	-0.01					
Dominican Rep.	-0.27	-0.36	-0.09	-0.10	Romania	0.89	0.97	0.74	-0.20					
Ecuador	-0.19	-0.05	0.99	1.00	Russia	0.66	1.00	1.00	0.97					
El Salvador	0.78	0.09	0.76	-0.00	San Marino	0.89	0.35	0.09	-0.69					
Grenada	0.24	0.58	0.54	0.71	Serbia	0.59	-0.27	0.43	0.28					
Guatemala	-0.36	-0.20	-0.67	-0.03	Slovakia	1.00	1.00	1.00	1.00					
Guyana	-0.05	-0.11	0.12	0.44	Slovenia	-0.23	-0.29	0.80	0.66					
Haiti	-0.19	-0.23	-0.09	-0.00	Spain	0.11	0.96	0.90	0.98					
Honduras	0.73	0.20	0.55	0.26	Sweden	1.00	1.00	1.00	1.00					
Jamaica	0.27	0.98	0.28	0.32	Switzerland	-0.01	-0.03	0.42	0.99					
Mexico	0.92	1.00	0.97	1.00	Tajikistan	-0.24	-0.09	-0.11	0.39					
Nicaragua	-0.16	-0.19	0.99	0.99	Macedonia	0.59	0.56	0.54	0.68					
Panama	0.58	0.85	-0.18	-0.30	Turkey	-0.42	-0.50	0.20	0.79					
Paraguay	-0.22	-0.32	-0.34	-0.26	Turkmenistan	0.36	0.21	0.68	0.70					
Peru	0.99	0.11	0.29	-0.06	Ukraine	1.00	0.98	0.44	-0.19					
St. Kitts & N.	0.85	0.75	0.94	0.65	Former USSR	NaN	NaN	NaN	NaN					
Saint Lucia	-0.09	-0.03	0.01	0.40	UK	-0.01	-0.01	-0.01	1.00					
St. Vinc & Gren	0.77	0.99	0.99	0.86	Uzbekistan	0.73	0.86	0.72	1.00					
Trinidad & Tob.	0.00	0.27	-0.07	-0.11										
USA	-0.37	0.72	0.91	0.91										
Uruguay	0.06	0.16	0.15	0.66										
Venezuela	-0.73	-0.89	-0.15	-0.47										

Table A.2: Vaccine Performance Index values VPI values for the parameterisation $\nu = 95\%$, $d = 2$, and $\tau = 2$ as shown in Figure 2.3.

B | Appendix B

B.1 Hierarchical linear model assumption checks

The linear model assumptions outlined in Section 3.3.2 are applied to the hierarchical linear models given in Section 3.3. As logistic hierarchical regression is implemented, we need only check the residuals for the random-effect coefficients β_{0j} in Equation (3.4).

In Figure B.1 below, we verify the assumption of normality of residuals ν_{0j} (in Equation (3.4)), the homoscedasticity of residuals, and check for no autocorrelation in the residuals. This is done for the hierarchical linear model for vaccine importance in Figure B.1 A and B (the same checks for the other three models for safety, effectiveness, and religious compatibility all similarly pass these checks). A QQ-plot of residuals ν_{0j} reveals residual normality (A). A plot of these residuals against predicted values shows homoscedasticity and a lack of autocorrelation in the range -3.0 to 4.0 (B). At values higher than 4.0, residuals tend to be negative for a small number of countries.




Figure B.1: **Validation of linear model assumptions for Equation 3.4.** Level-2 residuals ν_{0j} are tested for normality, homoscedasticity, and for a lack autocorrelation. (A) QQ-plot of the residuals ν_{0j} reveals residual normality. (B) A residual plot against predicted values reveals homoscedasticity and no autocorrelation in the range -3.0 to 4.0, with possibly a minor violation of these two assumptions above this range.

