Multi-scale immune selection and the maintenance of structured antigenic diversity in the malaria parasite *Plasmodium falciparum*

Author:

Thomas Holding

Supervisor:

Dr. Mario Recker

Submitted by Thomas Mitchell Holding, to the *University of Exeter* as a thesis for the degree of *Doctor of Philosophy in Mathematics* in January 2018

Signature: _____

Abstract

The most virulent malaria parasite, *Plasmodium falciparum*, makes use of extensive antigenic diversity to maximise its transmission potential. Parasite genomes contain several highly polymorphic gene families, whose products are the target of protective immune responses. The best studied of these are the PfEMP1 surface proteins, which are encoded by the *var* multi-gene family and are important virulence factors. During infection, the parasite switches expression between PfEMP1 variants in order to evade adaptive immune responses and prolong infection. On the population level, parasites appear to be structured with respect to their *var* genes into non-overlapping repertoires, which can lead to high reinfection rates. This non-random structuring of antigenic diversity can also be found at the level of individual *var* gene repertoires and *var* genes themselves. However, not much is known about the evolutionary determinants which select for and maintain this structure at different ecological scales. In this thesis I investigate the mechanisms by which multi-scale immune selection and other ecological factors influence the evolution of structured diversity. Using a suite of theoretical frameworks I show that treating diversity as a dynamic property, which emerges from the underlying infection and transmission processes, has a major effect on the relationship between the parasite's transmission potential and disease prevalence, with important implications for monitoring control efforts. Furthermore, I show that an evolutionary trade-off between within-host and between-host fitness together with functional constraints on diversification can explain the structured diversity found at both the repertoire and parasite population level and might also account for empirically observed exposure-dependent acquisition of immunity. Together, this work highlights the need to consider evolutionary factors acting at different ecological scales to gain a more comprehensive understanding of the complex immune-epidemiology of *P. falciparum* malaria.

Acknowledgements

I am immensely grateful for all the support, insight and distractions offered by my friends and colleagues in the Tremough House mathematics corridor. In particular, I would like to thank Mario for his supervision and patience. My thanks also go to Marina for leading the way - I got there in the end, and of course JJ and Warren for being fantastic office mates.

Madeleine, thank you for your support and understanding, and especially for putting up with me during the final weeks of writing up. Living with me during this process definitely can't have been easy, and I deeply appreciate it all.

I would like to thank my parents and my sister, Ellen, for all the love and encouragement they have always shown me. Without all of your support this thesis would not exist. Finally, I would like to thank W100, Z100, Webless, Despot, Blanca, Bald-Eyes, Little and Rollo for everything they've taught me.

Contents

List of Abbreviations

AIC   Akaike Information Criterion

CIDR   Cysteine-Rich Interdomain Region

CSA   Condroitin Sulphate A

DBL   Duffy Binding Like

DOI   Duration of infection

EIP   Extrinsic Incubation Period

EIR   Entomological Innoculation Rate

EPCR   Endothelial Protein C Receptor

ICAM-1   Intercellular Adhesion Molecule-1

PAM   Pregnancy Associated Malaria

PfEMP1   *Plasmodium falciparum* Erythrocyte Membrane Protein 1

SIER   Susceptible-Infected-Exposed-Recovered (modelling framework)

SIS   Suseptible-Infected-Susceptible (modelling framework)

SIR   Susceptible-Infected-Recovered (modelling framework)

SIRS   Susceptible-Infected-Recovered-Susceptible (modelling framework)

TC   Test Curve

Ups   Upstream Promotor Sequence

List of Symbols

Exposure independent transmission model:

S Proportion of host population that are susceptible

I Proportion of host population that are infectious

β Transmission rate

σ Recovery rate

μ birth/death rate

Temporary immunity model:

S Proportion of host population that are susceptible

I Proportion of host population that are infectious

R Proportion of the host population that are recovered

β Transmission rate

σ Recovery rate

μ Birth/death rate

α Rate of loss of temporary immunity

Temporary immunity model:

S_1 Proportion of non-immune host population that are susceptible

I_1 Proportion of host non-immune population that are infectious

S_2 Proportion of host semi-immune population that are susceptible

I_2 Proportion of host semi-immune population that are infectious

β_1 Transmission rate (non-immune)

σ_1 Recovery rate (non-immune)

σ_2 Recovery rate (semi-immune)

λ Rate of acquisition of semi-immune state

γ Proportion change in transmission rate for semi-immune hosts

μ Host birth/death rate

α Rate of loss of temporary immunity

Cumulative exposure model:

S_i Proportion of susceptible hosts that have i previous infections

I_i Proportion of infectious hosts that have i previous infections

R Proportion of hosts with immunity all circulating parasite diversity

β_i Transmission rate for hosts with i previous infections

σ_i Recovery rate in hosts with i previous infections

μ Host birth/death rate

n Parasite diversity (number of exposure levels])

α Strength of cross-immunity

λ_i Force of infection for a sub-population with i previous infections

ϵ_i Incidence rate for a sub-population with i previous infections

Dynamic diversity model:

A_{max} Total size of antigen space

a Denotes an antigen ID

g Denotes a gene

b	Mosquito daily biting rate
n	Number of antigen ID encoding bits
H	Number of human hosts
M	Number of mosquitoes
p_{trans}	Probability of mosquito becoming infected after biting an infectious human
p_c	Per-gene probability of mitotic recombination
p_s	Per-gene probability of meiotic recombination
α	Baseline per-antigen contribution to infection length
γ	Strength of cross-immunity
θ	Scaling coefficient for change due to mitotic recombination
I_{init}	Number of initially infected mosquitoes
A_{init}	Initial pool of antigens
S_{init}	Initial number of strains
μ_M	Mean lifespan of mosquitoes in days
R	Antigen repertoire size
$\mathbf{h}$	Host immune status vector
$\mathbf{s}$	Strain vector (defines antigen repertoire)
Ω	Host susceptibility
$\omega(\mathbf{h}, \mathbf{s})$	Duration of infection in host ($\mathbf{h}$) by strain ($\mathbf{s}$)

Extended dynamic diversity movel:

$\mathbf{g_A}$	Strain vector component for antigen group A
$\mathbf{g_B}$	Strain vector component for antigen group B
$\mathbf{h_A}$	Immune vector component for antigen group A
$\mathbf{h_B}$	Immune vector component for antigen group B

p_g	Per-gene probability of mitotic recombination resulting in a change in group
p_{clear}	Probability that an infection is terminated early (per antigen expressed)
N_A	Size of antigen space containing group A antigens
N_B	Size of antigen space containing group B antigens
N_A^*	Initial antigen pool for group A antigens
N_B^*	Initial antigen pool for group B antigens
Δb	Change in biting rate (seasonality)

Repertoire structure model:

l_0	Baseline antigen contribution to infection length
l_k	Contribution to infection length made by the k^{th} antigen
R	Repertoire size
N_A	Number of group A antigens in the initial antigen pool
N_B	Number of group B antigens in the initial antigen pool
A	Number of antigens in an antigen repertoire which are of group A
B	Number of antigens in an antigen repertoire which are of group B
$\mathbf{s_A}$	Antigen repertoire (portion containing group A antigens)
$\mathbf{s_B}$	Antigen repertoire (portion containing group B antigens)
$\mathbf{h_A}$	Host immunity to group A antigens
$\mathbf{h_B}$	Host immunity to group B antigens
γ_A	Strength of cross-immunity in antigen group A
γ_B	Strength of cross-immunity in antigen group B
λ	transmission rate
p_{inf}	Probability of infection
s	Denotes a strain

h   Denotes a host

Clustering strains (Chapter 5):

$S_{i,j}$ Similarity between parasites i and j

A_i Set containing the antigen IDs for each antigen in strain i's repertoire

C_j A set of strains contained in cluster j

k Total number of clusters

Chapter 1

Introduction

1.1 Overview

Malaria is one of the most important and devastating human diseases. Although incidence has been decreasing in recent decades, the World Health Organization estimates there were over 200 million cases causing over 400,000 deaths in 2015 alone (World Health Organization, 2016). Over two billion people remain at risk of infection across sub-Saharan Africa, South East Asia and South America (Guerra et al., 2008). The most virulent human malarial parasite is *Plasmodium falciparum*, which is responsible for most of the malaria related deaths in Africa. *P. falciparum's* transmission strategy makes extensive use of antigenic diversity enabling it to reinfect hosts throughout their lives and establish long infections. The work described in this thesis investigates the epidemiological impact of this antigenic diversity. I focus on the antigen *Plasmodium falciparum* Erythrocyte Membrane Protein 1 (PfEMP1), which exhibits diversity both at the scale of parasite populations and within individual genomes. PfEMP1 is a target of naturally acquired immunity and is a key determinant of virulence (described in detail in section 1.2.2). I consider the effect that PfEMP1 diversity, and the structure contained within this diversity, has on

transmission, prevalence, immunity and epidemiological patterns in disease severity. Additionally, the evolutionary processes which can select for observed structure in diversity and the epidemiological consequences of constraints on diversity are explored.

1.2   Biological background

P. falciparum is transmitted to the human host by female *Anopheles* mosquitoes, which transfer sporozoites when taking a blood meal. Once in the blood, sporozoites migrate to the liver and undergo asexual replication. After several days thousands of merozoites are released into the bloodstream, beginning the blood stage of infection. Merozoites invade erythrocytes and further replicate inside them, forming a cycle of invasion, replication and egress. A small proportion of merozoites differentiate into male and female gametocytes and remain in erythrocytes, thus evading host immunity. When a mosquito bites an infected human and ingests a blood meal, gametocytes are taken up with the erythrocytes. This begins the sexual phase of the parasite's life-cycle, the end result of which is the production of new sporozoites, which migrate to the mosquito salivary gland ready for onward transmission. The lifecycle of *P. falciparum* is summarised in figure 1.1.

Transmission intensity ranges from sporadic outbreaks and seasonal transmission to year-round transmission, with prevalence as high as 90% in some regions. Figure 1.2 shows the prevalence in children for over 3600 sites across sub-Saharan Africa. The relationship between transmission intensity and prevalence is highly non-linear. Work by Beier *et al.* (Beier, Killeen, and Githure, 1999) showed that rapid increases in prevalence occur from low to moderate transmission intensity but that
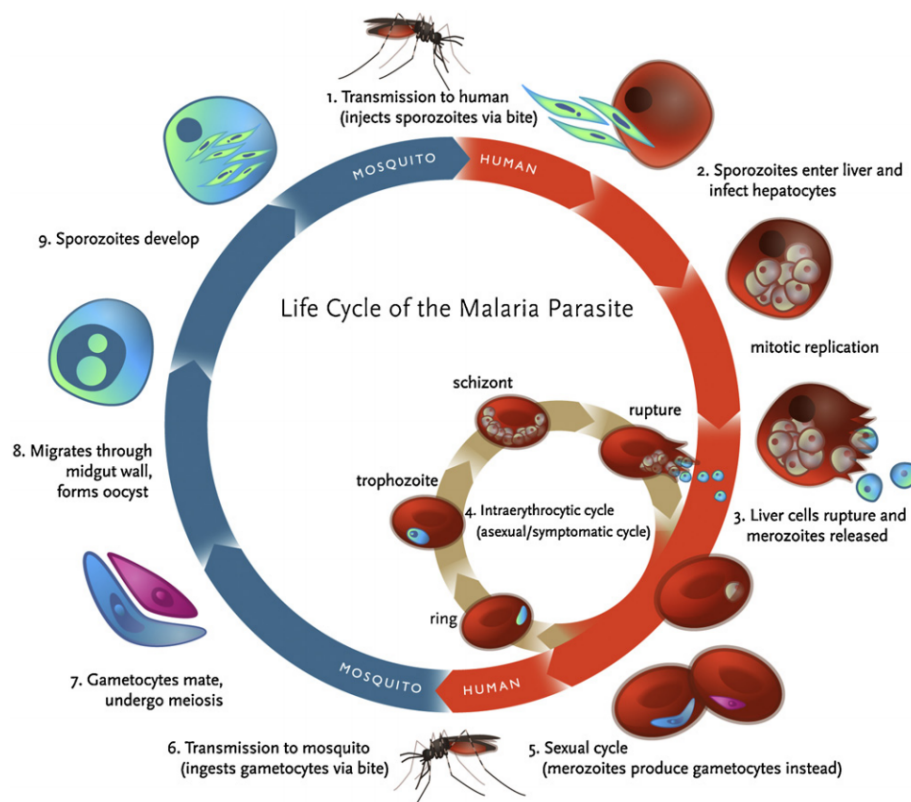
FIGURE 1.1: The life-cycle of *P. falciparum* requires transmission to both a human host and a an *Anopheles* mosquito. Once delivered to a human's bloodstream by a bite from an infectious mosquito, sporozoites migrate to the liver where they undergo rapid replication. After replication thousands of merozoites are released from the liver and re-enter the blood stream in order to begin a cycle of invading erythrocytes, replication and egress. Figure taken from the Klein *et al* (Klein, 2013).

prevalence plateaus at higher transmission intensity. The data from this study is reproduced in figure 1.3, and includes data from 30 field sites reported by different studies. For each location, sampling took place at least monthly for at least one year. Mosquitoes were collected variously by human-biting catching, pyrethrum spray catching and/or light traps. All samples were taken when no mosquito-control programmes were in operation at the time. The infection status of mosquitoes were determined by dissection or enzyme-linked immunosorbent assay. Human prevalence was determined by thick/thin film blood smear microscopy and diagnosed to the level of *Plasmodium* species.

With increasing transmission intensity there is also an increase in the occurrence
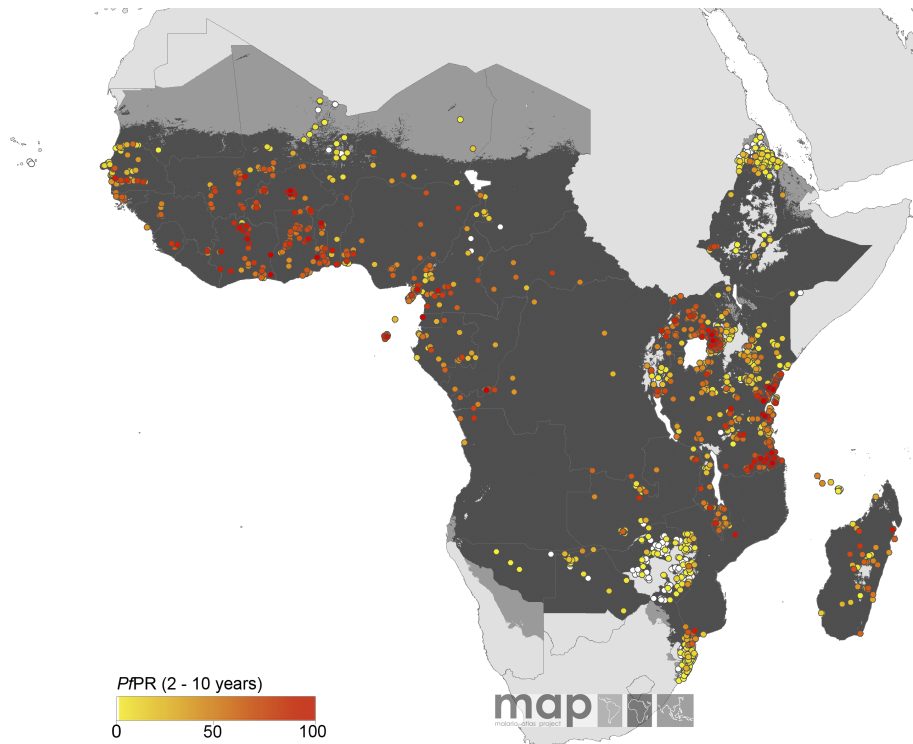
FIGURE 1.2: *P. falciparum* prevalence in children at over 3600 sites across sub-Saharan Africa. Taken from the Malaria Atlas project, Guerra et al (Guerra et al., 2008).

of concurrent infections, in which hosts can become infected by more than one genetically distinct parasite. This can result from multiple bites by different infected mosquitoes, or a single bite from a mosquito carrying multiple parasite genotypes. Concurrent infections are common even in relatively low transmission regions and may be the dominant mode of infection in high transmission areas (Smith et al., 1999c; Juliano et al., 2009; Juliano et al., 2010; Peyerl-Hoffmann et al., 2001; Mayor et al., 2003; Beck et al., 1997).
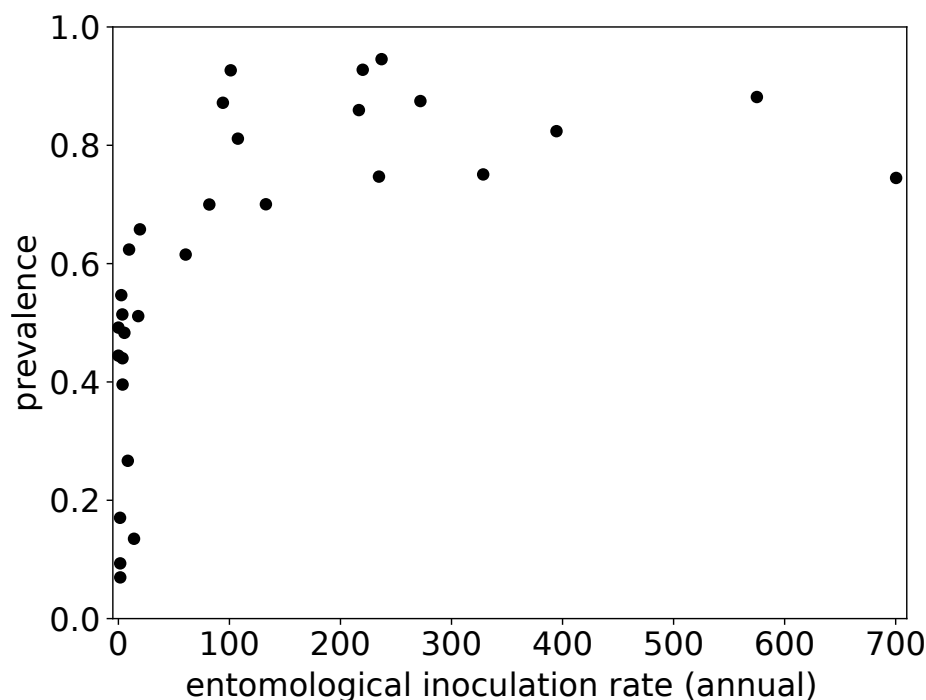
FIGURE 1.3: The relationship between transmission intensity and prevalence for African 30 field sites. Reproduced from Beier *et al* (Beier, Killeen, and Githure, 1999). Data was collated from different studies which each took place for a minimum of one year.

1.2.1 Burden of disease

Very high prevalence rates are possible because of the parasite's ability to reinfect individuals throughout their lives. While sterile immunity never develops and human hosts remain susceptible throughout their lives, it is clear that repeated exposure induces immunological changes in hosts that offer some protection from disease. The clinical presentation of malaria ranges from severe and life-threatening disease to asymptomatic infection. Symptoms include fever, headache, anaemia, nausea and muscle pain. Severe cases are characterised by complications including respiratory distress, severe anaemia, seizure, and coma. The gradual acquisition of anti-disease immunity through repeated exposure establishes a strong negative correlation between severity of disease and host age. This means that young children are at the greatest risk of severe malaria episodes and death, while older children

and adults frequently exhibit mild or asymptomatic infections (Marsh and Kinyan-jui, 2006; Langhorne et al., 2008). The distribution of severe, mild and asymptomatic disease with respect to host age is shown in age distribution of morbidity for an endemic region in figure 1.4 (from (Marsh and Kinyanjui, 2006)). This pattern is especially apparent in high transmission regions where hosts are exposed to many parasites in their first few years of life, while in regions of lower transmission in-tensity older hosts still remain at risk of severe disease (Snow et al., 1997; Okiro et al., 2009; Reyburn et al., 2005). Infants up to six months old are protected from symptomatic disease but can still experience asymptomatic infection. This protec-tion is usually attributed to the passive transfer of maternal antibodies (Hviid and Staalsoe, 2004). While anti-disease immunity is generally thought to be mediated by exposure, there is some evidence from human infection experiments (malaria ther-apy data) that age-related changes in the way the immune responds to the parasite could also play a role in the rate of immune acquisition (Baird, 1995).
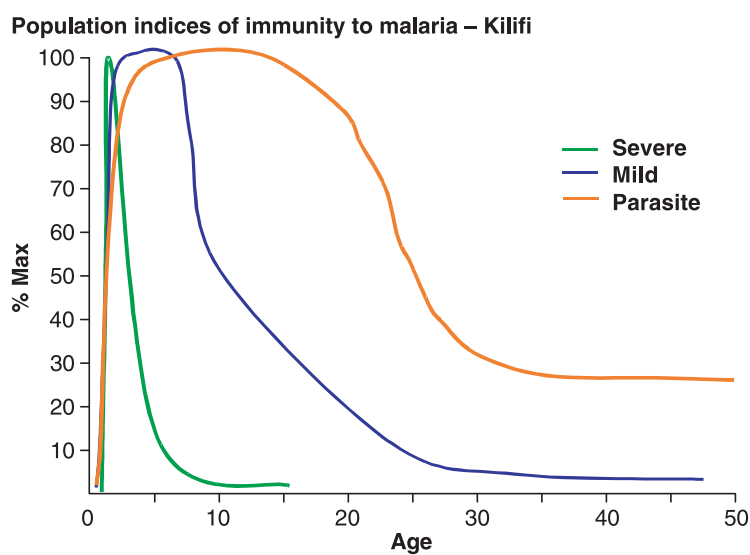
FIGURE 1.4: Anti-disease immunity is acquired gradually over repeated infections leading to a strong correlation between disease severity and host age in endemic regions. The figure emphasises the age at which maximum prevalence is observed for severe, mild and asymptomatic ('parasite') infections. Hence all prevalences are normalised by the maximum observed for each clinical outcome. The figure is taken from (Marsh and Kinyanjui, 2006) and was produced by Kevin Marsh and Sam Kinyanjui using data collected from multiple locations in Kilifi, Kenya.

1.2.2   PfEMP1

There has been considerable effort to explain the apparent lack of immunity and to understand the factors that lead to the age-dependent acquisition of anti-disease immunity. Recent research has emphasised the importance of several families of highly polymorphic proteins, which are expressed on the surface of infected erythrocytes. The best studied of these antigens is *Plasmodium falciparum* erythrocyte membrane protein 1 (PfEMP1), encoded by the *var* genes. PfEMP1 is broadly understood to be an important immune target of naturally acquired immunity, as well as a key virulence determinant. Its functional role involves facilitating cytoadhesion in order to evade the immune action of the spleen (Kraemer and Smith, 2006; Kyes, Kraemer, and Smith, 2007; Sherman, Eda, and Winograd, 2003). This can lead to sequestration of infected erythrocytes in the microvasculature of the host, leading to obstruction of blood flow and local inflammation.

Each *P. falciparum* genome encodes a repertoire of approximately 60 variants of PfEMP1, which are expressed in a synchronised and mutually exclusive manner during an infection. This process is called clonal antigenic variation. Exposure of the parasite's antigen repertoire to the immune system is thereby minimised, enabling the parasites to induce long infections (Scherf, Lopez-Rubio, and Riviere, 2008; Borst et al., 1995). The binding phenotype exhibited by infected erythrocytes is determined, at least in part, by the specific PfEMP1 variant expressed, which in turn determines the tissue/s in which sequestration occurs. PfEMP1 is considered an important virulence determinant because sequestration in different tissues is associated with different pathologies, complications and disease severity (Miller et al., 2002; Kraemer and Smith, 2006; Janes et al., 2011; Miller et al., 2002; Chen, Schlichtherle, and Wahlgren, 2000; Craig and Scherf, 2001; Sherman, Eda, and Winograd, 2003).