B.2 Vaccine effectiveness versus vaccine safety

In Figure B.2, we plot the difference between the proportion of respondents reporting positively (tend to agree or strongly agree) to “overall I think vaccines are effective” than to “overall I think vaccines are safe.” A similar result is found when comparing the difference in proportions of individuals replying positively to “overall I think vaccines are important” relative to “overall I think vaccines are safe”. However, unlike the importance-safety relationship, a handful of countries (Algeria, Nigeria, Ghana, Pakistan, Indonesia, and Bangladesh) are more likely to report that vaccines are safe than effective.




Figure B.2: **Differences in the proportion of people responding that they believe vaccines are effective but unsafe** The difference in proportion of people within each country who reply positively to vaccine effectiveness and positively to vaccine safety.

C | Appendix C

C.1 Sensitivity analysis and simulation study

Sensitivity analyses are performed to select appropriate hyper-parameters for the variance components of the hierarchical structural model in Equations (4.15) and to assess the robustness of item response parameter inference with respect to the non-informative Wishart prior in Equation (4.14).

C.1.1 Data generation

A total of 4,000 observations are generated across 10 items (with $K = 5$) from 20 groups each comprising of 200 individuals. Data is generated according to the hierarchical model outlined in Section 4.4.1. The random effects parameters β_j are first simulated through $\beta_j \sim \mathcal{N}(\gamma \mathbf{Z}_j, \Sigma)$, which are then used to obtain hesitancy values through $\eta_{ij} \sim \mathcal{N}(\beta_j \mathbf{X}_{ij}, \sigma^2)$. The item response parameters are generated from the exchangeable multivariate distribution $(\tilde{a}_m, \mathbf{b}_m)^T \sim \mathcal{N}(\boldsymbol{\mu}_{\text{item}}, \Sigma_{\text{item}})$ ($a_m = \exp(\tilde{a}_m)$) and the response data can be finally generated through a combination of the latent data and item response parameters via the graded response model (Equation (4.3)). Unless otherwise specified, the choices of model parameters are as follows: the parameter matrix γ is specified as,

$$\gamma = \begin{bmatrix} 0 & -\frac{1}{2} \\ 1 & \frac{1}{2} \\ 1 & -2 \end{bmatrix};$$

covariate data $\mathbf{X}_{ij}$ and $\mathbf{Z}_j$ are sampled from independent standard normal distributions; $\sigma^2 = 1$; $\Sigma = 2 \times \mathbf{1}$; and the item response parameters are generated from $\boldsymbol{\mu}_{\text{item}} = [0, 0, 0, 0, 0]$ and $\Sigma_{\text{item}}^{-1} = \mathbf{1}_{5 \times 5}$. As the identifying restrictions (see Section 4.3.3) on these item response parameters fixes the product of the discrimination parameters to one and the sum of the difficulty parameters to zero, the latent variables will have – approximately – a standard deviation of one and a mean close to zero. These data represent typical parameter values for

an IRT model.

C.1.2 Choosing individual- and country-level variance hyper-priors

Sensitivity analyses are performed to examine model robustness of variance component inference under varying hyper-parameters of the Gamma and Wishart priors for individual- and country-level variances respectively. We investigate choosing different hyper-parameters, α and β , for the prior distributions $\text{Gamma}(1, \alpha)$ and $\text{Wishart}(\beta \times \mathbb{1}, n_x + 1)$.

Inference robustness is tested by varying over individual-level variance values, $\sigma^2 = .01, .1, 1, 10$, and 20 , whilst choosing $\alpha = 1, .1$, and $.01$ to change the mean and variance of the distribution. Similarly, for the country-level covariance prior, the degrees of freedom is fixed at $n_x + 1$ (a minimally informative choice [246]), and $\beta = .01, .1$, and 1 is varied. These

	$\sigma^{-2} = .01$			$\sigma^{-2} = .1$			$\sigma^{-2} = 1$			$\sigma^{-2} = 10$			$\sigma^{-2} = 20$		
$\alpha =$	1	.1	.01	1	.1	.01	1	.1	.01	1	.1	.01	1	.1	.01
$\bar{\sigma}^{-2}$	.010	.010	.010	.096	.097	.097	0.80	0.90	0.90	1.99	8.82	14.5	1.86	13.89	19.59
$\text{sd}(\sigma^{-2})$	.000	.000	.000	.003	.003	.003	.06	0.07	0.08	.20	1.00	5.82	0.18	3.81	5.30