A consequence of antigenic variation is that a single parasite genotype is capable of inducing different pathologies depending on the particular phenotype/s of the PfEMP1 variants expressed, which in turn depends on the immune and/or physiological status of the host (Kraemer and Smith, 2003; Bull et al., 2000).

The task of identifying the specific properties and cytoadhesion phenotypes which are associated with severe disease has been the focus of much research over the last 20 years. The clearest example is pregnancy associated malaria (PAM), a severe form of malaria which can afflict pregnant women. PAM is associated with the expression of *var2csa*, an unusually well conserved *var* gene, whose transcription product binds placental chondroitin sulphate A (CSA) (Beeson et al., 2001; Salanti et al., 2004). Another example is that of rosetting, the binding of infected to uninfected erythrocytes, which is strongly correlated with increased risk of severe disease (Chen, Schlichtherle, and Wahlgren, 2000; Rowe et al., 1995). Other work suggests a link between severe and cerebral malaria and the ability of PfEMP1 variants to bind ICAM-1 (intercellular adhesion molecule-1) and EPCR (endothelial protein C receptor) (Tuikue Ndam et al., 2017). More generally, there is an association between risk of severe disease and the expression of a restricted sub-set of *var* genes that are also thought to be expressed preferentially in young children or immunologically naive individuals (Gupta et al., 1999; Bull et al., 1999; Bull et al., 2000; Kaestli et al., 2006; Nielsen et al., 2002).

Significant effort has been made by the malaria research community to understand the interaction between PfEMP1 expression, sequestration and susceptibility to disease. The release of a full parasite genome revealed that *var* genes can be broadly grouped based on their upstream promoter sequence (UpsA, UpsB, UpsC and UpsE) (Gardner et al., 2002; Lavstsen et al., 2003). Further analysis showed that there are differences in the chromosome location, transcription direction and

the number of domains between Ups groups (Lavstsen et al., 2003).

PfEMP1 exhibits a modular structure with between three and nine Duffy binding like (DBL) or cysteine-rich interdomain region (CIDR) domains (Smith et al., 2000; Rask et al., 2010; Kraemer and Smith, 2006; Kraemer et al., 2007; Kyes, Kraemer, and Smith, 2007). An example of domain level structure for a *var* gene is shown in figure 1.5. As the binding phenotype of a particular variant is ultimately determined by its domains, there has been some effort to understand whether different domain subtypes may predict the binding phenotype of PfEMP1. Particular domains, or often combinations of domains, have been linked to particular binding phenotypes and target different host receptors (Kraemer and Smith, 2006; Kyes, Kraemer, and Smith, 2007; Smith et al., 2000). For example, CIDRα has been found to bind CD36 (Baruch et al., 1997), whereas ICAM-1 binding has been associated with DBLβ-C2 (Mo et al., 2008). There is also evidence that the presence of DBLα1-like domains are associated with cerebral malaria (Kyriacou et al., 2006). The diversity in this domain level structure is what ultimately determines the range of phenotypes exhibited by different PfEMP1 variants (Claessens et al., 2012; Baruch et al., 1997; Mo et al., 2008)
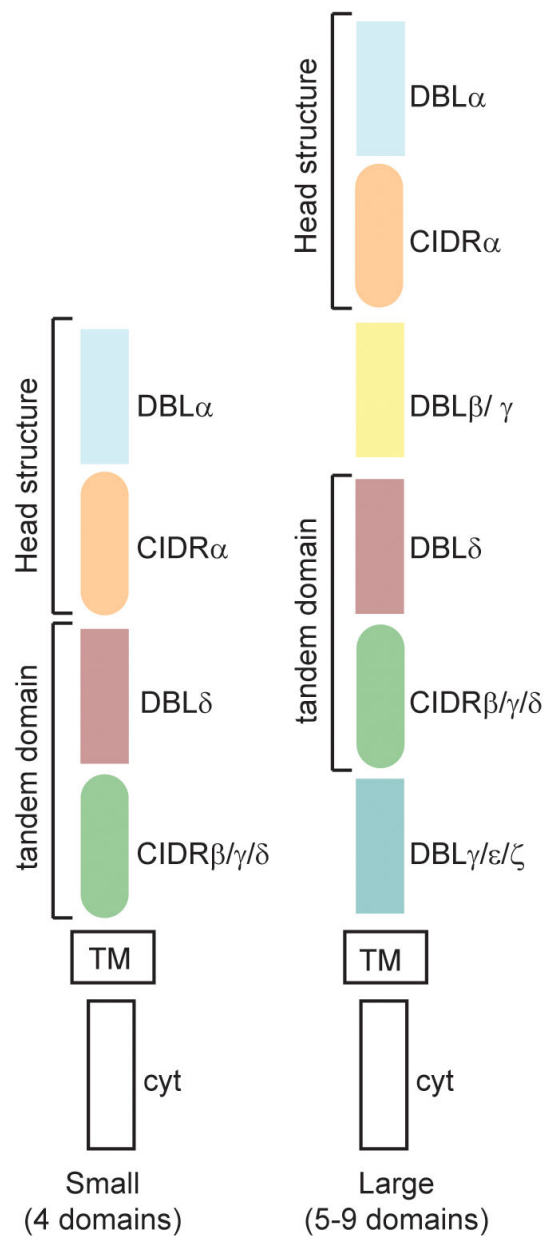
FIGURE 1.5: The domain level structure of the extra-cellular region of two example *var* genes, each consisting of a semi-conserved 'head' and several domains and tandem domain structures. Adapted from Smith *et al* (Smith, 2014)

1.2.2.1   Anti-PfEMP1 immunity

Naturally acquired immunity is primarily targeted at the blood-stage of the parasite's life cycle. Various antibody mediated mechanisms have been proposed, including the inhibition of erythrocyte invasion, opsonization and inhibition or reversal of sequestration (Langhorne et al., 2008; Mo et al., 2008; Udeinya et al., 1983). The relative importance of these mechanisms and the antigens involved in this process is an active area of research (Langhorne et al., 2008; Fowkes et al., 2010). Identifying specific antigen targets which are predictive of protection from disease has proved difficult, however, and it is becoming increasingly clear that anti-disease immunity is associated with the acquisition of a broad spectrum of antibodies to many different antigen variants.

The dual role that PfEMP1 plays in immune evasion and pathogenesis makes it a particularly promising candidate for understanding anti-disease immunity. An infection leads to the development of antibodies specific to the PfEMP1 variants to which the host has been exposed (Bull et al., 2000; Ofori et al., 2002), and susceptibility to clinical disease has been linked to hosts having a reduced ability to recognise specific variants expressed by the infecting parasite (Bull et al., 1998). This has given rise to the 'gap in the protective repertoire' hypothesis, in which hosts acquire a repertoire of protective antibodies in a piecemeal and predominantly variant-specific fashion (Bull et al., 1998; Bull et al., 1999; Bull et al., 2000). Susceptibility to disease is thus thought to be determined by the ability of a particular parasite strain to express antigens which correspond to gaps in the host's protective antibody repertoire. This hypothesis potentially explains two aspects of immunity: (1) that anti-disease immunity is acquired relatively quickly (since it is thought that severe disease is caused

by a restricted subset of PfEMP1 variants which may exhibit reduced sequence diversity), and (2) that hosts do not acquire sterile immunity in the face of extremely large PfEMP1 diversity.

There is evidence that anti-PfEMP1 immunity also involves some degree of cross-protection (Ofori et al., 2002; Giha et al., 1999b; Giha et al., 1999a). Additionally, it has been found that an ongoing sub-clinical infection is associated with a greater ability to recognise antigens from heterologous parasites (Bull et al., 2002; Smith et al., 1999a). This raises the possibility that the immune response generated by exposure to one strain is not entirely strain-specific. Cross-immunity may originate due to overlapping parasite repertoires in which whole antigens are shared between parasites (e.g. (Albrecht et al., 2006; Barry et al., 2007)), due to the presence of conserved or shared sequences between antigen variants.

Theoretical work has demonstrated that transient cross-immunity is capable of explaining the synchronised sequential expression of antigens and suggests an intriguing mechanism in which the parasite uses weak cross-immunity to drive antigen switching patterns to achieve the evasive behaviour necessary for chronic infections (Recker et al., 2004). However, the role of cross-immunity in naturally acquired immunity and between-host immune selection is not yet known.

1.2.2.2 Generation of PfEMP diversity

In biomolecular terms, diversity is generated primarily by two mechanisms: meiotic recombination between genetically distinct parasites inside the mosquito and mitotic recombination during the blood stage of infection (Rich, Hudson, and Ayala, 1997). The former allows whole genes to transfer between parasite lineages. Mitotic recombination, on the other-hand, involves the exchange of nucleotide sequences

between genes and is hypothesised to be a source of novel antigen variants (Ward et al., 1999).

The relative importance of mitotic and meiotic recombination and the evolutionary forces which maintain antigenic diversity are largely unknown. Conway *et al.* observed decreasing linkage disequilibrium from centromeres to telomeres, and argued that this suggests that diversity in the population is predominantly structured by mitotic recombination (Conway et al., 1999). Given that concurrent infections can be common (especially in high transmission regions) (Ntoumi et al., 1995; Smith et al., 1999b), mosquitoes which feed off hosts with multiple infections or which bite more than one infected host may act as hubs for recombination between parasite lineages. Mosquitoes may therefore play an important role in diversity generation (Koella, Sörensen, and Anderson, 1998; Klowden and Briegel, 1994). The effect of interactions between multiple parasite strains is not well understood, however, but may have a potentially important influence on infectivity and transmission (Morlais et al., 2015).

Other studies have highlighted the importance of mitotic recombination in generating new *var* gene diversity (Freitas-Junior et al., 2000; Claessens et al., 2014; Bopp et al., 2013). Mitotic recombination appears to occur in such a way as to maintain domain sequences and has been hypothesised to be the major source of diversity generation (Claessens et al., 2014). Furthermore, there is evidence that recombination occurs predominantly within Ups groups (Kraemer et al., 2007; Kraemer and Smith, 2003).

1.2.2.3 Population-level structure of diversity

The degree of diversity exhibited by PfEMP1 is difficult to characterise at the level of entire parasite populations because of the very large number of variants in circulation. All sequenced *var* genes to date have a DBLα domain at the 5′ end of the extracellular intron, and several studies have sequenced this domain in order to infer the population diversity (Barry et al., 2007; Day et al., 2017; Albrecht et al., 2010). These studies have reported immense diversity (Barry et al., 2007). Barry *et al.* (Barry et al., 2007), for example, was unable to place an upper estimate on global diversity because the rate of discovery of new variants did not decrease after sequencing over 1000 DBLα domains. At the local scale, however, over 400 samples were sequenced from a single village in Papua New Guinea, and approximately 200 distinct DBLα variants were found. This suggests that there is a spatial component to the population-level structure of diversity (Chen et al., 2011; Day et al., 2017; Nielsen et al., 2004; Aguiar et al., 1992). Caution must be used when interpreting these observations, as they were based on the sequence variation of a single domain and are difficult to map to immunologically distinct variants. The existence of conserved regions (e.g. as a consequence of the constraint to maintain functional binding to their target receptor) and/or the sharing of domains between variants (e.g. as a consequence of mitotic recombination) may mean that shared epitopes between variants are not uncommon (Bull and Abdi, 2016; Bull et al., 2008; Ward et al., 1999).

There is evidence that the population-level diversity of PfEMP1 is structured, with studies having shown little overlap between the *var* gene repertoires between any two parasites (Rask et al., 2010; Day et al., 2017; Chen et al., 2011; Barry et al., 2007; Aguiar et al., 1992). Rask *et al.* (Rask et al., 2010) analysed the *var* repertoires of

seven strains and found virtually no overlap. An exception to this was a small number of UpsE variants, which are known to be highly conserved and include *var2csa*, thought to be causative of pregnancy associated malaria. Similarly, Barry *et al.* (Barry et al., 2007) found virtually no overlap between DBLα sequences from repertoires globally, and only 7% overlap within samples taken from a single village. There is also good evidence that recovery from infection is associated with the acquisition of antibodies which recognise the PfEMP1 variants expressed by the infecting isolate (Marsh and Howard, 1986; Forsyth et al., 1989; Bull et al., 1999). Furthermore, sera from recovering children rarely recognised isolates from other children/infections. This strongly suggests that PfEMP1 diversity is a useful part of the parasite's immune evasion strategy, and also concurs with theoretical work which demonstrates that host immunity will select for non-overlapping antigen repertoires, resulting in structured diversity at the level of parasite populations (Gupta et al., 1996; Gupta and Anderson, 1999). This theoretical work is discussed in detail in section 1.3.3.1. Parasite populations can therefore be described as exhibiting a strain structure with respect to these antigen repertoires, in which strains minimise their immune interference with one another by minimising overlap in their repertoires.

Variation in the population level diversity of *var* genes across geographic regions may be linked to transmission intensity. For example studies based in low, moderate or seasonal transmission regions in South America have found much lower population level diversity in sequenced DBLα domains and greater overlap between antigen repertoires (Albrecht et al., 2006; Albrecht et al., 2010) than in regions with higher transmission (e.g. (Barry et al., 2007; Day et al., 2017)). Relatively high occurrence of antigen repertoire overlap was also reported for a low transmission region in Sudan (Giha et al., 1999b).

1.2.2.4 Gene and repertoire level structure

Recent work has revealed that *var* gene repertoires themselves are highly structured.
The proportion of UpsA, B and C genes which make up each repertoire is highly
conserved across geographically diverse parasite strains (Rask et al., 2010; Buckee
and Recker, 2012). The proportions of UpsA and UpsB/C genes in seven repertoires
are shown in Figure 1.6 (adapted from (Buckee and Recker, 2012)).
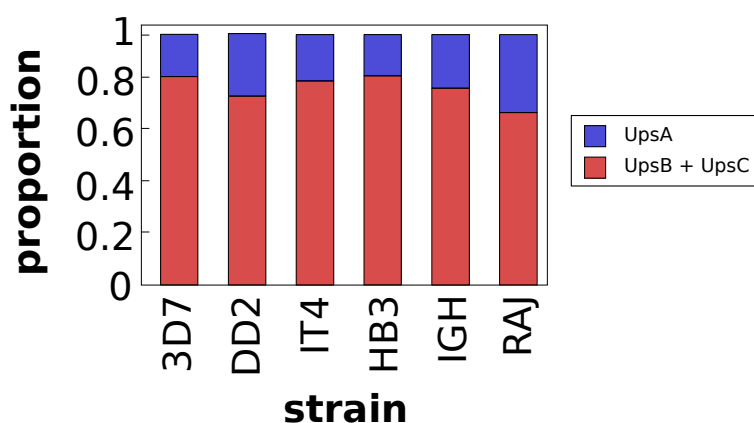
FIGURE 1.6: The observed repertoire partitioning structure in se-
quenced *var* gene repertoires. Adapted from (Buckee and Recker,
2012).

There is an association between the expression of UpsA *var* genes (and some-
times UpsB), severe disease and high parasite densities (Bull et al., 2005; Jensen
et al., 2004; Kaestli et al., 2006; Kyriacou et al., 2006), while UpsC *var* genes are
more likely to be expressed in semi-immune individuals during asymptomatic in-
fections. This suggests that the UpsA group includes variants which have evolved
to maximise their fitness in naive hosts, while other Ups groups may be specialised
to semi-immune hosts, in which immune evasion is weighted more strongly as a fac-
tor in parasite fitness (Kraemer and Smith, 2003; Kraemer and Smith, 2006). While
it is difficult to separate the association of UpsA expression with the confounding
factor that naive hosts may be more prone to high parasite density, evidence in sup-
port of this hypothesis has been steadily increasing. UpsA *var* genes appear to be

preferentially expressed in hosts with little exposure despite evidence for there being pre-disposition to express of centrally located (i.e. not UpsA) *var* genes in *in vitro* (Frank et al., 2007). This suggests that UpsA *var* genes might have a selective advantage within the host (Wang et al., 2009). Additionally, genetic studies have shown that UpsA genes exhibit less sequence diversity than other groups and also have a larger number of binding domains. Both of these properties might be expected in proteins that are optimised for binding affinity rather than immune evasion (Nielsen et al., 2002; Rask et al., 2010; Kraemer and Smith, 2003; Kraemer and Smith, 2006).

1.2.3   *P. falciparum* transmission and immunity

In *P. falciparum*, anti-disease immunity does not necessarily correspond directly to reduction in transmission, and while the acquisition of anti-disease immunity is relatively well characterised, there is considerable uncertainty regarding the impact of immunity on transmission, which could manifest through changes in the infectivity to mosquitoes or the duration of infection.

The first opportunity for immune interference with transmission comes as sporozoites enter the bloodstream from an infected mosquito bite and migrate to the liver (the pre-erythrocytic stage). Effective anti-sporozoite immunity has been induced under lab conditions using irradiated sporozoites (Hoffman et al., 2002). However, natural infection leads to incomplete pre-erythrocytic immunity and is generally considered to play a minimal role in naturally acquired immunity due to the relatively short window of opportunity in which free-moving sporozoites come into contact with the host's immune system.

Ultimately, the infectivity of a human to mosquitoes is dependent on the presence of gametocytes in the blood and the ability of ingested gametocytes to establish an infection in the mosquito. Gametocytes are typically present at low densities

compared to asexual parasite stages, which may help to minimise contact with the host immune system and thereby help to evade targeted anti-gametocyte immunity (Taylor and Read, 1997). While anti-gametocyte antibodies are produced by at least some hosts, it is not clear how effectively these inhibit gametocyte function (Smalley and Sinden, 1977). Exposure-dependent changes in immunity mean that both age and cumulative exposure are correlated with decreasing parasite density (Beadle et al., 1995). This suggests that naturally acquired immunity facilitates faster clearance of blood stage parasites and/or inhibits their replication. How this affects host infectivity to mosquitoes is less clear. In gametocytemic individuals, the ratio of gametocytes to asexual parasites increases in older individuals, despite an overall decrease in total parasite density (Drakeley et al., 2006a; Molineaux, Gramiccia, Organization, et al., 1980). Whether this is due to the parasite adapting its transmission strategy by changing the ratio of gametocyte differentiation in response to the immune status of the host or simply arises due to the selective destruction of merozoites is not known (Bousema and Drakeley, 2011).

Immune action resulting in a change in the duration of infection would also affect transmission. However, estimates of infection length vary widely from 1-2 weeks (Bruce et al., 2000) to several years (Sama, Killeen, and Smith, 2004; Smith and Vounatsou, 2003), reflecting the difficulty of accurately measuring duration of infection in the field. Most estimates range from several weeks to up to around 250 days (Felger et al., 2012; Bruce et al., 2000; Smith et al., 1999a; Bousema et al., 2011; Kitua et al., 1996; Bekessy, Molineaux, and Storey, 1976). The duration of infection can be difficult to estimate because of overlapping infections, recrudescence, asymptomatic infection and/or co-infection with other malaria species. It can be especially difficult to estimate in young children because they frequently exhibit high parasitaemia or severe disease, and as such the natural course of the infection is often

interrupted by treatment. Limitations in the sensitivity of detection methods can mean that parasites frequently appear to be cleared, only for the same genotype to re-emerge days or weeks later (Bekessy, Molineaux, and Storey, 1976; Bruce et al., 2000). This is further obscured by the parasite's ability to sequester in host microvasculature and synchronise egress, meaning that the density of blood-stage parasites fluctuates between extremes and may periodically drop below detection thresholds (Bruce et al., 2000; Felger et al., 2012). One recent cohort study estimated that as few as 30% of ongoing infections are detected using nested PCR and showed that detectability is generally lower in older age groups, which typically exhibit lower parasite densities (Felger et al., 2012).

In high transmission regions hosts frequently exhibit overlapping infections involving multiple parasite clones, making it difficult to determine when one infection ends and another begins. Isolates can be genotyped to determine the presence or absence of particular clones, but this is time-consuming and expensive. Furthermore, it is difficult to account for any effects that co-circulating strains may have on one another, for instance because of immunological interference or changes in parasite behaviour.

Several studies report that the duration of infection peaks in young children and is generally negatively correlated with age (Bruce et al., 2000; Felger et al., 2012). Bruce *et al.* (Bruce et al., 2000) used microscopy based detection methods and controlled for the low parasite densities between waves of parasitaemia. They estimated infection lengths of greater than 48 days in children under 5 years of age (the duration of the study period prevented tracking of infections for longer), which drops to around two weeks in individuals aged over 14. In contrast, Felger *et al.* (Felger et al., 2012) fit three statistical models to data, gathered using a combination of microscopy and nested PCR, from a high transmission region to predict the duration of infection

for different age groups. They found that infection length peaks in children aged between 5 and 9, and last for approximately 200 to 300 days, much longer than Bruce *et al.*. This fell to 120 days in adults aged 20+ years. Smith *et al.* estimated a positive correlation between exposure and duration of infection (Smith and Vounatsou, 2003); however, this study only considered children younger than 30 months and is therefore not directly in conflict with estimates by Felger *et al.*.

One complicating aspect is that older individuals with well developed immunity sometimes exhibit chronic infections, which may last many months or even for years. While the factors which determine when and how chronic infections are sustained are not clear, antigenic variation and the interaction of diversity and the host immune system is undoubtedly important (Druilhe and Perignon, 1997; Smith et al., 1999a; Recker et al., 2004).

In the context of onward transmission, the important factor is the presence of viable gametocytes in the blood of the host. Anti-gametocyte immunity, gametocyte longevity and density, and gametocyte detectability are therefore all potentially important considerations. Gametocyte clearance can lag behind that of merozoites, meaning hosts may remain infectious to mosquitoes after asexual parasites have been fully cleared. To complicate the situation further, limitations in gametocyte screening methodology mean even individuals with no detectable gametocytemia can be infectious to mosquitoes (e.g. (Boudin et al., 1993)). The malaria therapy dataset provides unique insight into the dynamics of untreated severe infections, and Eichner *et al.* (Eichner et al., 2001) used this to estimate gametocyte survival times. They found that there was a large variation, between 1.3 and 22.2 days, and gametocytemia status was sometimes periodic. The upper portion of this range suggests hosts may remain infectious to mosquitoes for up to three weeks following clearance of asexual blood stage parasites.

The uncertainty in estimates of duration of infection and host infectivity is problematic for understanding the relationship between transmission and prevalence. A number of studies have tried to estimate the relative contribution to transmission by different age groups, but the evidence is conflicting. Some reports suggest that the age distribution of infectivity to mosquitoes is strongly skewed towards young children and hence predict that transmission is disproportionately driven by younger age groups (Githeko et al., 1992; Graves et al., 1988; Coleman et al., 2004), while others report little difference between age groups (Drakeley et al., 2000; Boudin et al., 2004; Ross, Killeen, and Smith, 2006). To further confuse matters, densities of gametocytes which are too low to reliably detect with microscopy can still infect mosquitoes (Schneider et al., 2007). Mosquitoes fed on infected blood in which only asexual parasites could be detected have ended up being equally as infective as those where gametocytes were detected (Boudin et al., 1993; Coleman et al., 2004). This makes it difficult to characterise the relationship between gametocyte density and infectivity using microscopy methods (or even PCR (Nwakanma et al., 2008)). Ross *et al.* developed a statistical model to predict host infectivity to mosquitoes from asexual parasite density (Ross, Killeen, and Smith, 2006) and concluded that adults contribute similarly to transmission as younger age groups, and argued that while young children may have higher infectivity to mosquitoes, uneven biting frequencies between age groups lead to relatively even contributions across age groups. However, this model is based on malaria-therapy data in which all hosts were naive, and as such is cannot account for the effect of cumulative exposure and the possibility of anti-gametocyte immunity.