	$\Sigma_{..}^{-1} = .01$			$\Sigma_{..}^{-1} = .1$			$\Sigma_{..}^{-1} = 1$			$\Sigma_{..}^{-1} = 10$			$\Sigma_{..}^{-1} = 20$		
$\beta =$	1	.1	.01	1	.1	.01	1	.1	.01	1	.1	.01	1	.1	.01
$\bar{\Sigma}_{11}^{-1}$	.008	.011	.012	0.08	0.09	0.09	1.18	2.20	2.42	2.85	8.74	12.30	2.47	10.02	25.30
$\bar{\Sigma}_{22}^{-1}$	.010	.013	.014	0.18	0.25	0.26	0.58	0.75	0.78	2.28	5.35	6.86	2.60	11.98	34.90
$\bar{\Sigma}_{33}^{-1}$	.008	.011	.011	0.09	0.11	0.11	1.00	1.63	1.74	3.04	10.96	17.29	2.62	12.38	35.54
$\text{sd}(\Sigma_{11}^{-1})$	.003	.004	.004	0.02	0.03	0.03	0.37	0.70	0.78	0.89	2.90	4.47	0.81	3.25	10.14
$\text{sd}(\Sigma_{22}^{-1})$	.003	.004	.004	0.06	0.08	0.08	0.18	0.24	0.24	0.72	1.76	2.42	0.85	3.91	14.01
$\text{sd}(\Sigma_{33}^{-1})$	.002	.003	.003	0.03	0.03	0.03	0.31	0.52	0.56	0.96	3.75	6.74	0.85	4.03	14.71

Table C.1: **A simulation study investigates the inference accuracy when varying hyper-parameterisation of the variance and covariance priors** Inferring individual (top table) and country-level (bottom-table) variance and covariance (respectively) components (with standard errors) under different hyper-parameterisations for the Gamma and Wishart priors. The most robust estimates occur when $\alpha = \beta = 0.01$, with these parameters as described in the text.

parameters are investigated for $\Sigma^{-1} = .01 \times \mathbb{1}, 0.1 \times \mathbb{1}, 1 \times \mathbb{1}, 10 \times \mathbb{1}$, and $20 \times \mathbb{1}$.

In Table C.1, we show the inferred variance and covariance components under these parameterisations of prior distributions. We find that in both cases, $\alpha = \beta = 1$ results in prior distributions which are too heavily weighted towards low precisions (high variance), whereas $\alpha = \beta = .01$ and $\alpha = \beta = .1$ perform well across the range of precision values. As $\alpha = \beta = .01$ tends to be more consistent over a range of precision values, these values are chosen as hyper-parameters for the Gamma and Wishart priors (which have been used in Equations (4.15)).

C.1.3 Testing model robustness to extreme item response parameters

A range of item response parameters are generated according to the data-generation mechanism outlined in Section C.1.1. Item response parameters are generated under three distributions: $(\log(a_m), \mathbf{b}_{m,k})^T \sim \mathcal{N}(\boldsymbol{\mu}_{\text{item}}, \Sigma_{\text{item}})$, where $\boldsymbol{\mu}_{\text{item}} = (\mu, \mathbf{0})^T$ with $\mu = 0, 1$, and 2 , and $\Sigma_{\text{item}} = s\mathbb{1}$ with $s = 1, 2$, and 5 (for $m = 1, \dots, 10$ and $k = 1, \dots, 4$ (as described in Section C.1.1)). These sets of item response parameters represent what we might expect to be ‘typical’ parameters ($\mu = 0, s = 1$), ‘unlikely’ parameters ($\mu = 1, s = 2$), and ‘extremely unlikely’ parameters ($\mu = 2, s = 5$).

We find that the choice of Wishart prior in Equation (4.14) is appropriate in this case as it allows accurate inference under these broad range of possible parameter values. We show simulated versus inferred values under these three parameter regimes in Figure C.1.