1.2.4 Summary: *var* gene diversity explains *P. falciparum* epidemiology

The *var* genes, and the multi-scale structured diversity they exhibit, are central to explaining several important aspects of *P. falciparum* epidemiology. The extensive population-level diversity and population level structure allows parasites to reinfect hosts throughout their lives, while *in vivo* antigenic variation helps to establish long infections and ensure onward transmission. As a result anti-disease immunity is acquired gradually through repeated exposure to a diverse set of antigen variants. Furthermore, diversity in cytoadhesion phenotypes of different PfEMP1 variants help to explain the range of pathologies observed in *P. falciparum* infections. Severe disease is thought to be associated with the expression of a more restricted sub-set of antigen variants, which may help to explain why hosts acquire immunity against severe disease relatively quickly, but remain susceptible to infection throughout their lives.

Recent theoretical work has begun to unpick the evolutionary processes which may explain differences in structure, phenotype and selective expression of PfEMP1 (e.g. (Buckee and Recker, 2012; Recker et al., 2004; Recker et al., 2011)), while others have concentrated on the processes that lead to repertoire and population level structure (Recker and Gupta, 2005; Recker, Arinaminpathy, and Buckee, 2008; Artzy-Randrup et al., 2012; Bull et al., 2008; Buckee, Bull, and Gupta, 2009). It is clear, however, that more work is needed to fully understand not only the consequences but also the generation and maintenance of this diversity on *P. falciparum* epidemiology.

1.3   Models of malaria: Parasite diversity and immunity

Modelling has become an invaluable part of modern epidemiology by providing a way to objectively assess the consequences and internal logic of competing hypotheses and make predictions about reality. Gaps in our knowledge are highlighted when predictions disagree with observations, causing us to reassess our assumptions. This guides the focus of lab and field work and generates new hypotheses. Useful models are simplified versions of reality which capture the essential processes of the target system, while being amenable to analysis and interpretation by omitting non-essential complexity. With this in mind, the following sections provide a short overview of some important models of malaria epidemiology.

By the early 20th century *Anopheles* mosquitoes were well established as the vector for malaria between human hosts (Ross, 1911b). Work began on trying to characterise the relationship between vector populations, transmission and malaria prevalence with the aim of developing effective control strategies. This led Ronald Ross to develop the first malaria models (Ross, 1905; Ross, 1911a; Ross, 1911b). These early models captured some of the key aspects of malarial transmission in a quantitative framework for the first time and have remained influential in the field of malaria modelling to this day.

The Ross-MacDonald model is probably best described as a theory of mosquito-borne transmission rather than a single model, as discussed in depth by Smith *et al.* (Smith et al., 2012). It encompasses a set of processes and assumptions pioneered by Ronald Ross and built upon by George MacDonald (Ross, 1905; Ross, 1908; Ross, 1911a; Macdonald, 1956; Macdonald et al., 1957) and numerous others over the best part of a century. The Ross-MacDonald theory of mosquito-borne transmission describes the major ecological interactions, which facilitate transmission of the malaria

parasite, namely transmission from an infectious human to mosquito, growth and maturation of the parasite in the mosquito, transmission to a human host by a bite from an infectious mosquito and multiplication in the human host. These embody a baseline set of processes on which a great many other Ross-MacDonald style models have been based.

Ross-MacDonald style models share a similar set of assumptions which govern each step in this transmission cycle. For example, Ross' second model, published in 1911 (Ross, 1911a), describes a population of hosts and a population of mosquitoes, each of which are split into two epidemiological compartments: infectious and susceptible, as illustrated in figure 1.7. The model gives mathematical descriptions of the flux between these compartments, i.e. susceptible hosts can become infectious, infectious hosts can recover (returning to the susceptible compartment), and susceptible mosquitoes can become infectious. In more general terms, the host population can be describes as an SIS (susceptible-infectious-susceptible) model, while the mosquito population is modelled as SI (susceptible-infectious), as depicted in Figure 1.7. The populations are kept constant with deaths and births being equal, and new births entering the susceptible compartment of their respective population. The change in infectious humans (Z) and mosquitoes (Z') is given by

$$\frac{dZ}{dt} = k'Z'(p - Z) - qZ, \tag{1.1}$$

$$\frac{dZ'}{dt} = kZ(p' - Z') - q'Z'. \tag{1.2}$$

The susceptible population can be calculated by subtracting the infectious population from the total population, i.e. $Y = p - Z$ and $Y' = p' - Z'$ for humans and mosquitoes, respectively. Changes to the populations are governed by two

parameters: the 'happenings' rates which govern new infections (k' and k for human and mosquito infections respectively), and the removal rates (g and g' respectively). The removal rates correspond to the recovery rate in humans and death rate in mosquitoes (which are assumed never to recover). The total human and mosquito populations are given by p and p', respectively. Ross uses the term 'happenings' to refer to the rate of events occurring in a system. In this case events are new infections, so k and k' are rates encompassing factors which contribute to infection (such as biting rate and infectivity) for mosquitoes and hosts. k' and k are composite parameters incorporating the various factors which influence the rate of new infections and can be formulated as

$$k' = mabZ', \tag{1.3}$$

and

$$k = acZ, \tag{1.4}$$

where m is the ratio of mosquitoes to humans, a is the proportion of mosquitoes that feed on humans per day, b is the probability of an infectious bite resulting in an infection, and c is the probability a mosquito becomes infectious after biting an infectious human.

While simple, this model captures many of the key aspects of malaria transmission and provides a tool from which the dynamics of malaria transmission could be quantitatively understood for the first time. The transmission dynamics depend on the interaction between human and mosquito populations (and their infectious sub-populations), allowing hypotheses about the multi-host transmission cycle to be tested. Hosts can be reinfected throughout their lives, and the flux of new infectious individuals is determined by the feedback between the infectious reservoirs of the
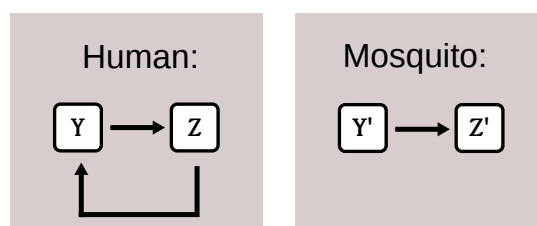
FIGURE 1.7: The epidemiological compartments of Ross' second model (Ross, 1911a). The human population is divided into infected (Y) and infectious (Z) subpopulations. The mosquito population is similarly compartmentalised (Y' and Z' respectively). Humans can recover from infection and return to the Y compartment, but once mosquitos become infected they remain infected for the duration of their lives.

two populations.

With the benefit of retrospective insight from over a century of malarial research Ross' models can appear to describe a rather oversimplified view of malaria epidemiology. The ratio of humans to mosquitoes is constant and humans receive bites at a constant rate without regard for seasonality or other factors. The probability that a bite from an infected mosquito leads to infection in a human does not vary with exposure and there is no attempt to account for antigenic diversity or the acquisition of immunity. Therefore, variation in the happenings rate in humans over time is entirely dependent on the proportion of infectious mosquitoes.

Equally, the progression of infection from inoculation to recovery (or death) over time is reduced to a simple binary 'infectious' versus 'not infectious' state; i.e. the happenings and loss rates encapsulates the average infectivity (from humans to mosquito and vice versa), mean duration of human infection, number of infected mosquitoes that die before becoming infectious, etc. without accounting for individual variation or stochasticity. From our perspective, the largest drawback is that there is no model of immunity or exposure-dependent variation in transmission. This makes it impossible to use this model to gain insight into these crucial processes, nor can the model be used to explain patterns of virulence or age-dependent

burden of disease.

Despite the model's simplicity Ross was able to identify some of the most fundamental characteristics of vector borne diseases. He predicted the existence of a critical threshold in the ratio of vectors to hosts, below which sustained transmission is impossible, and concluded that is therefore not necessary to kill every single vector in order to eliminate malaria (Ross, 1911a). Furthermore, he showed that there was a non-linear relationship between transmission and prevalence, noting that small increases in the vector population above this critical threshold causes a comparatively large increase in prevalence, which would eventually saturate such that any further increase in vector population has little effect on parasite prevalence (Ross, 1911a). These simple relationships and predictions are easy to take for granted given the wealth of cumulative knowledge and quantitative tools available to modern epidemiologists. However, they provided, for the first time, a quantitative tool which could demonstrate the relationship between many of the key aspects of malaria transmission, and which could be used to assess and design effective control programmes.

Half a century later George MacDonald published several extensions to Ross' framework, considering the effect of the extrinsic incubation period (EIP), mosquito longevity (Macdonald et al., 1957; Macdonald, 1956), and superinfection on the rate of transmission and equilibrium prevalence (discussed in detail in (Smith et al., 2012) and (Bellan, 2010)). This research highlighted the importance of demographic factors in mosquito populations by showing that because the EIP is on a similar order of magnitude as the vector species' life expectancy, many infected mosquitoes will die before they become infectious to humans. MacDonald also showed that Ross' models (which did not consider EIP) overestimated transmission and hence prevalence. MacDonald additionally applied the concept of the basic reproduction number (a

concept previously used in population ecology) to malaria transmission. Using this he could quantitatively demonstrate the effect of control measures targeting different parts of the transmission cycle, and use this to justify the application of strategies aimed at reducing mosquito longevity (Macdonald, 1956).

While neither Ross nor MacDonald considered immunity or parasite diversity in any detail, their framework has been extended by numerous researchers over the decades and produced a myriad of different models. Notable, only a small proportion of these models explicitly consider diversity. The following sections give an overview of some of the ways in which immunity and diversity have been treated. Rather than giving a comprehensive listing of malaria models, we illustrate various techniques and assumptions used to model parasite diversity and the acquisition of immunity using selected models as examples.

1.3.1 Exposure-mediated immunity

The Garki project was an initiative by the World Health Organisation in the 1970s, which aimed to understand the effect of residual insecticide spraying and mass drug administration on parasite prevalence in the African savannah. As part of the project, Dietz et al. (Dietz, Molineaux, and Thomas, 1974) developed a model to predict the effect of control measures which they fit using experimental data from field sites in Sudan and Nigeria. Building on the Ross-MacDonald framework, the model incorporated immunity as a binary 'non-immune' or 'semi-immune' state in hosts. This resulted in a model in which the host population is compartmentalised into seven epidemiological categories. These outline the main epidemiological states that human hosts can take, illustrated in figure 1.8, and include, uninfected (both non-immune and semi-immune), incubating (both non-immune and semi-immune),

recovering (non-immune and semi-immune) and infectious compartments. Infection of hosts is governed by a happenings rate h. Non-immune hosts become infectious after a period of incubation, while semi-immune hosts are assumed to be non-infectious but can nevertheless test positive for infection. Non-immune individuals remain infectious until they move to the non-infectious 'slowly recovering' state with fixed rate α_1. Hosts may then either recover and remain non-immune (at constant rate, R_1) or become immune at a rate of α_2 and they recover quickly (R_2).
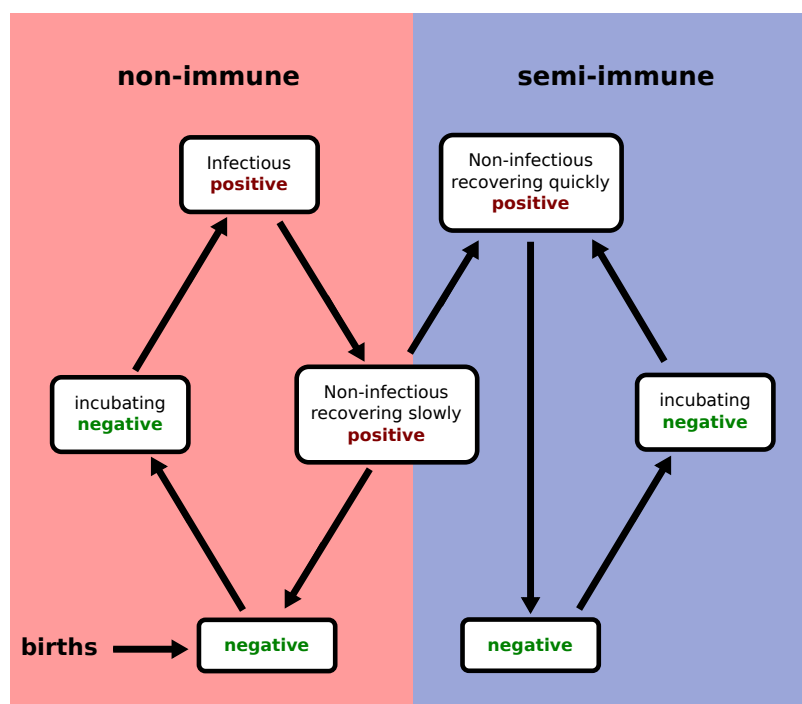
FIGURE 1.8: In the Garki model (Dietz, Molineaux, and Thomas, 1974) the human population is compartmentalised into seven sub-populations which capture the important changes in the epidemiological status of hosts. Immunity is handled as a binary state: non-immune or semi-immune. This determines host infectiousness, recovery time and the detectability of infection.

The main inputs required by the model are the birth/death rate for the region (population is assumed to be constant) and the vectorial capacity. The vectorial capacity is an expression of the vector-base determinants of transmission, defined by Dietz *et al.* as

$$C = \frac{ma^2 p^n}{-\ln p},\tag{1.5}$$

where m is the number of mosquito vectors per human, a is the average number of bites to humans by an individual mosquito per day, p is the probability of a vector surviving one day and n is the incubation period in the vector species. Each of these parameters must be measured or estimated for the region of interest in order to the calculate vectorial capacity. Seasonality can be incorporated by interpreting these parameters as functions of time. Other parameters of the model, such as recovery rates for immune and non-immune individuals, inoculation rates and the rate of acquisition of immunity were obtained by fitting the model to field data from Garki in Nigeria. The model generates a time series of parasite prevalence, which can be broken down to distinguish between the proportion of detected infections and the true prevalence, the proportion of hosts with immunity, and the age distribution of prevalence and immunity. The model was designed to be easily parameterised using existing field and laboratory methods for calculating vectorial capacity, and to be validated against observational prevalence data.

While the diversity in the parasite population is not explicitly modelled, differentiating between semi-immune and non-immune hosts allows the model to capture some of the mechanisms which arise from the continuous exposure to a diverse parasite population. Since infected hosts transition from non-immune to semi-immune with a fixed probability, there can be exposure-mediated variation in the recovery rate, susceptibility and the detectability of infection. Furthermore, the rate of acquisition of the semi-immune state is dependent on transmission intensity because only infected hosts may make the transition. Explicitly modelling the effect of exposure on these parameters allowed the authors to explain patterns in the age distribution

of prevalence and understand how control measures may impact on different age groups.

While the model was able to reproduce prevalence from the testing field sites to which it was fit, it is unlikely to generalise to other geographic regions in which the degree of diversity and parasite population structure may be different. The authors acknowledge that the parameters would need to be adjusted to estimate transmission in other regions; however, the simple handling of immunity as a binary state may further limit the ability of the model to fit different scenarios. In particular, since the age distribution of immunity (or at least anti-disease immunity) appears to be strongly dependent on transmission intensity, a simple two-state immunity may not offer sufficient resolution for settings with other transmission intensities. This may be especially true in lower transmission regions, in which hosts may be expected to have intermediate states of immunity for longer periods of their lives.

1.3.2   Immunity and the burden of disease

In 1956 MacDonald warned against the potential danger of partially effective control measures. He argued that while moderate control might reduce transmission, any reduction in prevalence could be balanced out in the long run by an increase in the infectious period due to lower population level immunity (Macdonald, 1956). In this view, any reduction in prevalence may only be temporary, with the true effect of control only becoming apparent when immunity equilibrates. While we now know that the epidemiological patterns of acquired immunity are considerably more complex than that assumed by MacDonald (Boudin et al., 1993; Eichner et al., 2001; Drakeley et al., 2006a; Drakeley et al., 2006b; Molineaux, Gramiccia, Organization, et al., 1980; Bousema et al., 2010; Bousema and Drakeley, 2011; Smalley and Sinden, 1977), there is clear evidence that protection from severe and life-threatening disease is

exposure-dependent. Ghani *et al.* (Ghani et al., 2009) used a mathematical model to explore the effect that periods of control have on the burden of disease. Specifically, the authors investigated the effect of endemic stability in *P. falciparum* malaria and a related phenomena in which decreases in transmission can cause a net increase in the burden of disease following natural or control induced changes in transmission. Endemic stability refers to a state where the incidence of clinical disease can be low despite high prevalence, which can occur when virulence is linked to the frequency of exposure (Coleman, Perry, and Woolhouse, 2001).

The authors adapted the Ross-MacDonald framework by splitting the infected human population into four categories based on infection status and disease severity: exposed, symptomatic infection, asymptomatic infection, and sub-patent infection. The rate of recovery and probability of the different levels of disease severity in hosts are modulated by the level of immunity. Immunity is modelled as two separate components: clinical immunity and 'parasite immunity'. Clinical immunity reduces the probability of an infection to cause symptomatic disease. Parasite immunity increases the recovery rate and applies to both symptomatic and asymptomatic infection. Explicitly separating the effect of immunity on morbidity from its effect on transmission allowed the model consider both aspects of immunity with greater flexibility.

The model was fit to data from Senegal and Dakar (Trape and Rogier, 1996) and was then used to understand the net effect of different control strategies on the average number of clinical episodes. Ghani *et al.* considered control strategies which reduce exposure (e.g. betnets), reduce susceptibility (e.g. vaccination) or increased recovery (e.g. prophylaxis). Endemic stability was observed under a wide range of parameters but could be disrupted intervention. The breakdown of endemic stability was characterised by an initial sharp drop in clinical episodes as the intervention

reduced transmission. This was followed by a slow rise in clinical episodes as immunity was slowly lost from the population until a new equilibrium was established. In some scenarios the disease burden could actually increase to higher equilibrium levels, despite sustained control. Although caution should be used when drawing quantitative conclusions from this work, given the considerable uncertainty surrounding the rates of acquisition and loss of immunity, it demonstrated that endemic stability in *P. falciparum* could mean that poorly designed control policies can yield lower than expected reductions in morbidity, or even lead to an increase in the disease burden over the long term.

Many models of *P. falciparum* transmission treat immunity in an abstract way, either by defining discrete immuno-epidemiological states that hosts move between with fixed probabilities or by defining immunity along a continuum as a function of cumulative exposure (e.g. (Dietz, Molineaux, and Thomas, 1974; Struchiner, Halloran, and Spielman, 1989). Relatively few models explicitly consider immunity at the level of antigens. The Ghani model (Ghani et al., 2009) (discussed above) takes an important step in recognising that anti-disease and transmission-reducing immunity may not be acquired together, nor through the same mechanisms. However, it does not explicitly consider antigenic diversity or parasite population structure. Like the majority of malaria models, antigenic diversity and parasite population structure are only considered as an aggregated property, which is encoded in the various rates or probabilities that guide host state transitions. This assumes that antigenic diversity is homogeneous across the parasite population and makes it difficult to investigate or test hypotheses about the effect of structure in parasite populations or individual genomes. In the next sections, we review models which explicitly consider antigenic diversity alongside immunity.

1.3.3 Modelling antigenic diversity

This section discusses previous theoretical work on the evolution of, and epidemiological consequences for, antigenic diversity as viewed at different ecological scales. Subsection 1.3.3.1 concerns antigenic diversity at the population level, while subsections 1.3.3.2 and 1.3.3.3 discuss antigenic diversity within an infection and within a single parasite genome respectively.

1.3.3.1 Structured diversity at the population level

Gupta *et al.* began to explore the dynamics which arise from strain structured populations of *P. falciparum*. Arguing that *P. falciparum* epidemiology can be best understood by considering the parasite population as a collection of largely independently transmitted strains (Gupta and Day, 1994), Gupta explored the transmission dynamics of multiple *P. falciparum* strains by extending a simple Ross-MacDonald framework to consider two strains that enter a susceptible population at the same time (Gupta, Swinton, and Anderson, 1994). The model therefore comprised four host compartments: susceptible, immune to strain 1, immune to strain 2, and immune to both. For convenience, hosts become immune upon exposure (preventing reinfection by the same strain) but in doing so transition through a period of infectiousness. Immunity is acquired upon infection, which prevents reinfection by the same strain, and also confers a degree of cross-protection to the other. This is modulated by a cross-immunity coefficient, which scaled the transmission rate for partially immune hosts. Immunity can only be lost in the population through host death. Mosquitoes are assumed to only become infectious with one or other strain, and, once infectious, remain so until they die. This framework allowed Gupta to characterise the effect of immune interference and/or differing transmission rates between strains on the

transmission dynamics and equilibrium prevalence in the host population.

The authors show that competition for susceptible hosts, mediated via cross-immunity, can produce a wide range of dynamic behaviours. In particular, strong cross-immunity can act to synchronise strains and lead to cyclic patterns independent of changes to vector populations, host migration or seasonal effects. Very strong cross-immunity causes the exclusion of the least transmissible strain, whilst under low cross-immunity strains do not synchronise and approximate the dynamics of independently transmitted agents.

Two years later Gupta *et al.* published another model, which addressed the question of how immune selection can lead to the emergence and maintenance of strain structure (Gupta et al., 1996). While this model was not specific to malaria, the authors demonstrate a fundamental mechanism, which applies generally to the evolution of a wide range of recombining infectious agents, including *P. falciparum*. Parasites in the model express n antigen-encoding loci, that can encode any of k alleles (without allele sharing between loci), leading to n^k possible strains. The central assumption in the model is a form of cross-immunity, in which exposure to one strain gives hosts some level of protection to any other strains which share one or more antigens. In the context of *P. falciparum*, this corresponds well to the 'gap in the protective repertoire' hypothesis.

Gupta *et al.* showed that immune selection with moderate levels of cross-immunity could lead to the system being dominated by a discordant set of strains. Discordant strains are strains which do not share any alleles and which therefore do not induce partial immunity to one another. This dominant set inhibits the invasion of other strains by effectively occupying the available antigenic space. The dominant set was therefore stable even when recombination was allowed to frequently generate non-dominant strains. This prediction of a stable structure in populations of *P. falciparum*

is partially supported by evidence from genetic studies, which show little overlap between antigen repertoires of different parasites, e.g. (Day et al., 2017; Chen et al., 2011; Rask et al., 2010).

While the model was not directly applied to *P. falciparum*, it demonstrates important dynamical behaviour which arises from antigen-level interactions between strains. Importantly, these cannot be understood without explicitly considering the structured diversity of a parasite population. Furthermore, Gupta *et al.*'s analysis suggests that control measures which induce partial immunity or reduce the period of infectiousness have the potential to alter strain structure (Gupta, Ferguson, and Anderson, 1998; Gupta and Anderson, 1999), and may thus have an impact on the long term effectiveness of interventions.

1.3.3.2   Antigenic diversity within an infection

Antigenic variation allows the flexible expression of different antigen variants by clonal *P. falciparum* parasites over the course of an infection. Recker *et al.* (Recker et al., 2004) demonstrated that transient cross-immunity offers a possible mechanism by which the expression of antigen variants could be synchronised within an ongoing infection. The central assumption was that antigens contain two types of epitope: major and minor. These differed in that major epitopes induced long-term immunity while minor epitopes induced short-term immunity. Antigens contained one unique major epitope in addition to a number of minor epitopes selected from a pool of antigen variants. Over the course of an infection transient immunity develops and wanes to the various minor epitopes depending on the antigen variant expressed at the time. This results in an ever-changing within-host fitness landscape for antigens, which in turn selects for/against the expression of particular antigen variants in a synchronised way. This occurred despite an equal probability of switching between

the expression of antigen variants, simply because differences in the immunity de-rived within-host fitness of variants meant that a single variant would out-compete others variants.

One of the properties of this immune driven mechanism for synchronised expres-sion is that when the strength (or duration) of immunity to minor epitopes increases there is an increasing trend toward long infections and low parasite densities. The authors point out that this may help to explain the propensity of older individuals to experience chronic infections with low parasitemia, e.g. under the assumption that cumulative exposure or age related differences in immune function lead to more effi-cient immunity against minor epitopes. Transient cross-immunity may therefore be able to explain not just within-host synchronisation of antigen expression, but also increasing chronicity of infection in older individuals with greater exposure.

Another aspect of *P. falciparum* epidemiology which arises naturally from this model is that of premunition. Premunition is the phenomena that the existence of low density parasitemia below the threshold for symptomatic disease can confer protection from disease (Bull et al., 2002; Smith et al., 1999a). In the context of this model, low density infections boost immunity to minor epitopes and offers some protection against acute infection by other strains.

1.3.3.3   Structured diversity within the genome

According to theory, population level immune selection selects for non-overlapping repertoires comprised of antigens with a small number of domains that minimise the chance of immune interference to one another. This is at odds with genetic studies which show that, in terms of the number of domains, antigen repertoires are struc-tured in such a way as to contain a mix of long and short PfEMP1 variants (Rask et al., 2010; Smith et al., 2000). This structure is well conserved over geographically

diverse parasite strains, suggesting that there is a functional advantage to the parasite in adhering to such a structure. Recent work by Buckee and Recker (Buckee and Recker, 2012) used an individual-based model to better understand the conditions under which these differences in the domain level structure of PfEMP1 variants are expected to evolve.

The model considered immunity to act in one of two ways. In the first, immunity induces sterile cross-immunity to whole antigens if they contain any domain to which a host had previously been exposed to. In the second, immunity acts at the domain level to inhibit cytoadhesion and thus to reduce the fitness of the antigen by reducing the duration of infection it can induce in the host. The authors tested the fitness of one and two domain antigens using each type of immunity. Under the first 'whole antigen' form of immunity, single domain antigens were selected for. This is because within-host fitness and population level ('between-host') immune selection will favour antigens with small numbers of domains as this minimises the chances of cross-interference. When immunity acts at the domain level by inhibiting specific domains however, multi-domain antigens can retain some binding efficacy in partially immune hosts. Multi-domain antigens were therefore preferentially selected for because they spread the risk of encountering hosts with prior immunity to all of their domains. Finally, when the fitness of these antigens were tested using both types of immunity acting together, a mixture of short and long antigens were selected for. This suggests that by utilising repertoires with mixed length antigens, parasites could be balancing a trade-off between immune selection acting at the whole gene level and immune selection acting at the domain level.

Next, the authors showed that there was a selective advantage for antigens to specialise in different functions: binding affinity or diversity. In other words, antigens with two high affinity domains (maximising within-host fitness) or two high

diversity domains (minimising between-host immune selection) domains were selected for, while antigens with mixed domain phenotypes were at a disadvantage. This supports the hypothesis that differences in the structure and diversity of different *var* genes may have evolved as a consequence of their specialisation to different host niches (see (Kraemer and Smith, 2003)), with long and conserved UpsA genes specialising in naive hosts, while shorter, more diverse *var* genes may be adapted to semi-immune host niches.

Finally, Buckee and Recker looked at the evolution of antigen repertoires containing two antigens using the same model. To convey the sort of flexible phenotype expression produced by antigenic variation, parasite fitness was determined by the highest fitness antigen. The evolutionary trade-off between within-host and between-host selection could select for antigen repertoires which were structured to contain both specialised forms of antigens. Parasites containing mixed repertoires could change strategy to maximise their fitness depending the immune status of the host.

These models demonstrate how interactions between the human immune system and antigenic diversity of the parasite population are fundamentally important to explain and explore the epidemiology and evolution of *P. falciparum*. From the role of immune selection in shaping whole parasite populations, to the interplay of within- and between- host fitness on the maintenance of domain structure within single genes, an understanding of these processes at multiple scales is a necessary prerequisite to developing a more complete theory of the evolutionary epidemiology of *P. falciparum*.

Summary and aims

The evolutionary processes which select for and maintain structured diversity in *P. falciparum*, both at the level of individual parasite genomes and in populations of parasites, are not well understood. Throughout this thesis, I argue that diversity is essential for understanding several of aspects of *P. falciparum* epidemiology. I investigate the potential for immune selection to generate and maintain antigenic diversity, and I seek to elucidate the epidemiological impact of these processes. The specific aims of this thesis are:

- Characterise the epidemiological implications of structured diversity in *P. falciparum* at the population level:

    - Elucidate the interactions between host immunity, antigenic diversity and prevalence by considering the effect of immunity on *P. falciparum* transmission under various assumptions about exposure-mediated inhibition of infection.

    - Understand key ecological and epidemiological factors, which maintain and generate antigenic diversity to further elucidate how changes in these factors influence population structure, antigenic diversity and disease prevalence.

- Investigate the evolutionary processes which lead to the maintenance of structured diversity within antigen repertoires of individual parasite genomes:

    - Understand how assumptions about the action of immunity and constraints on diversification affect the selection for phenotypically diverse antigen repertoires.

– Investigate the evolutionary and epidemiological mechanisms which can lead to the exposure-dependent expression of different PfEMP1 phenotypes.

The following chapters address different aspects of these aims. In Chapter 2 I explore the effect of antigenic diversity in parasite populations on the relationship between transmission and prevalence. I illustrate some of the limitations which arise from common assumptions about the effect of immunity on the parasite's transmission potential. Chapter 3 builds on this by moving away from the traditional treatment of diversity as a static property and instead allows diversity to be generated and maintained through the interaction between parasites and the host's immune system. This allows me to explore some of the evolutionary, epidemiological and ecological factors which determine site-specific levels of parasite diversity.

Chapter 4 is concerned with the evolution of structure within *var* gene repertoires. Specifically, I investigate the maintenance of partitioned antigen repertoires by considering the selection pressure that arises from immune interference, within-host fitness and constraints on diversification. Chapter 5 seeks to understand the interplay between population level diversity and structured diversity in parasite repertoires. Specifically, I consider how constraints on diversification affect the population structure of parasites that utilise partitioned antigen repertoires. I also investigate how structured antigenic diversity and antigen expression order can explain exposure-dependent distributions of infection pathology and acquisition of immunity. Finally, Chapter 6 discusses some of the wider implications of the main results of this work.

Chapter 2

Antigenic diversity, immunity and the transmission-prevalence relationship

2.1 Background

2.1.1 The transmission-prevalence relationship

The relationship between transmission and *P. falciparum* prevalence is driven by complex processes at different stages throughout the parasite's apicomplexan lifecycle. It is critical to understand this relationship to be able to effectively plan, execute, and assess the impact of control strategies (Shaukat, Breman, and McKenzie, 2010). The entomological inoculation rate (EIR) is often used as a measure of transmission intensity because it provides a direct measure of human contact with infectious agents (Shaukat, Breman, and McKenzie, 2010). It is defined as the number of infectious bites per human host received in a given time interval. Transmission can range from intermittent seasonal outbreaks to hyper-endemicity, with regional annual (EIR) reported from less than one to greater than 800, and parasite prevalence
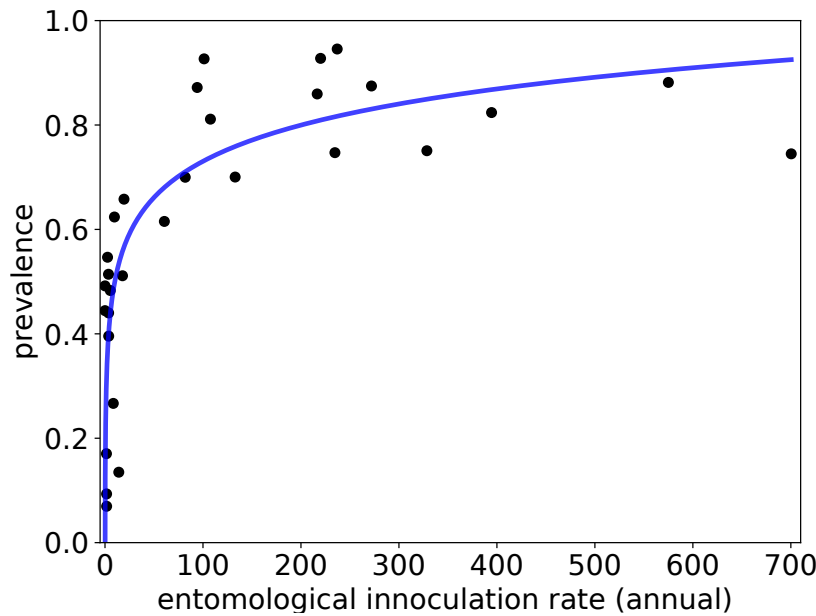
FIGURE 2.1: The basic relationship between transmission and prevalence involves a rapid increase in prevalence at low transmission intensity, which eventually plateaus at moderate to high transmission intensity. Figure has been reproduced from Beier et al (Beier, Killeen, and Githure, 1999) and is based on data from 30 African field sites.

which can range from near zero to above 90% (Beier, Killeen, and Githure, 1999).

The basic relationship between transmission intensity and prevalence is characterised by an extremely rapid increase in prevalence from low to moderate transmission intensity, which plateaus at high transmission intensity (Beier, Killeen, and Githure, 1999; Smith et al., 2005; Ebenezer, Noutcha, and Okiwelu, 2016). Figure 2.1, adapted from Beier et al (Beier, Killeen, and Githure, 1999), shows a logarithmic fit to data from 30 African field sites. The highly non-linear nature of this relationship means that large reductions in EIR are often generally required to make any appreciable impact on *P. falciparum* prevalence.

2.1.2 Diversity and *P. falciparum* epidemiology

The large antigenic diversity exhibited by *P. falciparum* (Barry et al., 2007; Day et al., 2017; Chen et al., 2011; Aguiar et al., 1992) and its interaction with the host immune system is central to many aspects of *P. falciparum* epidemiology. This diversity

underpins the parasite's ability to reinfect hosts throughout their lives and establish long infections by evading adaptive immune responses (Craig and Scherf, 2001; Kyes, Kraemer, and Smith, 2007; Kraemer and Smith, 2006). Furthermore, parasite populations appear to be structured into strains with little overlaps between their *var* gene repertoires (Kraemer et al., 2007; Rask et al., 2010; Day et al., 2017; Chen et al., 2011; Barry et al., 2007; Aguiar et al., 1992), which minimises immune interference (Gupta et al., 1996) and enables repeated infection by different and antigenically distinct parasites.

While infections induce the production of antibodies that recognise and protect against homologous parasites (Bull et al., 2000; Ofori et al., 2002), the extensive diversity in natural *P. falciparum* populations mean that hosts probably encounter novel antigen variants throughout their lives. As a result, naturally acquired immunity has been proposed to be acquired through the piecemeal accumulation of variant specific antibodies, whereby susceptibility to an infection is determined by the ability of parasites to express antigens that correspond to gaps in this protective repertoire (Bull et al., 1999; Bull et al., 2000). On the other hand immunity to severe disease is acquired relatively quickly. It has suggested that severe disease may be caused by a subset of antigen variants with reduced diversity (Gupta et al., 1994; Gupta et al., 1999; Langhorne et al., 2008; Doolan, Dobaño, and Baird, 2009).

There is considerable uncertainty about the effect of previous exposure on the transmission potential of subsequent infections. Estimates of the duration of infection (DOI) can vary widely from several weeks to over a year, but most studies suggest that it decreases with age and/or exposure (Felger et al., 2012; Bruce et al., 2000; Smith et al., 1999a; Bousema et al., 2010; Kitua et al., 1996; Bekessy, Molineaux, and Storey, 1976; Smith and Vounatsou, 2003). Furthermore, the duration of infection is

not necessarily equivalent to the infectious period. That is, we still lack a full understanding about the within-host dynamics of the transmissible stages of *P. falciparum* and how these are influenced by host immunity under repeated exposure.

The underlying processes which determine the relationship between transmission intensity and prevalence are complex. We argue that assumptions regarding how host immunity interacts with, and is influences by, parasite diversity are an essential and often neglected component in current models of *P. falciparum* epidemiology. Given that the parasite's ability to reinfect hosts, maintain long infections and establish acute versus chronic infections are determined by these interactions, it is important to gain a better understanding about how commonly used assumptions influence the predicted relationship between diversity, immunity and prevalence.

Modelling highly diverse and structured parasites, such as *P. falciparum*, is challenging however. Most previous theoretical work have simplified the interaction between diversity and immunity by treating diversity as an abstract quantity encoded in the probability or rate of acquiring some semi-immune state (e.g. (Dietz, Molineaux, and Thomas, 1974; Struchiner, Halloran, and Spielman, 1989)). Here we discuss three simple models of increasing complexity and explore the effect of adopting different assumptions about diversity and immunity. We examine the ability of each model to fit field transmission and prevalence data from 30 African field sites collated by Beier *et al.* (Beier, Killeen, and Githure, 1999). Specifically, we consider three simple ways in which immunity affects transmission: (1) no effect on transmission, (2) transient immunity, and (3) binary transmission-reducing immunity. Finally we describe a modelling framework which considers the epidemiological effects of gradually acquired immunity to diverse parasite populations. Given the considerable uncertainty in the impact of naturally acquired immunity on onward transmission, we tested the ability of this model to capture the characteristic

transmission-prevalence curve using a range of immuno-epidemiological assumptions.

2.2 Methods

We constructed a suit of deterministic modelling frameworks to explore different assumptions about the effect of immunity on transmission, and tested the ability of each to reproduce observed transmission-prevalence relationships. We considered a broad range of epidemiological parameters to reflect the uncertainty surrounding the effect of exposure on transmission.

2.2.1 Exposure independent transmission model

In the simplest model, transmission is assumed to be independent of exposure. Within this SIS model framework the human population is split into two compartments based on their infection status: susceptible and infected. Changes in the host population are described by the following set of equations:

$$\frac{dS}{dt} = -\beta SI + \sigma I + \mu(1 - S) \tag{2.1}$$

$$\frac{dI}{dt} = \beta SI - \sigma I - \mu I \tag{2.2}$$

where S and I are the proportion of hosts which are susceptible and infectious, respectively, β is the transmission rate, σ is the recovery rate ($\frac{1}{\sigma}$ gives the duration of infection in days) and μ is the birth and death rate (which assumes the population is constant). Mosquitoes are not expected to have an impact on the qualitative behaviour and are therefore not considered in this framework.

2.2.2 Temporary immunity model

In the second framework we assumed that recovery from infection results in transient immunity, which protects the hosts from immediate reinfection. This is modelled by adding a 'resistant' compartment, R, to the above framework (section 2.2.1). The average resistant period is defined by $\frac{1}{\alpha}$ and does not change with host age or exposure. The system is described by the following set of equations:

$$\frac{dS}{dt} = -\beta SI + \alpha R + \mu(1 - S) \tag{2.3}$$

$$\frac{dI}{dt} = \beta SI - \sigma I - \mu I \tag{2.4}$$

$$\frac{dR}{dt} = \sigma I - \alpha R - \mu R \tag{2.5}$$

2.2.3 Binary immunity model

In the third model, immunity is assumed to be acquired in a binary fashion by using a similar framework as described by Dietz *et al.* (Dietz, Molineaux, and Thomas, 1974). We assumed that immunity does not confer sterile protection but rather leads to changes in the recovery and transmission rates. Upon recovery there is a small probability that hosts transition into the a semi-immune state. Infection and recovery in semi-immune hosts are guided by different transmission and recovery rates, such that semi-immune hosts are always associated with a decrease in susceptibility but may either increase or decrease their recovery rate. The model comprises four epidemiological compartments depicted in figure 2.2 with equal birth and death rates to maintain a constant population. The transition between compartments is described as follows:
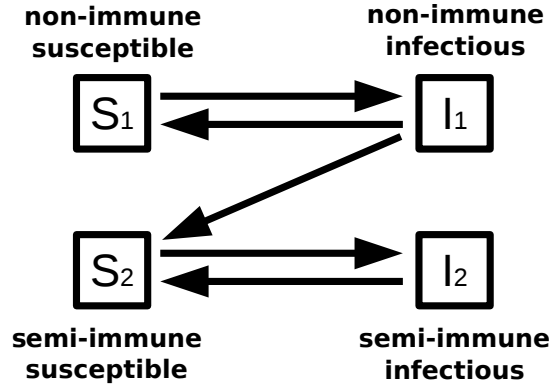
FIGURE 2.2: The fluxes between compartments in the binary immunity model. The host population is compartmentlaised according to their infection status and immune status. Susceptible non-immune hosts (S_1) can become infected (I_1). Upon recovery hosts either return to the non-immune susceptible compartment or, with a small probability (λ) acquire some degree of immunity and enter the semi-immune susceptible state (S_2). Semi-immune hosts can become infected move into the semi-immune infectious compartment (I_2). Semi-immune dynamics are governed by different transmission and recovery rates from non-immune individuals.

$$\frac{dS_1}{dt} = -\beta S_1 (I_1 + I_2) + \sigma_1 I_1 + \mu(1 - S_1) \tag{2.6}$$

$$\frac{dI_1}{dt} = \beta S_1 (I_1 + I_2) - \sigma_1 I_1 - \lambda I_1 - \mu I_1 \tag{2.7}$$

$$\frac{dS_2}{dt} = -\gamma \beta S_2 (I_1 + I_2) + \sigma_2 I_2 + \lambda I_1 - \mu S_2 \tag{2.8}$$

$$\frac{dI_2}{dt} = \gamma \beta S_2 (I_1 + I_2) - \sigma_2 I_2 - \mu I_2 \tag{2.9}$$

where S_1, and I_1 are the proportion of non-immune hosts which are susceptible and infectious. S_2 and I_2 are the proportion of susceptible and infectious semi-immune hosts, respectively. β is the baseline transmission rate, which in immune hosts is scaled by γ with $\gamma, \in [0,1]$. σ_1 and σ_2 are the recovery rates for non-immune and semi-immune hosts, respectively.

2.2.4 Cumulative exposure model

In our fourth model we assume that immunity is gained gradually as a function of cumulative exposure and population level diversity. As hosts become infected and recover they move along a 'ladder' of compartments, which represent different levels of host exposure (depicted in Figure 2.3). In this way cumulative exposure can be tracked in the host population, and transmission and recovery rates can be defined based on these exposure levels. The number of compartments in the model is given by $2n + 1$ where n is the number of exposure levels. This approach enables diverse parasite populations to be efficiently modelled by considering their cumulative effect on host immunity and immune-mediated changes in transmission. The advantage of this framework is that the complexity of the model grows linearly with n. This is in contrast to traditional multi-strain models (e.g. (Gupta et al., 1996)), which grow exponentially as diversity increases. The main drawback to modelling diversity through levels of exposure in our framework is that it constrains us to assume that the parasite population is well mixed, and any effect of heterogeneity in the parasite population is averaged out in the exposure of hosts. While stratifying the host population by age would potentially allow us to gain insight into important age-dependent epidemiological patterns, including age stratified prevalence and duration of infection, in the interest of computational tractability the host population was not stratified by age in addition to exposure. Cumulative exposure (i) can act as a proxy for host age, although differences in the transmission rate or duration of infection at different exposure levels means that caution should be used when interpreting cumulative exposure in this way.

We use S_i to indicate susceptible hosts with a history of exposure to i previous infections, and similarly I_i for infectious hosts with i previous infections (excluding
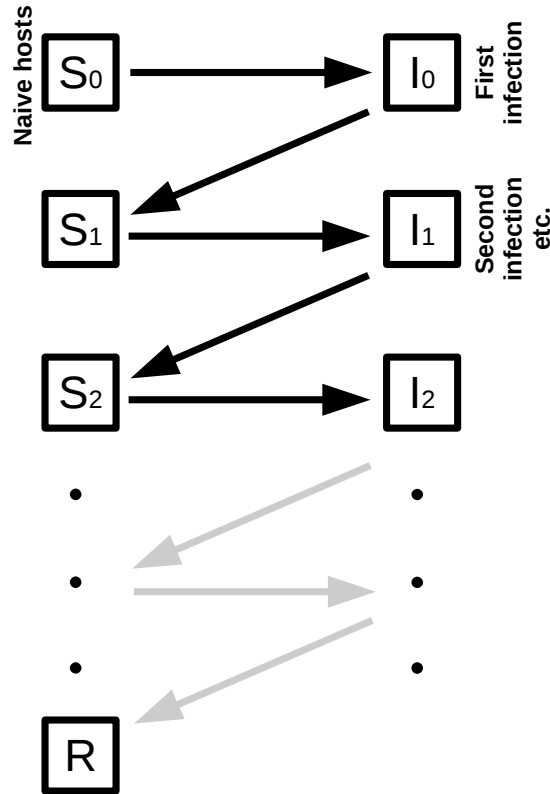
FIGURE 2.3: The compartments of the cumulative exposure model. As hosts become infected and recover they move through the exposure compartments (top to bottom). The cumulative exposure of a host (indicated by subscript) determines the transmission and recovery rate such that these parameters can be defined as arbitrary functions of exposure.

the current infection). Therefore $i = 0$ refers to hosts with no previous exposure. For the general case of n exposure levels the model is defined by the following equations:

$$\frac{dS_0}{dt} = -\beta_0 S_0 \sum_{j=0}^{n-1} I_j + \mu(1 - S_0) \tag{2.10}$$

$$\frac{dS_i}{dt} = -\beta_i S_i \sum_{j=0}^{n-1} I_j + \sigma_{i-1} I_{i-1} - \mu S_i \quad i = 1, 2, ...n \tag{2.11}$$

$$\frac{dI_i}{dt} = \beta_i S_i \sum_{j=0}^{n-1} I_j - I_i(\sigma_i + \mu) \tag{2.12}$$

$$\frac{dR}{dt} = \sigma_{n-1} I_{n-1} - \mu R \tag{2.13}$$

where S_i and I_i are the proportion of susceptible and infectious hosts, respectively, with an exposure history of i previous infections. R is the proportion of the population with complete sterile immunity. β_i and σ_i are the transmission and recovery rates for hosts with i previous infections. Finally, the birth and death rate is given by μ (note that deaths due to disease are ignored).

No closed form analytical solution is possible because changes in each susceptible compartment is coupled with all infected compartments, creating a feedback mechanism in which an increase in a single infected compartment increases the incidence in all susceptible compartments regardless of exposure. The force of infection for a given susceptible sub-population with cumulative exposure i is therefore given by

$$\lambda_i = \beta_i \sum_{j=0}^{n-1} I_j \tag{2.14}$$

This simple framework can be used to explore the effect of exposure on transmission and prevalence under a wide range of assumptions about immunity. The effect of immunity and exposure is encoded by defining functions for β and σ which map exposure level, i, to transmission and recovery rates. A similar framework has been used previously (Gupta, Swinton, and Anderson, 1994) to investigate different possible mechanisms as to how immunity could be acquired (e.g. in a strain-specific way, or through the gradual boosting of broadly acting immunity over repeated exposure), and how this would effect the age-distribution of disease and prevalence within a single population. Here, we focus on the effect that assumptions about immunity and diversity have on the transmission-prevalence relationship for populations over a range of transmission intensities. Furthermore, because our model focuses on host exposure rather than directly modelling host immunity to individual strains, we avoid making specific assumptions about the number of strains in

circulation, their antigenic composition or relation to one another. We treat n as a proxy for antigenic diversity such that after n infections it is assumed that a host has seen all the circulating strains in the parasite population.

The mean number of infections that hosts experience over their lifetime at equilibrium can be computed through the equilibrium prevalence. This is calculated by summing over the mean time between infections and mean duration of infection for each exposure category, i, until $i = n$ or the mean life expectancy is exceeded. In pseudo-code this is given as:

```
meanInfections = 0
daysRemaining = lifeExpectancy
for i in range [0, n):
    timeToInfection = 1 / (beta[i] * prevalence)
    durationOfInfection = 1 / sigma[i]
    timeTaken = timeToInferction+durationOfInfection
    if daysRemaining >= timeTaken:
        meanInfections += 1
        daysRemaining -= timeTaken
    else:
        meanInfections += (daysRemaining/timeTaken)
        return meanInfections
```

The incidence rate for a given sub-population with cumulative exposure i indicates the exposure-stratified contribution to infection:

$$\epsilon_i = I_i \sum_{j=0}^{n-1} \beta_j S_j. \tag{2.15}$$

We considered three ways in which immunity can affect transmission: the acquisition of transmission reducing immunity, cross-immunity acting to reduce transmission, and exposure dependent reduction in the duration of infection.