C.1.4 Full parameter inference and MCMC convergence

A full simulation study with the parameter values given in Section C.1.1 is conducted to evaluate the performance of the MCMC sampler and to ensure high accuracy in parameter recovery. In Figure C.2 we plot the generated versus observed values for the range of model parameters. Across all inferred parameters we note a very high accuracy in parameter recovery for discrimination (A), difficulty (B), fixed-effect (C), random-effects (D), and variance (E) parameters. Across all parameters, 96% of the inferred parameter estimates fall within the 95% HPD intervals in line with expectations.

A total of 126 model parameters were inferred (excluding the latent hesitancy values) and only eight (6%) were found to have a Geweke p -value less than 0.05. This percentage is in line with what we would expect by chance if all the Geweke p -values followed the null distribution. Nonetheless, inspecting these eight parameters shows good convergence (Figure C.3).




Figure C.1: **Robustness of the Wishart prior for item response parameter inference**
The Wishart prior over item response parameters is a suitable choice. A wide range of possible item response parameters are accurately inferred under this prior. (A) Simulated (true) versus inferred (pred) values for the discrimination parameters. (B) Simulated versus inferred values for the difficulty parameters. (C) Simulated versus observed values across all parameters. Black points with 95% HPD intervals represent the simulated (pred) data, and coloured dots represent the simulated (true) values (magenta = ‘likely’, red = ‘unlikely’, blue = ‘extremely unlikely’ – as described in text.)




Figure C.2: **Inferred versus simulated parameter values shows robust inference under our implemented MCMC scheme** Inferred versus simulated parameter estimates shows good parameter recovery precision. Predicted versus true values are shown for (A), difficulty (B), fixed-effect (C), random-effects (D), and variance (E) parameters.




Figure C.3: MCMC Gibbs samples for parameters whose Geweke statistic failed the test for convergence Gibbs samples for all eight parameters (see main text) whose Geweke statistic indicated no convergence. Despite these significant p -values, the fraction of un-converged time-series is what we would expect by chance and these time-series show good convergence.

C.2 Verifying model assumptions for Bayesian hierarchical IRT

C.2.1 Testing and verifying the conditional independence assumption

The conditional independence assumption for IRT and LCA states that responses to individual items are conditionally independent given latent attribute. This was stated for IRT in Equation (4.5) as,

$$P(Y_{1i}, Y_{2i}, \dots, Y_{Mi} | \eta_i) = \prod_{m=1}^M P(Y_{im} | \eta_i), \quad (\text{C.1})$$

where the dependence on the item response parameters has been removed. This strong form of conditional independence requires testing conditional independence on all subsets of items, which can be cumbersome to compute. A weaker form of conditional independence exists, which requires conditional independence between all pairs of items (thus fewer tests must be satisfied),

$$P(Y_{im}, Y_{il} | \eta_i) = P(Y_{im} | \eta_i) P(Y_{il} | \eta_i), \quad (\text{C.2})$$

for $\forall m, l$ ($m \neq l$) [288]. When satisfied, this pairwise independence is often indicative of whether the stronger form is met, as although higher order dependencies are possible they are often unlikely [288].

In seeking to establish conditional independence, it is instructive to consider the mutual information (MI), a information theoretic measure of dependence between variables which makes no distributional assumptions [190]. The MI quantifies the amount of information obtained about one random variable X through knowledge of the other Y (here we use X and Y as our vaccination response survey data),

$$I(X; Y) = \sum_{y \in Y} \sum_{x \in X} P_{X,Y}(x, y) \log_2 \left(\frac{P_{X,Y}(x, y)}{P_X(x) P_Y(y)} \right).$$

When X and Y are independent, it is easy to see that $I(X; Y) = 0$. The mutual information can be readily extended to test conditional independence between two variables. The conditional mutual information (CMI) is given by

$$\begin{aligned} I(X; Y | Z) &= \mathbb{E}_Z(I(X; Y) | Z) \\ &= \int P_Z(z) \sum_{y \in Y} \sum_{x \in X} P_{X,Y|Z}(x, y | z) \log_2 \left(\frac{P_{X,Y|Z}(x, y | z)}{P_{X|Z}(x | z) P_{Y|Z}(y | z)} \right) dz, \end{aligned}$$

where it is assumed that X and Y are discrete random variables (corresponding to observed multivariate survey response data) and Z is a continuous random variable (here Z is the

inferred latent vaccine hesitancy). Analogously to MI, two variables are conditionally independent when $I(X; Y|Z) = 0$.