Acquisition of transmission-reducing immunity

To investigate the effect of exposure-mediated transmission-reducing immunity we defined the transmission rate using a simple linear function of cumulative exposure given as

$$\beta_i = (1 - \frac{i}{n})\beta_0, \tag{2.16}$$

where the number of exposure levels, n, is taken to be a measure of the population level antigenic diversity, β_0 is the baseline transmission rate and i is the cumulative exposure (number of prior infections) of the host. The rate of change in transmission rate is thus a function of the diversity in the population.

Cross-immunity

To model cross-immunity which could arise, for example due to shared epitopes or overlapping antigen repertoires, we defined a transmission rate using an inverse exponential function of exposure, such that hosts experience a disproportional decrease in susceptibility for early infections, i.e.

$$\beta_i^* = e^{-i\alpha}, \tag{2.17}$$

where α controls the rate of decay in transmission rate and can therefore be thought of as the strength of cross-immunity. The acquisition of cross-reactive immunity is dependent on the cumulative exposure i, but independent of n. Cross-immunity is

assumed to act in conjunction with the exposure immunity described above, such that the final transmission rate for a host with i prior infections is given as

$$\beta_i = \beta_i^*(1 - \frac{i}{n})\beta_0, \tag{2.18}$$

The special case where $\alpha = 0$ leads to a linear change in transmission rate with exposure and is equivalent to a piecemeal acquisition of immunity.

Exposure-dependent duration of infection

The third scenario considers the effect of exposure on the rate of parasite clearance. We assume that gradually developed immunity decreases the duration of infection. This is implemented using an exposure-dependent recovery rate, σ_i, defined as a function of total diversity (n) and cumulative exposure (i):

$$\sigma_i = \frac{i}{n}(\sigma_{n-1} - \sigma_0) + \sigma_0, \tag{2.19}$$

where σ_0 is the baseline recovery rate and σ_{n-1} is the maximum recovery rate. The change in recovery rate with each exposure therefore scales with parasite diversity. The change in σ_i with exposure is determined by the difference between σ_0, σ_n and by n, while the start and end points are determined by σ_0 and σ_n respectively.

2.2.5 Parameter range

We considered a broad range of parameter values to encompass most estimates of the duration of infection and transmission intensity found in the literature (summarised in table 2.1). The models described above do not explicitly consider vectors and hence we approximate daily transmission rates through annual EIR by assuming that every infectious bite results in a new infection, i.e. $\beta_0 = \frac{EIR}{365}$.

TABLE 2.1: Parameter range considered across all four modelling frameworks. Note that not all models use all of these parameters.

Parameter	Description	Baseline value [range]
β	transmission rate	$[0.001, 2.0]$
$\frac{1}{\sigma}$	duration of infection	$[7, 365]$ (days)
$\frac{1}{\mu}$	host life expectancy	$60[30, 90]$ (years)
$\frac{1}{\alpha}$	temporary immune period (temporary immunity model only)	$[1, 365]$ (days)
λ	rate of semi-immune acquisition (binary immunity model only)	$[0, 1]$
n	number of exposure categories (cumulative exposure model only)	$[1, 501]$
α	strength of cross-immunity (cumulative exposure model only)	$[0, 0.5]$

We note that while humans do not naturally acquire sterile immunity to *P. falciparum*, hosts which experience n infections in our cumulative exposure model are protected from further infections. The purpose of the model is to provide a tool which can be used to assess how parasite diversity, and the interaction of parasite diversity and immunity, affect the transmission-prevalence relationship. Therefore, we chose values of n which include both low and high diversity ($n = [1, 501]$) to characterise the effect of diversity over a broad range of conditions. Furthermore, it was important that the upper range of n could capture systems in which most hosts do not reach the end of the 'ladder' and hence do not gain sterile immunity. The chosen upper bound of $n = 501$ represents a compromise between biological validity and computational tractability. Figure 2.4 shows the proportion of hosts with sterile immunity using the 'worst-case' combination of parameters at maximum diversity (when $n = 501$) for some preliminary simulations.
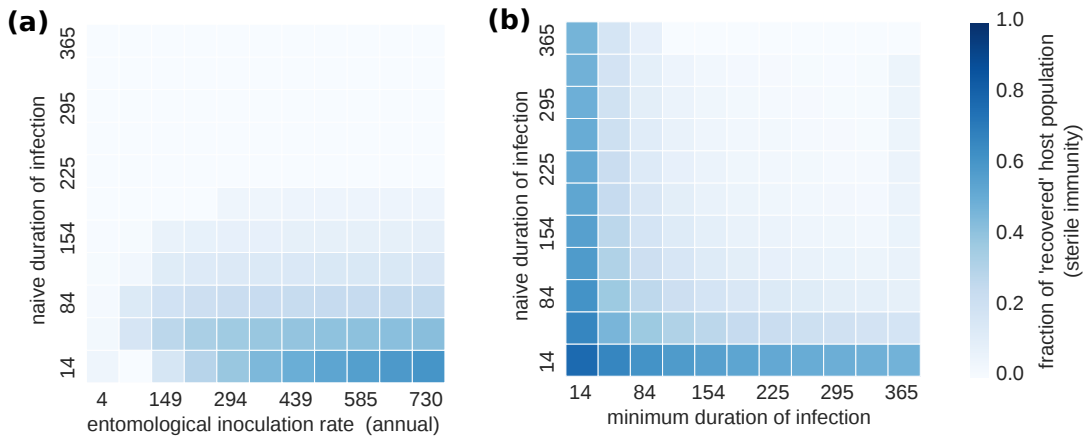
FIGURE 2.4: Preliminary simulations showing the proportion of hosts which acquire sterile immunity (enter the 'recovered' compartment, see figure 2.3) using maximum diversity ($n = 501$). Host life expectancy is assumed to be 60 years. Other parameters have been selected to maximise the recovered fraction. **(a)** Piecemeal immunity, which prevents transmission of strains which hosts have previously been exposed to. **(b)** Immunity acting to reduce duration of infection with repeated exposure (annual entomological inoculation rate is constant at 730). Recovered fraction is shown for combinations of the naive duration of infection ($\frac{1}{\sigma_0}$) and minimum duration of infection ($\frac{1}{\sigma_{n-1}}$), which is only achieved when hosts have experienced $n - 1$ infections. Since sterile immunity has not been observed in natural transmission settings, the proportion of the host population which gain 'recovered' status (see 2.3) should ideally be low.

2.2.6 The transmission-prevalence relationship

In order to quantitatively assess the ability of each modelling framework to accurately capture the transmission-prevalence relationship, parameters were fit to field data from Beier et al (Beier, Killeen, and Githure, 1999). This was done using a method called differential evolution (Storn and Price, 1997) to approximate the global minimal sum of squared errors. The high density of data points in the low (<30) EIR range may bias the fits to this region more strongly. It was decided not to attempt to correct for this because this region represents the part of the transmission-prevalence curve, which is most important in the context of malaria control; i.e. it is the region in which control can cause meaningful change. Akaike information criterion (AIC) was calculated for the best fit from each model framework by assuming assuming residuals came from a Gaussian distribution with constant variance.

The increase in computational complexity in the cumulative exposure framework for large n made fitting to the full Beier dataset impractical. In this case we constructed a simplified transmission-prevalence curve which captures the main features of the relationship by selecting four data points. The selected points are indicated in red in figure 2.5. This provides a computationally tractable method to approximate errors during the fitting process. The best fitting set of parameters were then used with the full Beier dataset to calculate a final sum of squared errors and allowed comparisons across frameworks.
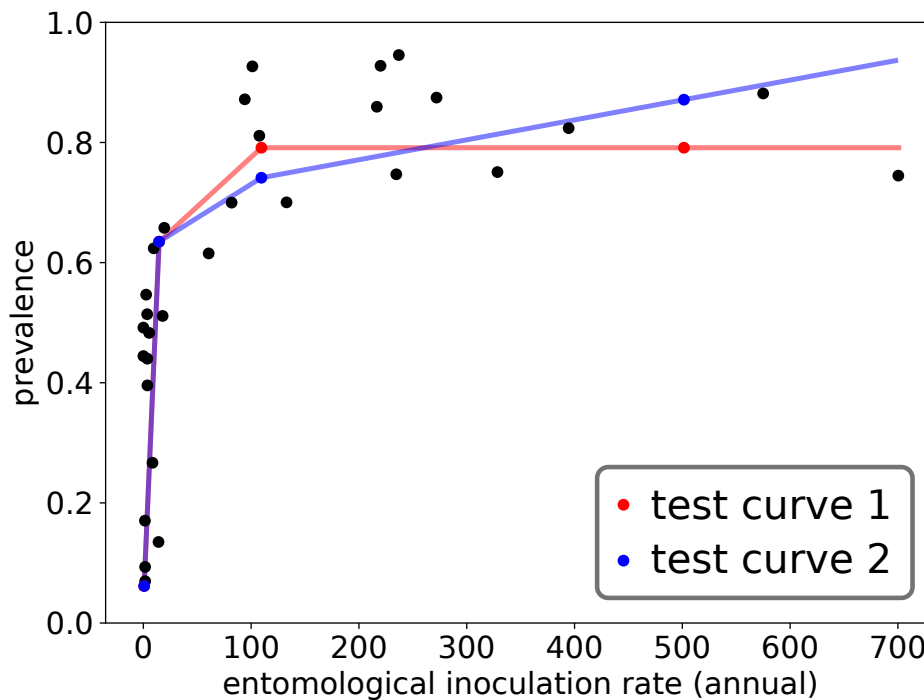
FIGURE 2.5: The transmission-prevalence relationship is approximated by four test points in order to make fitting the cumulative exposure model computationally tractable. Two sets of four test points were considered (shown in red and blue). Transmission and prevalence data from 30 field sites across sub-Saharan Africa are shown in black (Beier, Killeen, and Githure, 1999).

The test points were chosen by eye to produce two test curves which capture the important features of the transmission-prevalence relationship. Note that test points do not include specific data points in the Beier dataset. This is intended as a qualitative representation of the general relationship. Differences in the particular test

points selected will undoubtedly yield small changes in the predicted parameters. However, as the focus of the study is not to find the specific parameters which fit this particular dataset, but rather to identify processes which may be important in shaping the transmission-prevalence relationship more generally, small differences in predicted parameters are not of great interest.

The second test curve is used as a means to test the robustness of the model quality to changes in the test curve. Beier *et al.* has shown previously that the transmission-prevalence relationship can be approximated with a logarithmic function (Beier, Killeen, and Githure, 1999), and so the second test curve was specifically selected to be similar to the fit used by Beier.

The use of only four points in each test curve results in a very sparse dataset. However, since equilibrium prevalence can only increase with increases in transmission intensity, simulated transmission-prevalence curves are constrained to monotonically increasing functions of transmission intensity. This limits the paths that can be taken between test points, and there can be no local minima or maxima. Hence, relatively few test points are needed to define a the transmission-prevalence relationship well. Four test points is the minimum required to capture the two most important characteristics of the relationship: a steep initial increase in prevalence followed by a plateauing.

2.3 Results

2.3.1 Exposure-independent transmission

The exposure independent transmission model assumes that cumulative exposure does not lead to any change in the recovery rate (σ) or transmission rate (β) and that hosts can be reinfected throughout their lives. When we fit the model to estimate σ, the best fits are found when the duration of infection is around 190 or 240 days (figure 2.6a). This simple framework can easily generate the initial region of the transmission-prevalence curve where prevalence increases rapidly with transmission rate. However, as can be seen from figure 2.6b, it struggles to accurately reproduce prevalence at the higher range of EIRs. As transmission rate increases prevalence increases asymptotically to $1 - \mu$. Equilibrium prevalence was relatively insensitive to host life expectancy ($\frac{1}{\mu}$) over the range considered ([30, 90]). However, lower life expectancies did result in marginally better fits because this meant that more infected hosts were lost due to death and replaced with naive hosts. This limited the maximum prevalence as $\beta \to \infty$. The effect was very small for the range in host life expectancy considered. Hence, since host life expectancy is well known and on the order of decades not weeks or years, this framework cannot accommodate both the steep rise in prevalence seen at low transmission and the plateauing at moderate to high transmission intensity.

2.3.2 Temporary immunity

Next we added temporary immunity to the model, which acts to protect hosts from reinfection for a short period of time. The level at which prevalence plateaus was predominantly determined by the duration of transient immunity, controlled by α. Longer periods of immunity (smaller α) limits the proportion of infectious hosts at
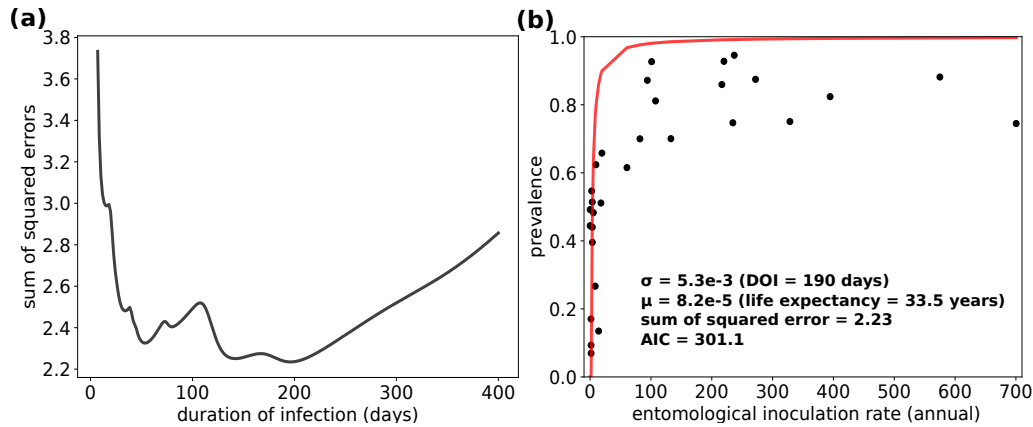
FIGURE 2.6: The exposure independent transmission model is a simple SIS framework and cannot explain both the steep rise in prevalence at low transmission intensity and the mid-high plateauing at higher intensities. **(a)** Sum of squared errors for different durations of infection (DOI), $\frac{1}{\sigma}$. Host lifespan is $\frac{1}{\mu} = 33.5$ years. The error landscape has several nearly equal minima which result in a poor overall fit in all cases. **(b)** Simulated prevalence (red) plotted with the data (black) for the best fit. The best fitting parameter values, sum of squared error and Akaike information criterion (AIC) are shown as annotations.

equilibrium. We fit the model for σ, μ and α, and both the initial rise and plateau regions of the transmission-prevalence relationship could be reproduced well. This resulted in smaller sum of squared error than were achieved for the exposure independent transmission model (1.87 versus 2.23), and appears to match the general transmission-prevalence relationship more closely. Figure 2.7a shows the error landscape of simulated transmission-prevalence curves using different σ and α and highlights the interaction between infection length and duration of immunity. The lowest error region predicts long infectious periods greater than 300 days (facilitating the rapid rise in prevalence for small increases in transmission intensity) and moderate duration of immunity (safeguarding a fixed minimum proportion of the host population from infection at equilibrium). The best fit, shown in figure 2.7b, tends to underestimate prevalence in high intensity settings. Furthermore the best fits are found at the extreme ends of the considered parameter ranges for μ and σ. These are biologically unrealistic values and perhaps reflects that assumptions inherent to
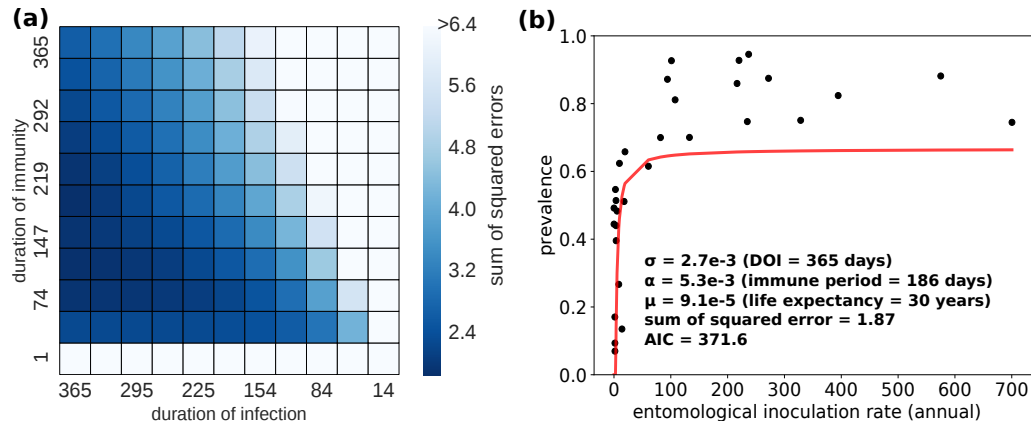
FIGURE 2.7: Temporary immunity can place an upper limit on prevalence at high transmission intensity. The error landscape for duration of infection ($\frac{1}{\sigma}$) and duration of immunity ($\frac{1}{\alpha}$) is shown in **(a)**, where the best fitting region of parameter space corresponds to short infections and long lasting temporary immunity. The simulated prevalence for the best fit is shown in **(b)**. Best fitting parameters, sum of squared error and Akaike information criterion (AIC) are shown as annotations.

the SIRS framework, namely that the existence of sterile immunity (albeit temporary sterile immunity) is at odds with biological observations. While it is possible that some degree of sterile immunity (transient or long term) is acquired to a sub-set of the parasite population, we know that hosts do not gain total sterile immunity.

2.3.3 Binary immunity

The third framework considered a binary model for immunity in which recovery from infection is associated with a small probability (λ) of attaining a semi-immune state that reduces host susceptibility to future infections. The duration of infection for immune hosts (σ_2) was more loosely constrained to allow for the possibility of immunity leading to longer or chronic infections, with infections potentially lasting between 14 and 730 days. Fitting the model for σ_1, λ, and γ resulted in the smaller sum of squared error (0.89) than the previous two frameworks. The model could simulate the general prevalences-transmission relationship relatively well, capturing both the initial steep rise and plateauing of prevalence (figure 2.8d).

For an intuitive understanding, consider that the two halves of the model, non-immune and semi-immune, can be viewed as two linked SIS models. The net transmission-prevalence curve is determined by mixing the effects of both, with the non-immune half being most influential in low transmission settings and the semi-immune half dominating in high transmission settings. λ can therefore be thought of as controlling the mixing of these two models. Rapid increases in prevalence can be achieved at low transmission intensity because non-immune hosts can experience long infections. σ_2 and γ give the model the flexibility needed to curtail this rapid growth in prevalence at higher transmission intensity. By mixing the effects of two SIS models in a way in which the relative importance of each depends on transmission intensity, the model is able to fit a wide range of theoretical transmission-prevalence curves, and has no problem fitting the field data (figure 2.8d). One limitation is that because the model is essentially an SIS model, prevalence increases monotonically with transmission intensity, even when virtually all of the population is semi-immune, and will approach 100% prevalence at high transmission intensity.

To explore the limitations and trade-offs between the parameters used by this framework, we visualised error landscapes between pairs of parameters. The lowest sum of squared errors were found when there is a long duration of infection in non-immune hosts (responsible for the rapid initial increase in prevalence) and a short duration of infections in semi-immune hosts (slowing the transmission-prevalence gradient at high transmission intensity). This is shown in figure 2.8a. Different combinations of duration of infection and transmission rate scaling coefficient (γ) in semi-immune hosts can produce transmission-prevalence curves which appear to match the data similarly well (2.8b). This suggests that the important aspect is the net reduction in the transmission rate from semi-immune hosts, i.e. when either $\frac{1}{\sigma}$ or γ are low. The model is best able to approximate the prevalence plateau in high

transmission settings when both $\frac{1}{\sigma_2}$ and γ are low. In order to fit the data well the rate of immune acquisition (λ) is restricted to low values; however, some degree flexibility is possible if the transmission-reducing effect of immunity is low (figure 2.8c).

While we allowed for the possibility that immunity may lead to chronic infections, this resulted in models which captured the transmission-prevalence relationship poorly. However, our model did not consider that chronic infections may be associated with reduced infectiousness to vectors or may be only one of several clinical outcomes in semi-immune hosts.
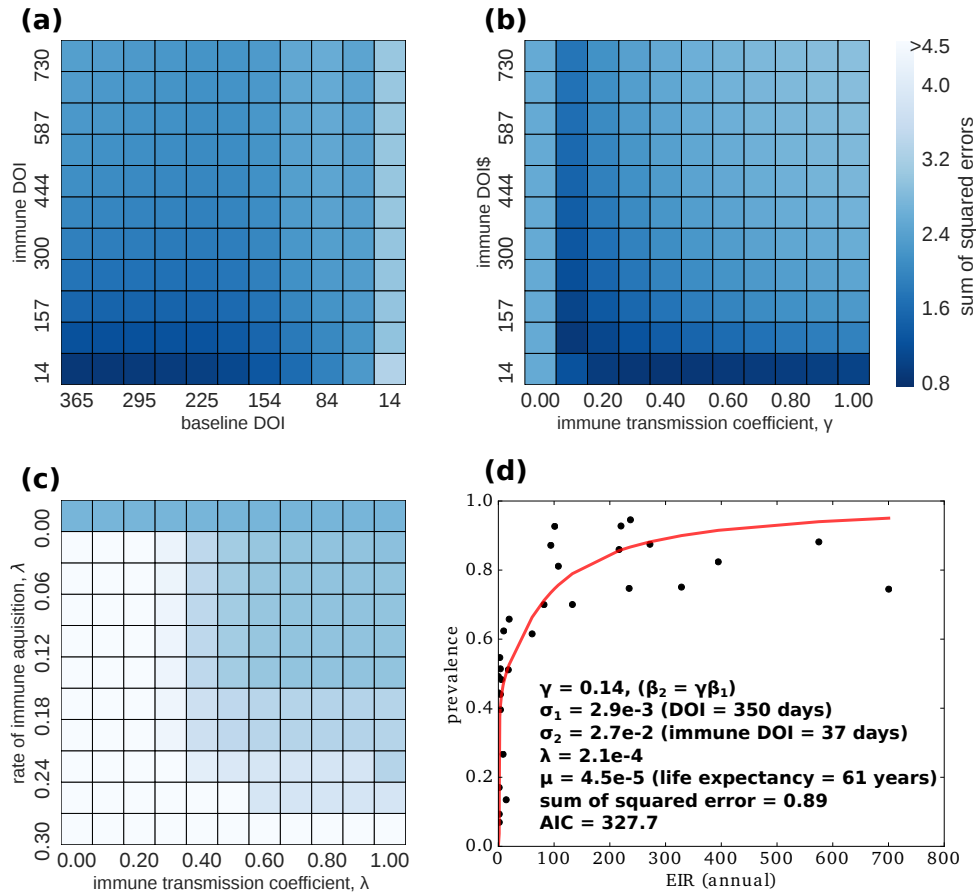
FIGURE 2.8: Error landscapes and the simulated transmission-prevalence curve for the binary immunity model. The effect of interactions between parameters on the sum of squared error is shown for immune $(\frac{1}{\gamma_1})$ and non-immune duration of infection $(\frac{1}{\gamma_2})$ **(a)**, immune duration of infection $(\frac{1}{\gamma_2})$ and the transmission scaling factor in immune hosts (α), **(b)**, and the rate of immune acquisition λ versus γ **(c)**. For each heatmap other parameters are fixed at those of the best found fit (shown in **(d)**). The best fitting parameters, sum of squared error and Akaike information criterion (AIC) are shown as annotation.