Estimation of this quantity (in the case of IRT) relies on discretising the continuous random variable Z and approximating the integral through numerical integration techniques. The CMI is used to verify the conditional independence assumption through the G -statistic,

$$G = 2N\hat{I}(X; Y|Z), \quad (\text{C.3})$$

which is asymptotically χ^2 -distributed under the null with $\nu = n_{\text{item}}^2 n_{\text{bins}}$ degrees of freedom, where n_{bins} is the number of bins used in the discretisation process and n_K is the number of possible survey responses ($n_K = 5$) [289].

This G statistic is implemented for all pairs of survey questions given the inferred hesitancy values. For each pair of vaccination questions, we accept the null distribution when the binning process is sufficiently small. Thus we find that each pair of vaccination questions is conditionally mutually independence given vaccine hesitancy and our assumption is verified.

C.2.2 Testing successful burn-in and convergence

To test successful burn-in and convergence of our implemented MCMC scheme, we examine the Geweke statistics for each inferred parameter in the posterior distribution as described in Section (4.4.4).

For model $\mathcal{M}_4$, 5.5% of Geweke p -values are less than 0.05. Further, we find that these p -values follow a uniform distribution, suggesting that we should accept the alternative hypothesis of sufficiently burnt-in and converging chains across our parameter set (Figure C.4). Similar distributions are found across all other models. Visual inspection of those samples with p -value < 0.05 show good convergence.

C.2.3 Verifying linear model assumptions

Linear model assumption checks of residual normality, homoscedasticity, and uncorrelatedness are now shown for model $\mathcal{M}_4$ which was introduced in Section 4.4.3. These checks were initially outlined in Section 3.3.2. (We note that similar checks of the linear model assumptions for other models outlined in Section 4.4.3 reveal similar behaviours to those presented here for model $\mathcal{M}_4$).

In Figure C.5 A a QQ-plot of the residuals from the individual-level regression in Equation (4.13) reveals a slightly non-normal distribution in the tails of the residual distribution. In Figure C.5 B, we plot residuals against predicted values (due to the large number of residuals in this analysis, residuals are binned at regular intervals and box-and-whisker plots at each




Figure C.4: **Distribution of Geweke p -values of all inferred parameters in Model $\mathcal{M}_2$** The distribution of Geweke p -values appears to follow a uniform distribution which suggests good convergence of MCMC chains of model parameters.

bin are used to show the variance at a given point). Figure C.5 B reveals homoscedasticity and uncorrelated residuals over the majority range of the predicted values; however, there seems to be some violation of these assumptions at the upper extremes of the predicted values range. Item response theory maps discrete values onto a continuous scale. In the case when the number of items is limited (such as in our analysis, where we consider only four items), discretisation effects will manifest in the latent variable (see for example the distribution of individuals' hesitancy in Figure 4.2 A, for example). These effects could result in violation of these linear model assumptions. We note that although these model violations appear minor, a discrete modelling approach to vaccination survey items may be more suitable – see Chapter 5.

In Figure C.6 we show that the country-level regression for $\mathcal{M}_4$ (see Section 4.4.3) obeys the linear model assumptions. In Figure C.6 A a QQ-plot reveals normality of the country-level residuals and Figure C.6 B shows that the residuals are uncorrelated and homoscedastic.