2.3.4 Cumulative exposure model

Although the previous three models assumed some form of exposure-dependent changes in susceptibility or transmissibility, they did not explicitly account for antigenic diversity. For example, in the binary immunity framework diversity diversity is effectively incorporated into the rate at which individuals transition from non-immune to semi-immune. This precludes any in-depth analysis of the effect of exposure and parasite diversity. In order for the binary immunity model to generate realistic transmission-prevalence curves, the difference in the transmission rates and duration of infection between non-immune and semi-immune hosts had to be very large (shown in figure 2.8d). In this section we build on the binary immunity model by explicitly modelling an arbitrary number of exposure levels. The epidemiological consequences of gradually acquired immunity can therefore be investigated.

2.3.4.1 Basic behaviour

To demonstrate the basic behaviour and output produced by the cumulative exposure model we first made the simplest set of immune assumptions, which we refer to as the baseline scenario. Here, we assume that the transmission rate and duration of infection are independent of cumulative exposure, i.e. $\beta_i = \beta_0$ and $\sigma_i = \sigma_0$ for $i = 1...n - 1$. After n infections, hosts develop complete sterile immunity. n therefore determines a hard threshold in cumulative exposure beyond which hosts are no longer susceptible. While humans do not acquire sterile immunity to malaria this serves as a simple baseline from which to understand more complicated behaviour.

Figure 2.9a-c shows the time-series from a simulation initialised with a small proportion of the population infected ($I_0 = 0.001$) and an otherwise fully naive host population ($S_0 = 0.999$). The total prevalence, the proportion susceptible and the

proportion immune exhibit dynamics similar to that of an SIS or SIR model. The compartments for different levels of cumulative exposure peak successively over time as the initial susceptible population are infected and recover (figures 2.9b and c). There is an increase in the magnitude and width of peaks in susceptible compartments around 750 days which coincides with the crash in the total proportion infected (seen in figure 2.9a). The corresponding increase in the average time between infections causes hosts to accumulate in these compartments. Eventually a static equilibrium is reached after approximately 4000 days, at which point compartments representing hosts with lowest exposure contain the greatest proportion of the host population. This can be seen most clearly in figures 2.9d and e, which shows a snapshot of compartment values at equilibrium, based on the same parameters as figure 2.9a-c except for $n = 100$ to better illustrate the distribution of the host population between compartments. Hosts are distributed in this manner as a natural result of each compartment receiving a flow of hosts from the compartments at the adjacent lower exposure level, forming a chain in the supply of hosts. Since transmission and recovery rates are identical between exposure levels, the only variable factor between exposure levels which influences the inward flux of new hosts at equilibrium is the number of hosts in the upstream compartment.

For the particular case where $n = 1$ our model is exactly equivalent to the canonical SIR model. As n increases, prevalence converges to that of the SIS framework. This is shown for different values of n in Figure 2.9f, where the lower and upper dotted lines show prevalence for SIR and SIS models, respectively.
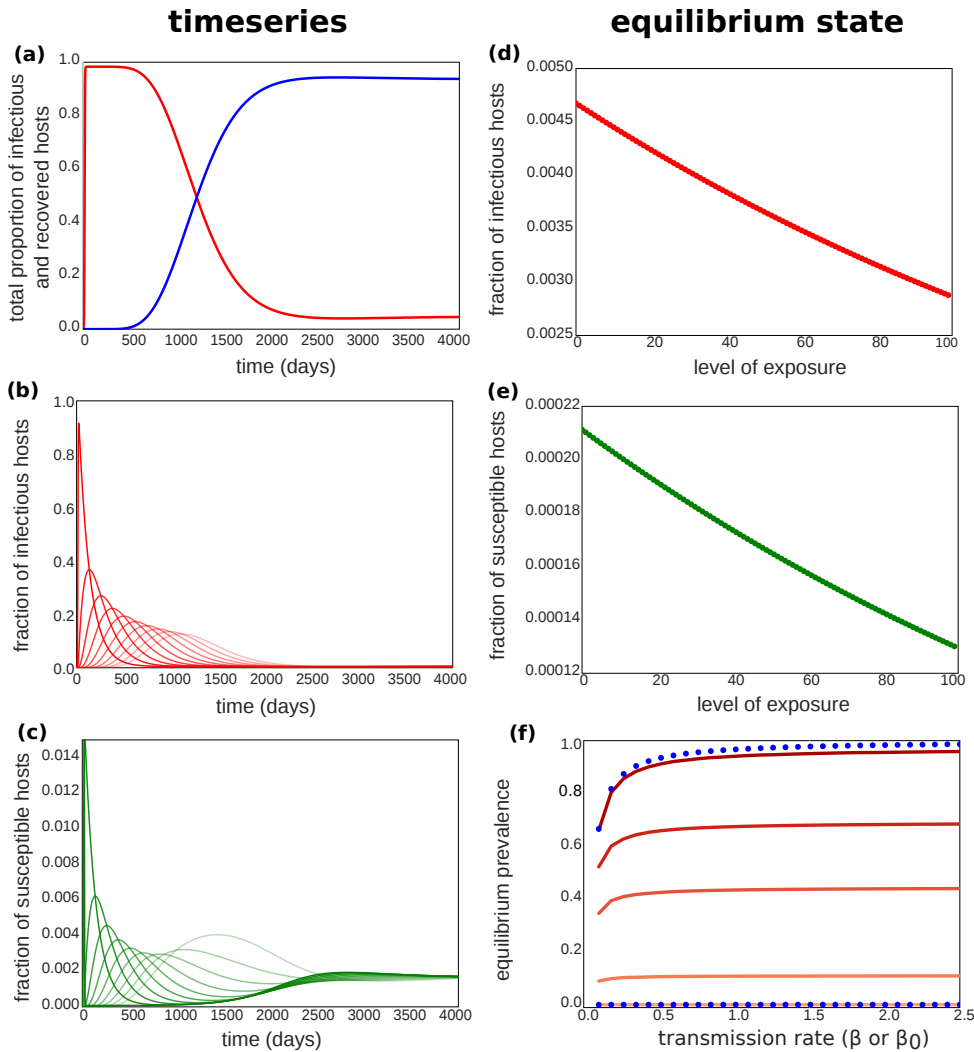
FIGURE 2.9: Timeseries (left hand side) and equilibrium compartment values (right hand side) of simulations using the cumulative exposure model when transmission and recovery rates are independent of exposure (i.e. $\beta_i = \beta_0$ and $\sigma_i = \sigma_0$). **(a)** Total infected (red) and recovered (blue) proportion of the host population over time. **(b)** and **(c)** show timeseries of the exposure stratified proportion of the host population which are infectious and susceptible, respectively (note the difference in y-axis scales). Cumulative exposure is indicated by line shading, with bolder lines indicating lower exposure. Parameters for **(a)** to **(c)** have been chosen to clearly illustrate the temporal dynamics of a typical simulation; $\beta_0 = 0.5$, $\sigma_0 = 0.0083$ (mean infection length of 120 days), $n = 10$. **(d)** and **(e)** show the equilibrium values of each infectious and susceptible compartment and were produced by a different set of simulations. The same parameters are used as before, except that $n = 100$ to better illustrate the decreasing trend in equilibrium values when cumulative exposure is greater. **(f)** Transmission-prevalence curves for a different values of n. In the special case where $n = 1$ the model is exactly equivalent to an SIR framework (lower blue dotted line). For large n the proportion of hosts which reach the higher exposure levels and fully immune state becomes very small and the model is arbitrarily equivalent to an SIS model (upper blue dotted line). From bottom to top $n = 1, 100, 500, 1000, 3000$.

2.3.4.2 The effect of antigenic diversity, *n*

When we assume that immunity acts in a piecemeal fashion hosts gradually develop immunity through exposure over multiple infections. Hosts with greater exposure are assumed to be protected against a larger proportion of the circulating strains. Hosts which experience *n* infections become completely immune because they have been exposed to (and subsequently developed immunity to) all circulating strains. Figure 2.10a shows the effect of diversity and the baseline transmission rate on prevalence. Unsurprisingly, larger diversity and/or higher transmission rates lead to higher equilibrium prevalences. Similar levels of prevalence can be obtained for different combinations of baseline transmission rate and diversity. High diversity is therefore able to compensate for lower transmission intensity.

Prevalence plateaus at higher values when diversity is larger (figure 2.10b) because greater exposure is required to acquire immunity to an equivalent proportion of the circulating diversity. At very low β_0 prevalence becomes insensitive to *n*, because hosts experience too few infections for there to be any significant advantage to the parasite population in being more diverse. We refer to this situation as being 'diversity saturated'. The inverse scenario, 'transmission saturation', occurs when the transmission rate is high enough to ensure that most hosts are exposed to most of the circulating strains over their lifetime. Therefore, increasing transmission cannot increase prevalence any further. In our model prevalence is therefore limited by different mechanisms at different transmission intensities, and diversity is an important factor in determining the point at which one mechanism dominates over the other. This is illustrated by the red (diversity saturated) and blue (transmission saturated) shaded regions for the simulated transmission-prevalence curves in figure 2.10b.
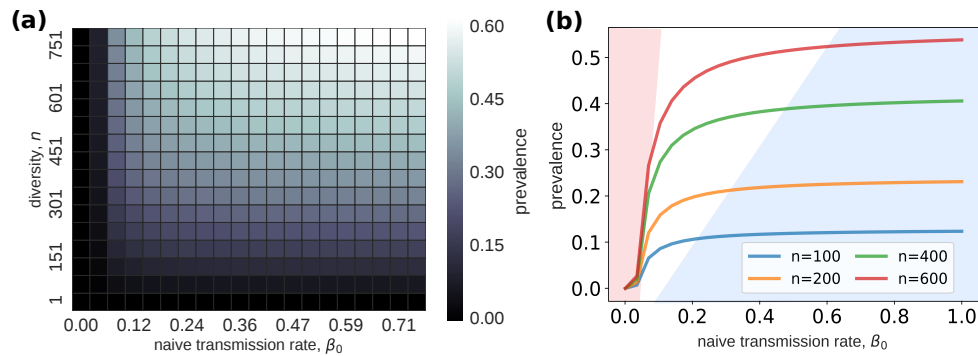
FIGURE 2.10: The effect of parasite diversity on the prevalence-transmission relationship when immunity is acquired in a piecemeal fashion. **(a)** Shows the equilibrium prevalence for different baseline transmission rates and degrees of parasite diversity. Diversity saturation can be seen in left-hand columns, which are insensitive to changes in parasite diversity. Transmission saturation can be seen in the bottom rows, which are insensitive to changes in transmission rate. (b) Demonstrates the effect of greater diversity in delaying the onset of transmission saturation, which in turn leads to higher maximum prevalence. Red shaded areas roughly illustrate diversity saturation, while the blue shaded area illustrates regions which are transmission saturated.

Next we considered cross-immunity which acts such that immunity generated during one infection reduces host susceptibility to heterologous parasites. This scenario differs from that of piecemeal immunity in that the transmission rate does not depend solely on the total diversity in the system (n) but on the ability of the host to generate cross-reactive immune responses, controlled by α (see figure 2.11). For non-zero α the greatest reduction in transmission rate occurs as a result of the first infection and each subsequent infection causes progressively smaller reductions. β_i therefore decreases asymptotically to zero. When α is low, cross-reactive protection is acquired slowly and hosts experience a greater number of infections over the course of their lives (assuming there is sufficient diversity in the parasite population). Higher α leads to rapid acquisition of broad immunity, limiting the number of infections that hosts experience and making high levels of diversity redundant. Figure 2.12a and b shows this threshold effect on the sensitivity of the system to diversity for weak and strong cross-immunity, respectively. Strong cross-immunity
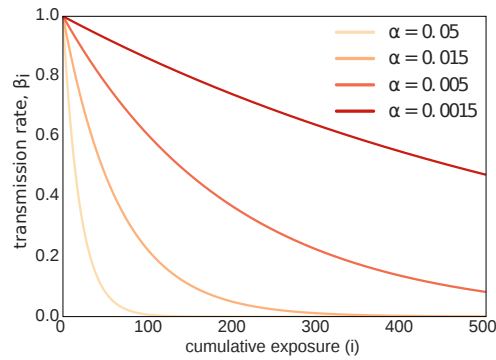
FIGURE 2.11: The effect of transmission-limiting cross-immunity on transmission rate for different strengths of cross-immunity. Larger α corresponds to stronger cross-immunity, which causes transmission rate to decrease more rapidly with cumulative exposure. $\beta_0 = 1.0$.
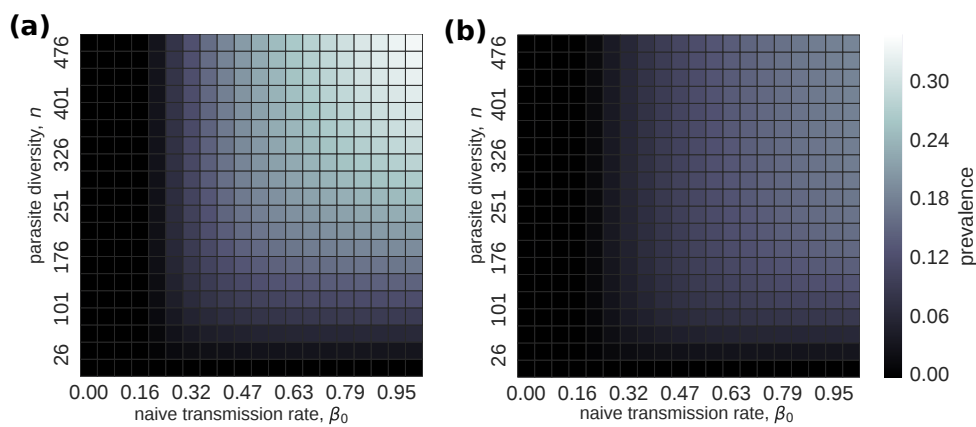
FIGURE 2.12: The effect of diversity on the transmission-prevalence relationship when exposure leads to some degree of cross-immunity. **(a)** Low cross immunity; $\alpha = 0.01$. **(b)** High cross immunity; $\alpha 0.025$. Higher cross-immunity confers protection to hosts after fewer infections. As a result prevalence is limited in the higher cross-immunity scenario (right hand pane), even when transmission intensity is very high. Duration of infection is constant at 30 days and host lifespan is set to 60 years.

prevents greater diversity from being effectively utilised by the parasite population (top right region of figure 2.12b).

In the third scenario we considered that immunity reduces the duration of infection by limiting the ability of parasites to fully utilise their antigen repertoires. We made the implicit assumption that immunity is gained in a hierarchical fashion, which essentially restricts the number of novel antigens that can successfully contribute to an infection but does not necessarily prevent infection. This was modelled by defining the recovery rate as a linearly increasing function between the naive rate
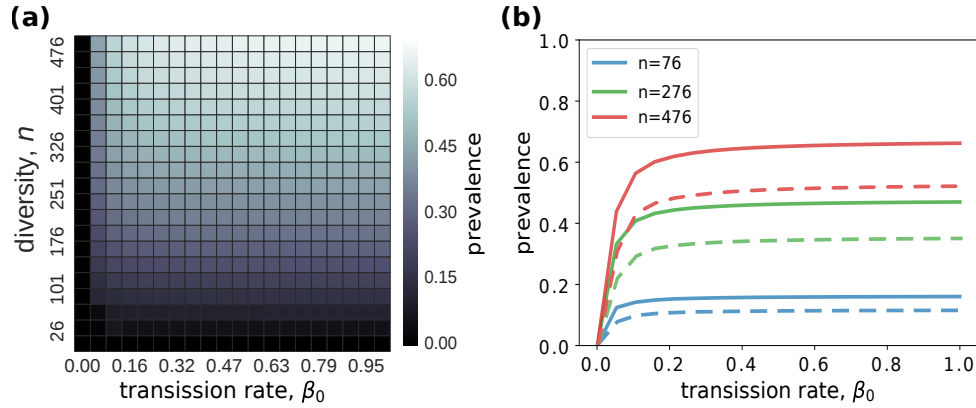
FIGURE 2.13: The effect of diversity on the transmission-prevalence relationship when immunity acts to increase the rate of parasite clearance. **(a)** Prevalence is insensitive to changes in diversity in regions of low transmission intensity (diversity saturation) and insensitive to changes in transmission intensity when diversity is low (transmission saturation). **(b)** σ_n scales the strength of immunity and determines the plateau point for prevalence (solid lines: $\sigma_n = \frac{1}{100}$, dashed lines: $\sigma_n = \frac{1}{40}$).The degree of parasite diversity, n, is indicated by colour.

(σ_0) and maximum exposure (σ_n). The $\frac{d\sigma_i}{di}$ is proportional to n and leads to rapid decreases in the duration of infection (the reciprocal of σ_i).

We find that prevalence responds in a qualitatively similar way to changes in the transmission rate and level of diversity as in the piecemeal scenario. Whether exposure leads to increasing recovery rates or decreasing transmission rates, in both scenarios it results in immune changes which curtail transmission. Prevalence is limited by transmission and/or diversity saturation, resulting in regions of parameter space in which prevalence is insensitive to changes in diversity (left side of figure 2.13a) and regions which are insensitive to changes in transmission intensity (bottom of figure 2.13a).

2.3.4.3 The transmission-prevalence relationship

The ability of the cumulative exposure model to reproduce field data for transmission and prevalence was assessed for each immune scenario by generating transmission-prevalence curves and calculating the sum of squared errors against a simplified

transmission-prevalence relationship (see Methods). We found that each set of immune assumptions resulted in models which were equally capable of generating good approximations of the observed data; however, there were large differences in the epidemiological parameters predicted by the fits. A good approximation must be able to produce both the initial period of rapid increase in prevalence at low-to-moderate transmission intensity and allow prevalence to plateau at higher transmission intensity. To help understand the constraints and implications of making different immune assumptions we visualised the effect of interactions between key epidemiological parameters on the sum of squared error (figures 2.14, 2.15, 2.16). This is discussed for each scenario below.

When immunity is acquired in a piecemeal fashion, prevalence plateaus at high transmission intensity if the system becomes transmission saturated. Diversity is therefore an important factor determining the point at which transmission saturation occurs, with higher diversity delaying its onset, as shown above. This results in a so-called pareto front in which there is a band of approximately equivalent goodness of fit for different combinations of n and σ_0 shown in figure 2.14a. There was an interaction between these two parameters in which high diversity can compensate for short DOI. In other words, transmission saturation can be induced because hosts experience a few long infections or many short infections. The best fit for this scenario predicted $n = 123$ and a DOI of 149 days (figure 2.14b). However, the existence of many near-equivalent fits along this pareto front suggests that we should not focus too much on these specific values. Instead it should be emphasised that there is a continuum of possible combinations of diversity and DOI, such that higher diversity would imply lower DOI for a given level of prevalence (and *vice versa*).

When we assumed that immunity acts to restrict the ability of parasites to utilise their full antigen repertoires and thereby to reduce duration of infection in hosts with
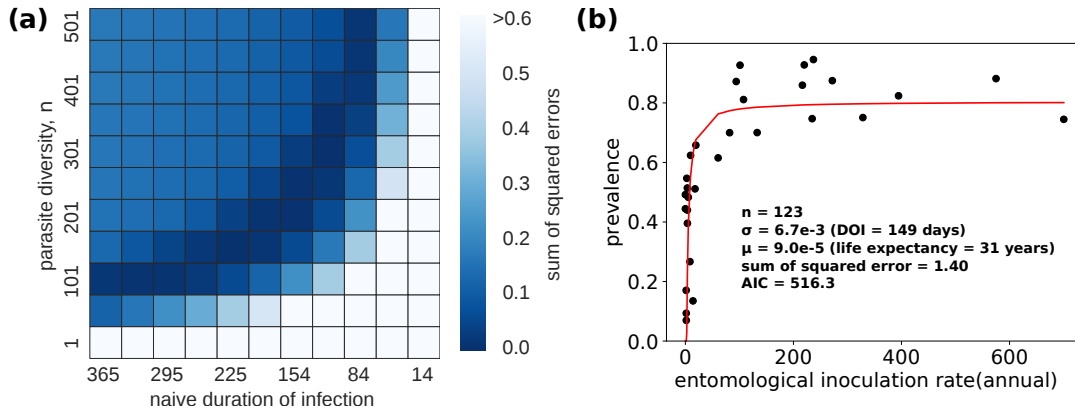
FIGURE 2.14: Fitting of the cumulative exposure model for the piece-meal acquisition of immunity scenario. **(a)** Error landscape showing regions of parameter space which result in good fits (low sum of squared error). The ability of either parasite diversity or duration of infection (DOI) to influence the point at which diversity saturation occurs demonstrates a trade-off in which either long infections or high parasite diversity lead to equivalently good fits. Host life expectancy is 60 years. **(b)** A simulated transmission-prevalence curve using the best fitting parameters under the piecemeal acquisition of immunity scenario. Best fitting parameters, sum of squared error and Akaike information criterion (AIC) are indicated as annotations. Black points show the Beier (Beier, Killeen, and Githure, 1999) data set.

greater exposure, we find a qualitatively similar trade-off between baseline recovery rate and diversity (figure 2.15a). Small σ_0, small σ_n and large n reduce the protection afforded to hosts from each infection event and therefore increase the prevalence for a given transmission intensity. This leads to a second trade-off and a pareto front involving all three parameters. For a given parasite diversity, low σ_0 (high baseline DOI) can be mitigated by lower σ_n because this increases protection gained with each infection. In other words, equivalently good fits can be attained for this scenario by long baseline infections, which induce strong immune protection in hosts, or shorter naive infections, which are associated with lower protection from future infection (figure 2.15b).

The best found fit is shown in figure 2.15c. Again, we caution against giving too much weight to the specific parameters predicted by this fit. Instead, the important point to take from these results is the existence of trade-offs which arise from the interactions between immunity, diversity and transmission.

(a)

(b)

(c)

n = 222
σ_0 = 7.5e-3 (immune DOI = 132 days)
σ_n = 1.1e-2 (naive DOI = 88 days)
μ = 6.9e-5 (life expectancy = 40 years)
sum of squared error = 1.51
AIC = 631.4
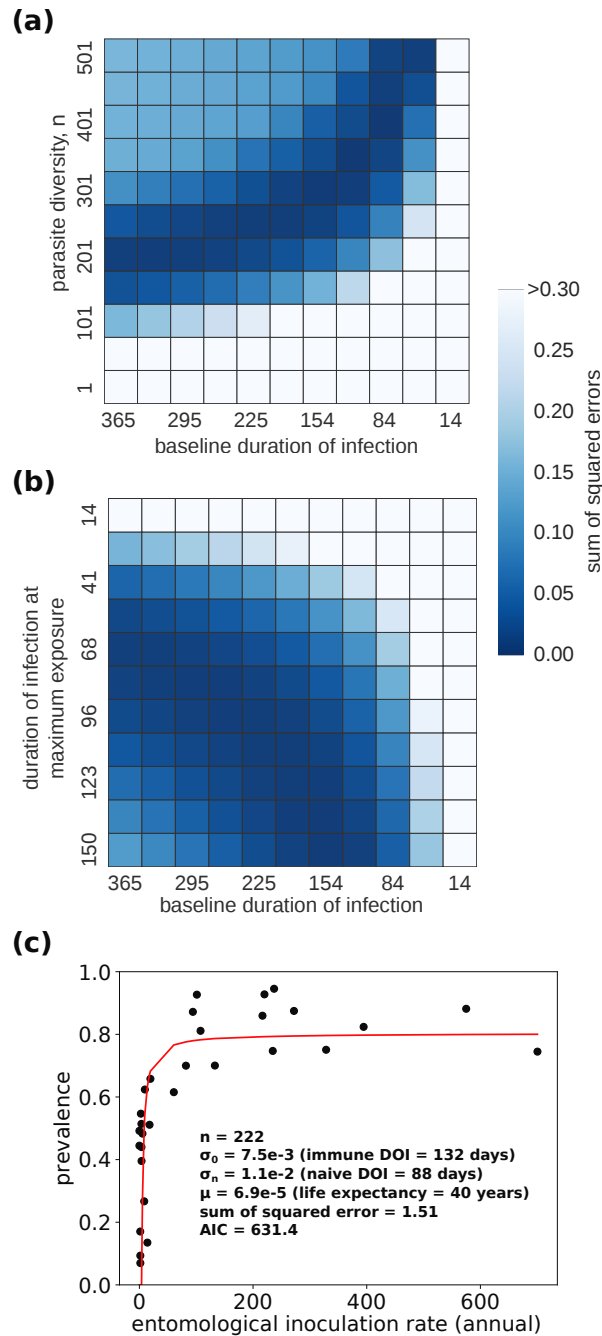
FIGURE 2.15: When immunity acts to restricts the number of antigen variants which can be expressed during an infection, greater exposure results in shorter duration of infection (DOI). The cumulative exposure model was fit using this scenario. **(a)** The interaction between parasite diversity and baseline DOI: Shorter DOI can be compensated for by greater diversity, which effectively scales the rate of immune acquisition. Parameters: $\sigma_n = \frac{1}{88}$, host life expectancy is 60 years. **(b)** The interaction between baseline DOI and DOI at maximum exposure. The difference between these two parameters acts in a similar way to diversity, scaling the rate of immune acquisition. Parameters: $n = 222$, host life expectancy is 60 years. **(c)** Simulated transmission-prevalence curves using best fit parameters. Black points show the Beier dataset (Beier, Killeen, and Githure, 1999). Best fitting parameters, sum of squared error and Akaike information criteria (AIC) are shown as annotations.