Figure C.5: **Validation of linear model assumptions for individual-level residuals from Model $\mathcal{M}_4$** The individual-level residuals are tested for normality, homoscedasticity, and for a lack autocorrelation. (A) QQ-plot of the residuals reveals non-normality at the tails of the distribution. (B) A residual plot against predicted values shows homoscedastic residuals and no autocorrelated residuals across the majority range of the predicted values. There is some violation of the model assumptions due to discretisation effects discussed in the text.




Figure C.6: **Validation of linear model assumptions for country-level residuals from Model $\mathcal{M}_4$** The country-level residuals ν_{0j} are tested for normality, homoscedasticity, and for a lack autocorrelation. (A) QQ-plot of the residuals ν_{0j} reveals residual normality. (B) A residual plot against predicted values reveals homoscedasticity and no autocorrelation.

C.3 JAGS code for Bayesian hierarchical IRT models

Some readers may find our JAGS code useful for the Bayesian hierarchical IRT model we introduced in Chapter 4. We provide an outline of it here. We note the following: `n.group` is the number of countries ($J = 72$ in our case); `g.idx` is a vector marking the individuals within a group; `n.item` is the number of items ($M = 4$); `n.cat` is the number of item response categories ($K = 5$); `disc` and `diff` are the discrimination and difficulty parameters, respectively; `diff.raw` are the difficulty parameters before the model identification process; `eta.tau` is the individual-level precision ($1/\sigma^2$); `beta.tau` is the country-level precision matrix (Σ^{-1} in Equation (4.13)); `xi.mu.hyp` and `xi.tau.hyp` are the hyper-priors for the item response parameters (see Section 4.4.2 noting that the precision matrix is again used); and `beta` and `gamma` are the random- and fixed-effect parameters. Unless otherwise stated, values in parantheses correspond to the variables in Section 4.4.3.

```
##### model #####
# polytomous graded-response IRT model
for (j in 1:n.group){
  for (i in (g.idx[j]+1):g.idx[j+1]){
    for (m in 1:n.item){
      Y[i, m] ~ dcat(pr[i, m, ])
      Q[i, m, 1] <- 1
      for (k in 2:n.cat){
        logit(Q[i, m, k]) <- disc[m]*(eta[i] - diff.raw[m, k])
      }
      pr[i, m, 1:(n.cat-1)] <- Q[i, m, 1:(n.cat-1)] - Q[i, m, 2:n.cat]
      pr[i, m, n.cat] <- Q[i, m, n.cat]
    }
    # individual-level regression
    eta.mu[i] <- inprod(X[i, ], beta[, j])
    eta[i] ~ dnorm(eta.mu[i], eta.tau)
  }
  # country-level regression
  beta.mu[1:n.x, j] <- gamma %*% Z[1:n.z, j]
  beta[1:n.x, j] ~ dmnorm(beta.mu[1:n.x, j], beta.tau)
}

# item difficulty and discrimination
for (m in 1:n.item){
  xi[m, 1:n.cat] ~ dmnorm(xi.mu, xi.tau)
```

```

# identify the location
diff.raw[m, 2:n.cat] <- ifelse(m==1, sort(xi[1, 2:n.cat])
  - xi.mean, sort(xi[m, 2:n.cat]) - xi.mean)
diff.raw[m, 1] <- -10^30
diff.raw[m, n.cat+1] <- 10^30
}
diff[1:n.item, 1:(n.cat-1)] <- diff.raw[1:n.item, 2:n.cat]
xi.mean <- mean(xi[1, 2:n.cat])

# log-normal distribution for discrimination
for (m in 1:(n.item-1)){
  disc[m] <- exp(xi[m, 1])
}
# identify the scale
disc[n.item] <- 1/prod(disc[1:(n.item-1)])

##### prior information #####
eta.tau ~ dgamma(1,0.01)
beta.tau ~ dwish(0.01*beta.tau.hyp, 5)
xi.mu ~ dmnorm(xi.mu.hyp, xi.tau)
xi.tau ~ dwish(xi.tau.hyp, xi.nu.hyp)
for (x in 1:n.x){
  for (z in 1:n.z){
    gamma[x, z] ~ dnorm(0, 0.1)
  }
}

```