When we considered cross-immunity, hosts gain immunity more rapidly causing the onset of transmission saturation at lower transmission rates. The net effect of this is that prevalence plateaus at lower levels. The model was fit for σ, n and α (strength of cross-immunity) using the simplified transmission-prevalence curve. The best fits were produced by parameters which predict no cross-immunity ($\alpha = 0$, equivalent to the piecemeal scenario) with n and σ equivalent to that found for the piecemeal scenario (compare figures 2.16a with 2.14b); both predict only moderate levels of diversity.

To understand the consequences of how larger population diversity would affect the ability of the model to fit transmission-prevalence data we fixed $n = 300$ and refit the model (figure 2.16b). Without cross-immunity, this scenario would overestimate prevalence at high transmission intensity or require very short infections. Cross-immunity limits the amount of diversity that can be usefully utilised by the parasite, since hosts eventually reach a level where they experience low risks of infections despite only having been exposed to a subset of circulating strains. The net result of this is that the maximum prevalence is lower, even in regions with high transmission intensity (figure 2.16d). Cross-immunity thus sets a limiting threshold on the advantage that can be gained from diversity beyond which additional diversity does not significantly affect prevalence. This can be seen from the interaction between cross-immunity and diversity in figure 2.12a and b in which stronger cross-immunity makes prevalence insensitive to higher parasite diversity. Figure 2.16c shows the changes in goodness of fit for different combinations of n (diversity) and α (strength of cross-immunity). There is a fairly narrow range of feasible values involving low-moderate levels of cross-immunity and moderate-to-high degrees of diversity. These results suggests that high parasite diversity could be accommodated with the addition of relatively weak cross-immunity. We further found a trade-off between the
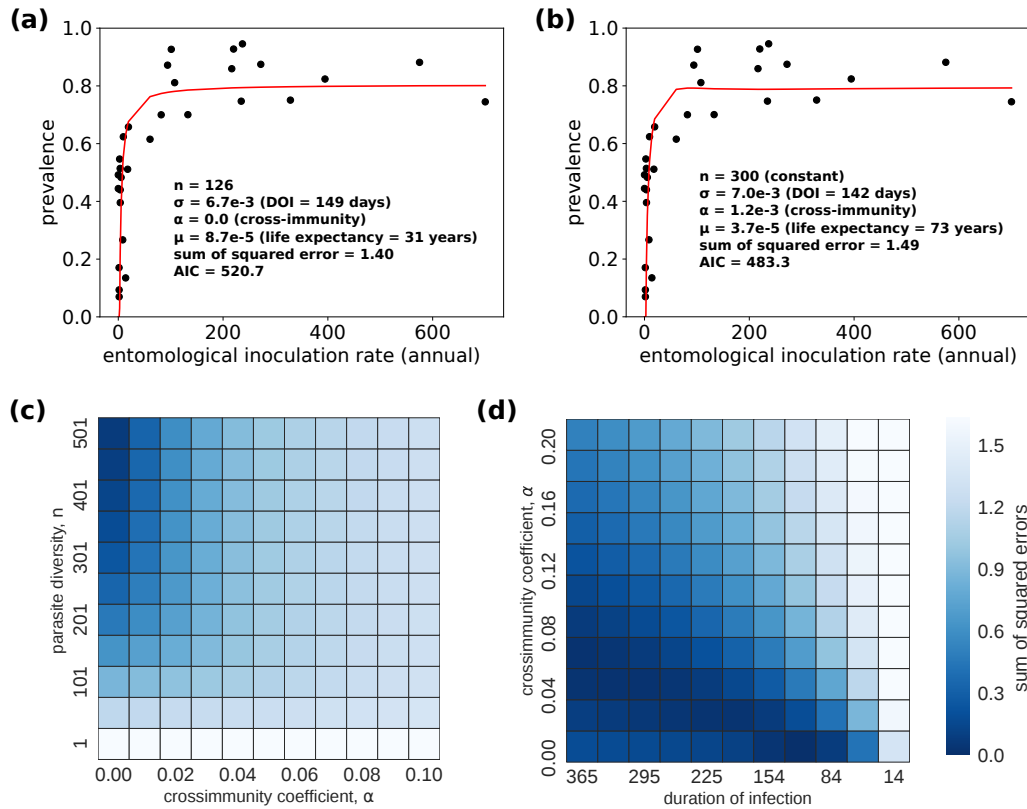
FIGURE 2.16: Fitting the cross-immunity scenario. **(a)** Simulated transmission-prevalence curves for the best fitting parameters. **(b)** Simulated transmission-prevalence curve for the best fitting parameters under the constrains that $n = 300$. Best fitting parameters, sum of squared error and Akaike information criteria (AIC) are shown as annotations. **(c)** The effect of interactions between strength of cross-immunity (α) and parasite diversity (n) on sum of squared error. Even moderate levels of cross-immunity limit prevalence strongly by inducing a transmission saturated (diversity limited) state. Severely limited prevalence leads to poor fits. **(d)** The interaction between strength of cross-immunity (α) and duration of infection shows that longer infections can compensate for stronger cross-immunity. Other parameters: $\frac{1}{\sigma} = 60$, host life expectancy is 60 years.

strength of cross-immunity and the duration of infection, in which greater duration of infection could be compensated for by increasing the strength of cross-immunity (figure 2.16d).

2.3.5 Comparison of modelling frameworks

AIC (Akaike information criterion) was calculated for each modelling framework, shown in table 2.2. In the case of the first three 'simple' models, AIC is lowest for the exposure independent model (301.1) and highest when temporary immunity was

used (371.6). The binary immunity model, arguably the most biologically realistic of the three, had an AIC of 327.7. The cumulative exposure model fits resulted in higher AIC for all scenarios, with the lowest AIC in the simplest piecemeal acquisition of immunity scenario (516.3 for test curve 1).

Given the relatively simple shape of the transmission-prevalence relationship and that Beier *et al.* demonstrated it could be approximated well with a simple logistic function (Beier, Killeen, and Githure, 1999), it is not surprising that the simplest framework lead to the lowest AIC. It is important, however, to consider this in the context of the study aims: to explore the implications of making different immune assumptions in models of transmission and prevalence, and to understand the biological processes which may be important in driving the transmission-prevalence relationship. While AIC suggests that the exposure independent framework is the best framework, the main assumption of this model, that exposure leads to no change in the immune status of hosts, is clearly at odds with the biological reality.

Of the models which incorporate exposure mediated changes in the immune status of hosts, the binary immunity framework has the lowest AIC, despite needing to estimate 5 parameters compared to the 2 to 4 parameters for the cumulative exposure scenarios. While it may be that the simpler binary immunity assumption is a good approximation of immunity to *P. falciparum*, it is also possible that this reflects the limitations of using the simplified the test curve to fit the cumulative exposure scenarios. Fitting used a simplified curves generate from four test points, while the sum of squared error and AIC were calculated based on the full dataset, which would be expected to adversely affect the fit to the full dataset. To illustrate the potential effect of the particular choice of test curve on the AIC, we repeated the fitting process with an alternative choice of test curve (test curve 2, see section 2.2.6). The best fits for test curve 2 for each scenario can be seen in figure 2.17 alongside the equivalent fits

TABLE 2.2: Akaike information criterion (AIC) for each fitted model.

Model framework	Test curve (TC) used	AIC
Exposure independent	n/a	301.1
Temporary immunity	n/a	371.6
Binary immunity	n/a	327.7
Cumulative exposure model:		
Piecemeal immunity	TC1	516.3
Exposure-dependent duration of infection	TC1	631.4
Cross-immunity	TC1	520.7
Cross-immunity ($n = 300$)	TC1	483.3
Piecemeal immunity	TC2	434.1
Exposure-dependent duration of infection	TC2	493.0
Cross-immunity	TC2	765.6
Cross-immunity ($n = 300$)	TC2	483.3

using test curve 1. While the resulting fits produced by each test curve appear relatively small, there can be quite dramatic differences in the AIC, with lower values generally reported for TC2. Interestingly, the relative AIC between the cumulative exposure scenarios differs, with the exposure mediated duration of infection and cross-immunity scenarios switching order between TC1 and TC2. It is therefore difficult to draw firm conclusions as to which framework / scenario is best. These results do, however, highlight the need for more data and a better understanding of the biological and immunological processes involved. An advantage of the cumulative exposure model is that permits investigation into the way successive exposure events can influence the transmission-prevalence relationship.
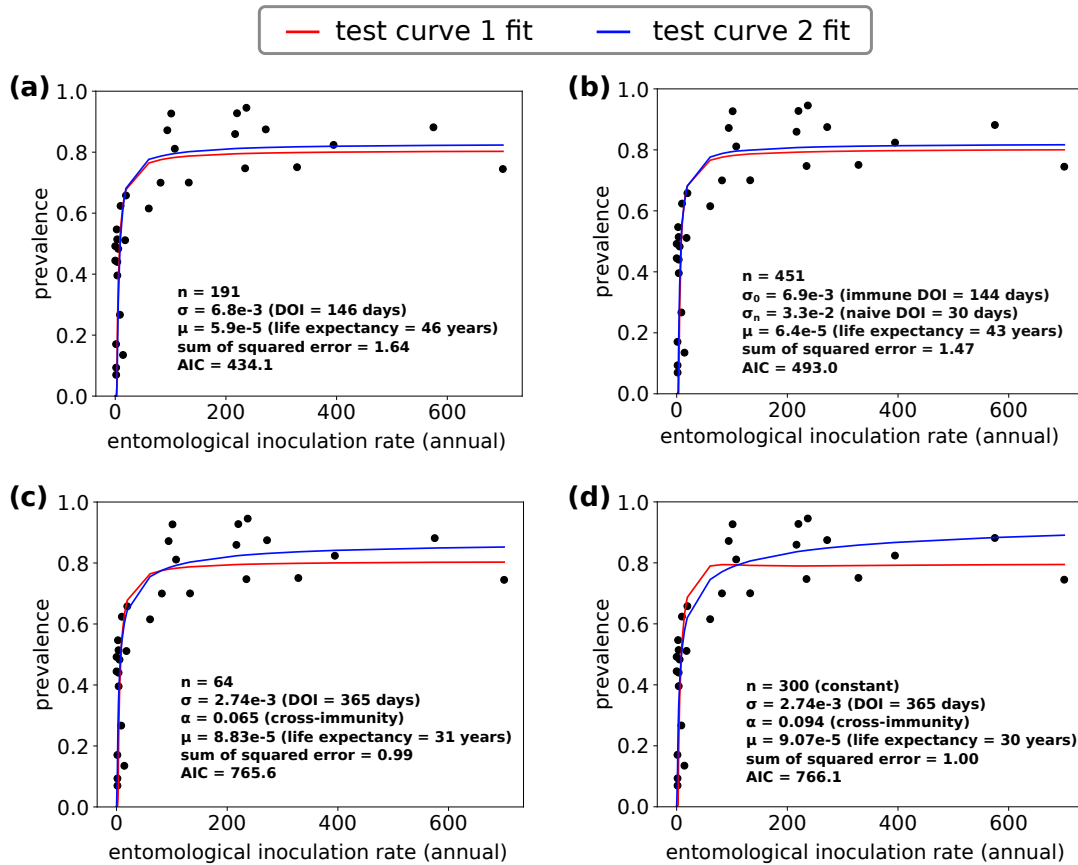
FIGURE 2.17: Comparison between fits using test curve 1 and test curve 2 for each of the cumulative exposure model scenarios: **(a)** variant-specific immunity, **(b)** exposure-dependent duration of infection, **(c)** cross-immunity, and **(d)** cross-immunity with parasite diversity fixed ($n = 300$). The best fit using test curve 1 is shown for comparison. Best fitting parameters, sum of squared error and Aikaike information criterion all refer to the test curve 2 fit. The Beier dataset are shown by black points (Beier, Killeen, and Githure, 1999).

2.4   Discussion

Host immunity to *P. falciparum* malaria is thought to be acquired gradually through repeated exposure to a diverse parasite population. While the epidemiological effect of immunity on disease severity is becoming clearer, there remains considerable uncertainty surrounding the effect of acquired immunity on transmission. Yet the relationship between transmission intensity and prevalence is strongly influenced by the assumptions we make about immunity. We used three simple deterministic models to explore the effect that common assumptions about immunity have on the transmission-prevalence relationship. In the first of these models we assumed there was no transmission-reducing effect of immunity and found that while this could easily generate the rapid increase in prevalence with transmission rate seen in the data, there was no mechanism to limit prevalence at higher transmission intensities. In the second model temporary immunity was shown to limit prevalence even at high transmission intensities. This is because a proportion of the host population was protected from infection at any give time, dependent on the mean duration of immunity. This framework reproduced both the initial low-moderate intensity and high intensity part of the relationship fairly well.

The third model assumed that recovery from infection is associated with a small chance of developing partial transmission-reducing immunity and reproduced both the initial rise and plateauing parts of the transmission-prevalence curve well. While this binary action of immunity is probably unrealistic, it represents a simplification of the gradual acquisition of immunity into two exposure levels: non-immune and semi-immune, whereby semi-immunity is acquired with some low probability. One consequence of this 'low resolution' exposure is that the model tended to predict infection length at the extreme ends of the parameter ranges. Very long infections

(350 days) were required in non-immune hosts to reproduce the initial steep rise in prevalence, but quite short infections (37 days) were predicted in semi-immune hosts to curtail transmission and limit prevalence at high transmission intensity.

In a fourth and more complex framework we explicitly modelled the effect of gradually acquired immunity in parasite populations with different degrees of diversity. In this framework, successive infection events explicitly lead to cumulative changes in host immunity, with parasite population diversity defining the number of exposure levels. This provided a flexible modelling framework in which the gradual acquisition of immunity could be modelled in more detail. We considered three immune mechanisms and their effect on the transmission-prevalence relationship: (1) The piecemeal acquisition of antibodies which recognise and prevent infection by serologically similar strains, (2) broadly acting cross-immunity, in which infection to one strain not only prevents reinfection by that strain but reduces susceptibility to other strains, (3) antigen-specific immunity which limits the number of novel antigens that can be expressed in subsequent infections, leading to a reduction in the duration of infection.

All three of our immune scenarios were able to fit the observed transmission-prevalence relationship well, making it difficult to make specific conclusions about the effect of immunity on transmissibility. However, we were able to characterise key interactions between parameters which lead to particular trade-offs and to identify regions of parameter space which lead to similarly good fits. One important trade-off which was apparent in all three immune scenarios was between diversity and infection length. Large parasite diversity could mitigate the need for extremely long infections, with equivalent prevalence generated from many short-lived infections or fewer longer-lasting infections. Our understanding would therefore benefit greatly from further work to characterise the extent of serologically distinct diversity

in natural parasite populations over a range of transmission settings.

While it is difficult to directly map the exposure levels in our model to the diversity measures used in genetic studies (e.g. sequence diversity in DBLα domains in PfEMP1 (Barry et al., 2007; Aguiar et al., 1992; Chen et al., 2011; Day et al., 2017)), our results highlight some of the qualitative effect of trade-offs between epidemiological parameters and diversity. In particular we identified two mechanisms by which prevalence can be limited: transmission saturation and diversity saturation. Transmission saturation was primarily responsible for the plateauing of prevalence in simulated transmission-prevalence curves, suggesting that diversity is a key factor in determining this relationship. We expect that the epidemiological response to changes in transmission intensity will respond differently depending on whether the local disease ecology is in a state of transmission or diversity saturation (or somewhere in-between). In particular, our results suggest that settings with high transmission intensity (i.e. in the plateau region of the transmission-prevalence curve) may be diversity limited, because this can explain the apparent insensitivity of prevalence to changes in transmission intensity. In these high transmission regions, variation in prevalence may be predominantly explained by factors relating to parasite diversity or population structure, rather than differences in the entomological inoculation rate. In low to moderate transmission settings, variation in parasite diversity is expected to have less influence on prevalence because, at least in our model, prevalence in these systems was primarily limited by transmission intensity. While this study focused on how immune mediated factors could limit prevalence, another possible source of plateauing before 100% prevalence at high transmission intensity could be host genetic factors, such as the resistance conferred by the sickle cell trait. Additionally limitations in the ability to detect infections with low parasite density could explain some of this plateauing.

Diversity in real parasite populations may well be significantly larger than we considered in this study (Barry et al., 2007). Nevertheless, we demonstrated that the way in which diversity and immunity interact have important consequences for understanding the processes that might underpin observed patterns in prevalence. Several studies have documented immense PfEMP1 diversity at global and local scales (Barry et al., 2007; Day et al., 2017; Chen et al., 2011; Aguiar et al., 1992), and it is understood that immune selection is responsible for the maintenance of such diversity. We made a number of simplifying assumptions, which included the assumption that all strains and/or antigens were equally abundant in the parasite population and that antigens could neither be lost from the population nor generated/imported. This meant that the model could offer no insights into the effect of immune selection on the evolution and maintenance of diversity, parasite population structure or the dynamics in the abundance of individual strains or antigens.

Given the potentially important influence that diversity has in shaping the transmission-prevalence relationship, it is pertinent to consider the processes by which this diversity is generated and maintained. Both inter- and intragenic recombination contribute to the generation of antigenically diverse parasites (Rich, Hudson, and Ayala, 1997; Rich, Ferreira, and Ayala, 2000; Conway et al., 1999; Taylor, Kyes, and Newbold, 2000; Freitas-Junior et al., 2000; Mu et al., 2005; Jiang et al., 2011). The degree to which each of these are important and how changes in the disease ecology may affect the maintenance of this diversity are not well understood. however the simple treatment of diversity as a static property could therefor mask important epidemiological processes. Identifying the factors which influences the degree of antigenic diversity that can be maintained in *P. falciparum* populations will aid our understanding of how parasite populations respond to natural or human-induced

changes in transmission intensity. In the next chapter we continue this line of investigation by explicitly allowing antigens and repertoires of antigens to evolve in response to immune selection.

Chapter 3

Diversity as a dynamic epidemiological property

3.1 Introduction

In considering the effect of antigenic diversity on *P. falciparum* transmission and epidemiology, it is necessary to consider the processes which generate and maintain this diversity. Diversity in PfEMP1 is generated through a combination of meiotic and mitotic recombination. Meiotic recombination, which takes place inside the mosquito, works to generate new *var* repertoires and may help to break up combinations of *var* genes which are maladaptive (e.g. if they contain cross-reactive epitopes to one another). Diversity generation through meiotic recombination may be amplified by frequent superinfection in humans, or by mosquitoes feeding from hosts infected by different parasite clones (Koella, Sörensen, and Anderson, 1998). Mitotic recombination involves the exchange of nucleotide sequences between genes and results in the generation of new *var* genes encoding distinct PfEMP1 variants. Intragenic recombination is typically associated with high rates of deleterious recombinants. In PfEMP1 this may be partially mitigated by the protein's modular structure, which facilitates the movement of intact binding domains or tandem groups of

domains, between *var* genes to preserve the domain level structure responsible for functional binding (Ward et al., 1999).

We previously developed a series of simple deterministic models to explore how the interaction of immunity and diversity influences the transmission-prevalence relationship. Models of *P. falciparum* transmission often treat diversity as a static property in which hosts develop some level of immunity after a certain threshold of exposure or with some small probability (Eckhoff, 2012; Dietz, Molineaux, and Thomas, 1974; Struchiner, Halloran, and Spielman, 1989). However, the processes which generate and maintain diversity are inextricably linked to the transmission process. Meiotic recombination requires transmission to a mosquito, while transmission to human hosts provide opportunities for novel recombinant antigens to be generated through mitotic recombination. We argue that diversity should more accurately be considered as a dynamic property, which arises from the underlying transmission setting. As such, diversity should also respond to temporal fluctuations in transmission intensity. However, the epidemiological consequences of treating diversity in this dynamic manner are largely unknown.

In this chapter we developed a stochastic individual-based model to simulate inter- and intragenic recombination, which allowed new antigen variants and antigen repertoires to be generated in an *ad hoc.* manner. Consequently, diversity emerged in our model as a dynamic property of the underlying transmission dynamics and was not hard-coded. The model considered host-parasite and vector-host interactions at the antigen level, which allowed us to investigate of the evolutionary forces that shape antigenic diversity.

3.2 Methods

In designing the dynamic diversity model, we outlined the following necessary properties:

- The ability to generate antigenic diversity in an *ad hoc.* manner:

    - Multiple infection can lead to meiotic recombination in mosquitoes, which can then transmit a recombinant parasite strain.

    - Mitotic recombination can occur in the human host and must be able to generate novel antigen variants, which can be transmitted to mosquitoes.

- Support for different immune assumptions with respect to how exposure affects the host's immune state (and therefore the immune selection on the parasite population).

    - The ability to map from exposure to immunity. This requires a distance metric for computing similarity between antigens.

- Output timeseries for human and mosquito prevalence, population level immunity, EIR and parasite diversity.

The dynamic diversity model consists of a population of individually modelled humans (hosts) and a population of individually modelled mosquitoes (vectors).

3.2.1 Mosquito definition

Mosquitoes are modelled individually and can be either uninfected or infected by a single strain. Mosquitoes die of natural causes and are immediately replaced with a new uninfected individual to maintain a constant population. The age distributed probability of death in mosquitoes is modelled using a logistic function (figure 3.1),

$$p_{death}(d) = \frac{0.25}{1 + e^{-0.5*(d-\mu_M)}} \qquad (3.1)$$

where mosquitoes have a mean lifespan of μ_M, and d is the mosquito's age in days.
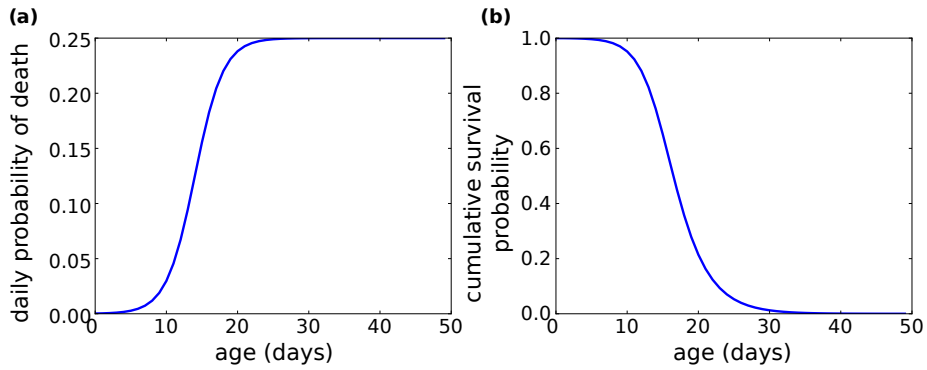
FIGURE 3.1: Mosquito demographics. (a) Age distributed daily probability of death. (b) Mosquito survival probability.

3.2.2 Host definition

Hosts are modelled individually and may have up to two active infections. The reason for constraining hosts to a maximum of two concurrent infections is to construct the simplest possible scenario in which meiotic recombination can take place (see description of recombination below). The immune status in hosts is tracked by an immune status vector, which describes the level of immunity to each antigen type, i. The immune status of a host is defined by a column vector of size A_{max}, referred to as $\mathbf{h}$,

$$\mathbf{h} = \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ \vdots \\ h_{a_{max}} \end{bmatrix} \qquad h_i \in [0,1] \qquad (3.2)$$

where h_i represents the degree of immune protection that the host has to antigen type i. $h_i = 1$ indicates complete immunity to antigen type i and $h_i = 0$ complete susceptibility. Upon death hosts are immediately replaced by a newborn naive individual to maintain a constant host population. We ignore any protection associated with maternal antibodies in infants. The daily probability of host death is defined using an exponential function (figure 3.2):

$$p_{death}(a) = e^{-\lambda y} - 1 \tag{3.3}$$

where $\lambda = -1.5e^{-6}$ and y is the host age in years. We assume a constant probability of death for all days over the course of a year.
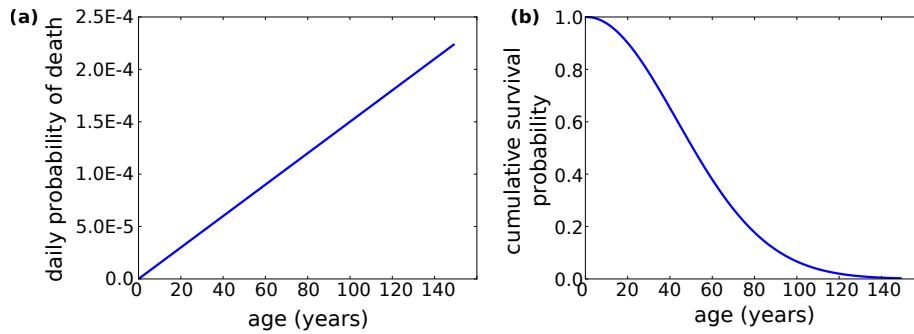
FIGURE 3.2: Host demographics. (a) Age distributed daily probability of death. (b) Host survival probability (annual).

3.2.3 Antigen representation

Antigens are encoded by a four byte binary number. This bit sequence determines both the genotype and antigen type. We refer to 'genotype' or 'gene' as the entire sequence, while 'antigen type' or 'antigen ID' refers to a variant which induces a distinct immunological response in the host. The n least significant (trailing) bits only affect the genotype, while the $32 - n$ most significant (leading) bits encode and fully define the antigenic type. There are a maximum of 2^{32-n} possible antigen types, each with 2^n possible genotype encodings. An example 1-byte gene is shown in

figure 3.3 for $n = 3$. Bits encoding the antigen type are shown as dark gray boxes

and bits which encode only the genotype are light gray.

Antigen ID = 10011_2 = 19_{10}
Genotype = 10011010_2 = 154_{10}

FIGURE 3.3: A simple 1-byte example representation of an antigen (in the model four bytes are used). All 8 bits are used to encode the antigen's genotype, while only 5 bits contribute to the antigen type (shown in darker gray). There are therefore multiple possible genotypes for each antigen type.

The antigen type (antigen ID), a, of a gene is calculated by performing a bit-wise

right shift by n and wrapping around the interval $[0, A_{max})$, where A_{max} controls the

total number of antigen types, i.e.

$$a = \lfloor \frac{g}{2^n} \rfloor \bmod A_{max} \tag{3.4}$$

Wrapping antigen space using the modulo operator can result in an unequal

number of ways to represent each antigen type, with some antigen types having one

extra representation than others unless $\lfloor \frac{(2^{32}-1)}{2^n} \rfloor \bmod A_{max} = 0$. However, this is un-

likely to affect our results in any significant way because there is large redundancy in

antigen representation for the considered ranges in n and A_{max}. For example, in the

extreme case, where $A_{max} = 50000$ and $n = 7$ some antigen types will have 85900

possible encodings while others have 85899, which this is not anticipated to affect

the output in any measurable way.

3.2.4   Parasite representation

Parasites are implemented as a set of R antigen genes. It is mathematically con-

venient to construct an antigen vector, **s**, which is a column vector of length A_{max}.

Each element $s_i \in \mathbf{s}$ indicates the presence (1) or absence (0) of each antigen type i in a parasite's antigen repertoire:

$$\mathbf{s} = \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ \vdots \\ s_{A_{max}} \end{bmatrix} \qquad s_i \in \{0,1\}, \quad \sum s_i = R \tag{3.5}$$

3.2.5 Transmission

Mosquitoes bite hosts at a constant rate, b, hence the number of feeding events for a mosquito follows a Poisson distribution. An example probability density function using a daily bite rate of 0.2 is given in Figure 3.4. Note that mosquitoes are assumed to be able to find any host at any time, and failure in feeding or host seeking behaviour is incorporated into the daily bite rate.
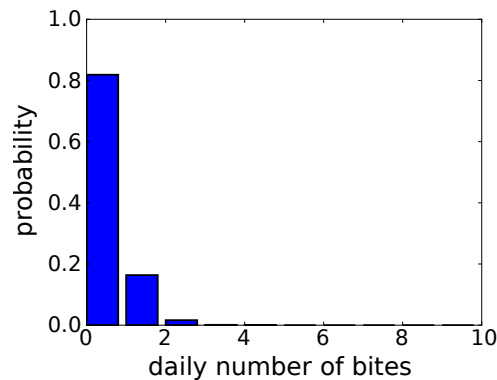

FIGURE 3.4: Probability density function for number of hosts bitten per mosquito, with daily bite rate $b = 0.2$.

When an infectious mosquito feeds, a copy of its infecting parasite is transmitted to the host if the host is not already at its maximum capacity for concurrent infections. Unless the host has full immunity to every antigen type in the parasite's

antigen vector a new infection is instigated. Latent periods are ignored in human hosts.

If an uninfected mosquito feeds on an infected host it becomes infected with a fixed probability, p_{trans}. Unlike hosts, mosquitoes can only be infected by a single strain. If the host has concurrent infections the mosquito becomes infected with a recombinant strain via sexual recombination. This represents a simplified scenario and assumes that the mosquito, if infected, is always infected by both strains. However, this allows meiotic recombination to occur without making complex assumptions about the within-mosquito dynamics of different strains. While latent periods are ignored in humans, the short life expectancy of mosquitoes mean that the incubation period in mosquitoes can have a much larger effect on transmission in mosquitoes. We therefore assumed that mosquitoes become infectious after a constant incubation period of 10 days. Mosquitoes are assumed to remain infectious until they die of natural causes.

3.2.6 Within host dynamics

Parasites are assumed to sequentially express their entire antigen repertoires during infections. The duration of infection is calculated by summing the contribution to infection made by each of these antigens, weighted by the host's immune status, i.e.

$$\omega(\mathbf{h}, \mathbf{s}) = \alpha(\mathbf{1} - \mathbf{h})^\mathsf{T}\mathbf{s} \qquad (3.6)$$

where α is the contribution to the duration of infection made by an antigen to which the host has no immunity, $\mathbf{s}$ is the parasite's antigen vector and $\mathbf{h}$ is the host's immune status vector. The duration of infection is therefore directly determined by the immune status of the host.

We make the simplifying assumption that the host's immune status changes immediately to reflect exposure to all antigens in the infecting strain's repertoire. Hence, the success of future transmission events reflect these changes even if a host has not cleared the ongoing infection yet. If a host has concurrent infections, each strain is assumed to act independently of one another with the exception of this indirect immune interference.

We consider two different assumptions as to how exposure transforms a host's immune status. Firstly, strict variant specific immunity is implemented such that hosts develop complete immunity specific to each antigen expressed by an infecting strain, i.e.

$$\mathbf{h}_{t_2} = \phi(\mathbf{h}_{t_1} + \mathbf{s}) \tag{3.7}$$

where $\mathbf{h}_{t_1}$ is the immune vector before infection by a strain with strain vector s, and $\mathbf{h}_{t_2}$ is the transformed immune vector, which reflects the changes resulting from exposure to s. ϕ is an element-wise function which ensures $0 \leq h_i \leq 1$ for all elements, i.e.

$$\phi(\mathbf{x}) = \begin{cases} 1 \ \text{if} \ x_i \geq 1 \\ x_i \ \text{if} \ 0 < x_i < 1 \qquad \text{for all } x_i \in \mathbf{x} \\ 0 \ \text{if} \ x_i \leq 0 \end{cases}$$

An example of an immune status transformed by variant-specific immunity is given in figure 3.5a.

We also consider cross-immunity, which acts such that exposure to antigen a confers protection to other antigens depending on the distance in antigen space from a. This is implemented by an additive transformation of the immune status vector
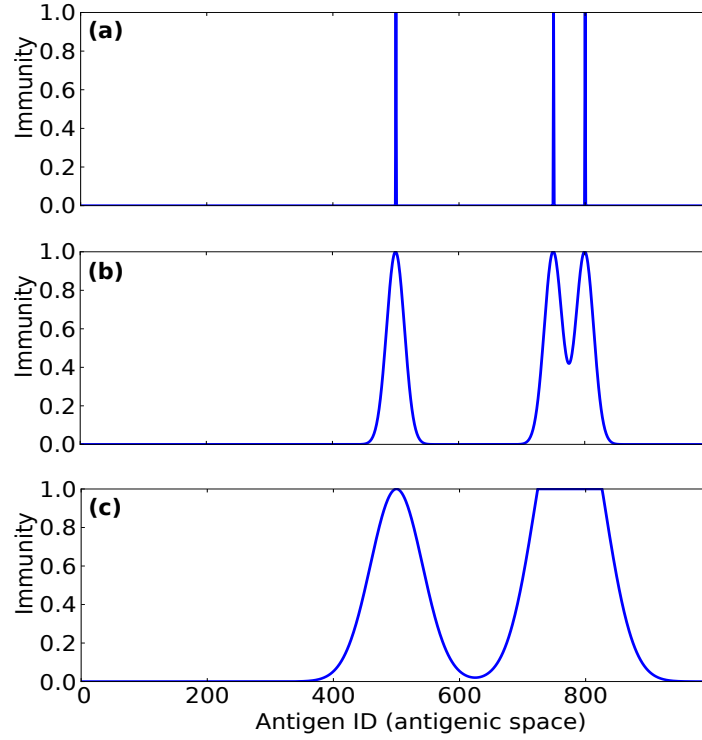
FIGURE 3.5: Host immune status after exposure to three antigens (antigen IDs: 500, 750, 800) under strict variant specific immunity (a), weak cross-immunity (b), and strong cross-immunity (c). Antigen space has been restricted to 1000 for illustrative purposes.

using a Gaussian distribution with a mean corresponding to the antigen type, and the variance determining the strength of cross-immunity. Figure 3.5b and c show an example of how exposure to three antigens would transform the immune status of a host under moderate and strong cross-immunity, respectively. The transformation of a host's immune status with exposure to antigen a and cross-immunity is calculated as

$$\Delta h_{i,a} = \beta \frac{1}{\sqrt{2\gamma n}} e^{-\frac{(i-a)^2}{2\gamma}} \qquad \text{for } i \in \{1, 2, \ldots, n\} \tag{3.8}$$

where γ controls the strength of cross-immunity, n is the size of antigen space, and a is the antigen to which the host is being exposed to. β is a normalising coefficient, defined as

$$\beta = \sqrt{2\gamma n} \tag{3.9}$$

Inherent in this scheme is the assumption is that similarly numbered antigen types are structurally similar with respect to the epitopes presented to the immune system. In other words, we assume antigen similarity is equatable to the difference between antigen type numbers.

Finally, the immune status of the host after infection can be calculated by summing the changes in the immune status for each antigen expressed:

$$\mathbf{h}_{t_2} = \phi(\mathbf{h}_{t_1} + \sum_{a=1}^{n} s_a \Delta \mathbf{h_a}) \tag{3.10}$$

where $\mathbf{h}_{t_1}$ and $\mathbf{h}_{t_2}$ is the immune status vector before and after infection respectively. n is the size of antigenic space, s_a is the antigen vector at antigen ID a and $\Delta \mathbf{h_a}$ is the change in immune status due to antigen a. ϕ is the previously described piece-wise function which ensures each element of $\mathbf{h}$ remains in the interval $[0, 1]$

3.2.7 Recombination

Intergenic (meiotic) recombination involves the exchange of whole antigens between repertoires. In our model this occurs when a mosquito feeds from a host with concurrent active infections. The new repertoire is generated by taking genes from each locus from the first parent strain with probability $1 - p_s$ and genes from the second parent with probability p_s, resulting in a recombinant repertoire made from a combination of whole genes from the two parent repertoires. Intergenic recombination between genes at one locus is thus assumed to occur independently of the other loci.

Intragenic (mitotic) recombination is able to generate new antigens types. It is expected that newly emerged parasites resulting from this processes would make up

only a small fraction of parasites in the host and have little influence on the ongoing infection. Additionally, we expect the probability that mosquitoes become infected by such a minority to be much lower than for that of the more numerous original clone. For simplicity, we do not model these recombination events at the within-host level. Instead, we allow for a small probability, p_c per gene, that a recombinant gene is copied to the transmitted strain in the mosquito. p_c therefore incorporates the probability that a gametocyte of a recombinant parasite is taken up during the blood meal as well as the rate of intragenic (mitotic) recombination. This has the computational advantage in that the recombinant strain only needs to be calculated upon successful transmission and not during ongoing infections.

The mitotic recombination operator is implemented by adding a number, r, to the original antigen representation (i.e. to the full genotype). r is proportional to the difference between the original and donor antigen numbers, i.e.

$$r = \kappa\theta(a_i - a_j) \tag{3.11}$$

where κ is a random number drawn from a uniform distribution with interval $[-1, 1]$, θ scales the magnitude of change, and a_i and a_j are the antigen type for original and donor antigens, respectively. Hence, the new recombinant gene is represented by g_i^*, given as

$$g_i^* = g_i + r \tag{3.12}$$

where g_i is the full genotype representation of the parent antigen.

This scheme avoids having to simulate the low-level biological mechanisms, such as the insertion, deletion or exchange of domains or nucleotide sequences between genes, leading to a computationally efficient model which emulates the main

features of recombination: (1) recombination between antigenically similar donor genes is likely to produce recombinant genes which are similar to the donor genes, (2) recombination between antigenically dissimilar genes is more likely to result in a recombinant antigen which differ from the original genes greatly, and (3) silent recombination events can alter the genotype but not the antigenic type, accumulate over time and can eventually lead to changes in antigen type.

An overview of the two steps in the transmission process and the generation of diversity is shown in Figure 3.6.
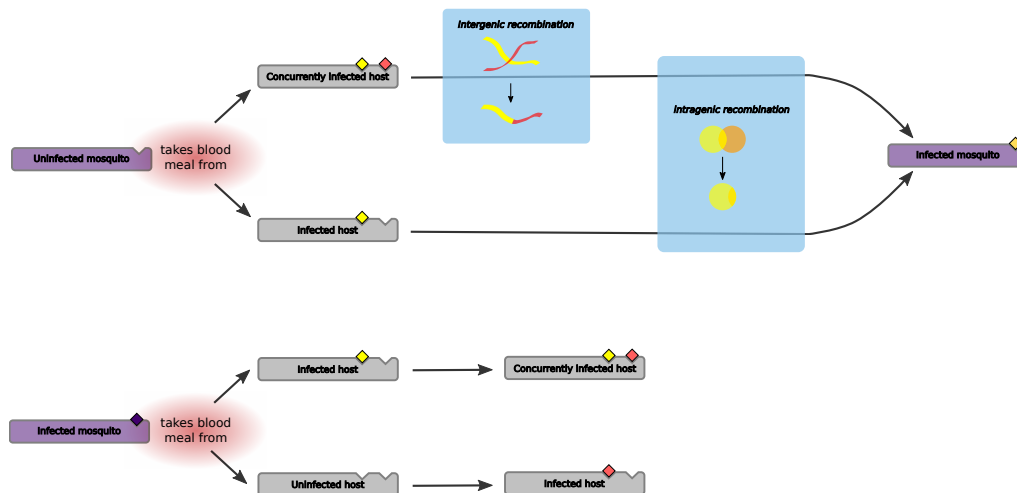
FIGURE 3.6: The possible transmission pathways and modes of re-combination used in the model. Uninfected mosquitoes can become infected by biting a host infected with one strain, or by biting a concurrently infected host. Mitotic recombination may occur in either case, but only when biting a concurrently infected host can the mosquito become infected with a recombinant strain produced by meiosis. Hosts can be infected with up to two strains concurrently, and infected mosquitoes can infect hosts which have zero or one active infections.

3.2.8   Initialisation

The model is initialised with the populations at demographic equilibrium. A small number of mosquitoes are infected, while all hosts are assumed naive. As a result, a burn-in period is required to allow the system to reach an equilibrium with respect to infections and immunity. A fixed number of mosquitoes, I_{init}, are infected by the following procedure. First, a pool of antigens of size A_{init} are generated from a uniform distribution over the whole genotype space. Next S_{init} strains are created by randomly sampling from the antigen pool. Finally, mosquitoes are infected with these initial strains, such that an approximately equal number of mosquitoes are infected by each initial strain until a total of I_{init} mosquitoes are infected.

3.2.9   Model output

Diversity

We use two metrics of parasite diversity: circulating antigen diversity and Shannon diversity index (Shannon entropy). Circulating antigen diversity is defined as the number of antigen variants in circulation in the parasite population divided by the total size of antigen space (A_{max}). It therefore takes a value in the interval $[0, 1]$ and indicates the proportion of antigenic space being utilised by the parasite population. The Shannon entropy is calculated as

$$H_s = \sum_{a=1}^{A_{max}} -\mathbf{p}_a \ln(\mathbf{p}_a) \tag{3.13}$$

where A_{max} the total number of possible antigens and p_a is the proportion of total antigens that antigen type a comprises. Shannon entropy takes into account the proportions of the different antigens in the population as well as the number of distinct antigen variants in circulation. Large Shannon entropy indicates that there are

either a large number of antigen types with uneven frequencies, or a smaller number of evenly distributed antigen types. Calculating circulating antigen diversity in addition to Shannon entropy helps to distinguish between these two cases. For example, if circulating diversity decreases but Shannon entropy stays approximately constant, it indicates that while the number of antigens in the parasite population are decreasing, the abundance of those that remain is becoming more equal.

Immunity

Two measures of immunity are used: absolute immunity and relative immunity. Absolute immunity is the average host immunity, normalised by A_{max}:

$$I_{absolute} = \frac{1}{H} \sum_{i=1}^{H} (\frac{1}{A_{max}} \sum_{a=1}^{A_{max}} h_{i,a})$$

(3.14)

where $h_{i,a}$ is the immunity to antigen type a that host i has acquired, H is the total number of hosts and A_{max} is the size of antigenic space. Absolute immunity is constrained to the interval $[0, 1]$ and indicates the extent to which hosts have acquired immunity to all possible antigens. Relative immunity is calculated by normalising by the number of circulating antigen variants (circulating antigen diversity) instead of A_{max}. This provides a quantitative indication of how well adapted the parasite population is compared to the level of immunity in the host population.

Host susceptibility

When relative immunity is greater than one, it indicates that the average host has immunity to a greater number of antigens than are currently in circulation and is suggestive of a poorly adapted parasite population. However, other possibilities exist. Hosts may be immune to a large set of extinct antigen types but relatively

susceptible to contemporary antigens. We therefore use another metric, host susceptibility, to distinguish between these two cases. Host susceptibility is a measure of how susceptible the host population is to the currently circulating antigenic diversity, or conversely how well adapted the parasite population is to the current immune state of the host population. Host susceptibility is computationally slow to calculate, however, and is only calculated when it is important to establish the degree that immunity protects against contemporary circulating antigens. It is determined by first calculating the mean host susceptibility vector, given by

$$\mathbf{S}_a = \frac{\sum_{i=1}^{H}(1 - \mathbf{h}_{i,a})}{H}, \tag{3.15}$$

where $\mathbf{h}_{i,a}$ is the immune status corresponding to antigen type a in host i, and H is the total number of hosts. Next, the mean host susceptibility vector is multiplied by the probability of encountering each antigen type (estimated by the proportion of each antigen type in the parasite population, $\mathbf{P}$). This provides a measure of risk that each antigen type posed to the average host. Finally, the sum of these risks is calculated to give an overall indication of the susceptibility of hosts to the current parasite population, i.e.

$$\Omega = \mathbf{P}\mathbf{S}^\mathsf{T}. \tag{3.16}$$

3.2.10   Summary of parameters

Table 3.1 summarises the main parameters, their baseline values and the ranges considered. For many parameters the large range reflects uncertainty in the literature. A sensitivity analysis can be found in Appendix A).

In this modelling framework, *P. falciparum* epidemiology can be viewed as a collection of interacting processes. These are summarised, along with the parameters
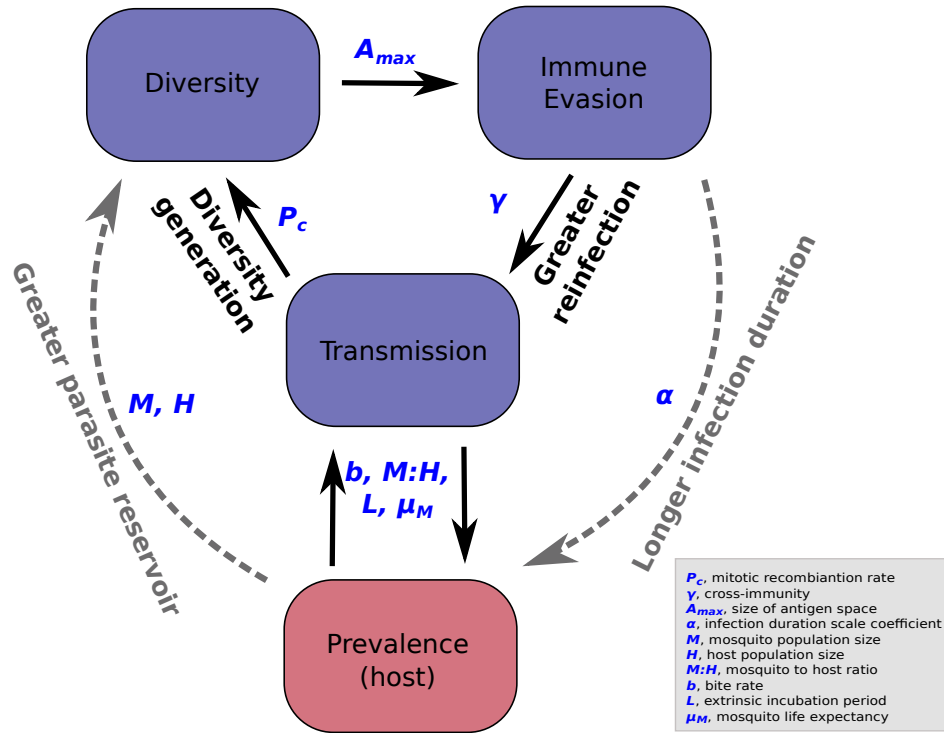
which directly influence them, in figure 3.7.

FIGURE 3.7: Interaction between the main processes in the model create a system of feedback loops. The main sources of feedback which this study focuses on are shown by the black arrows. Dashed gray arrows show other potential sources of feedback.. The model parameters which directly affect the strength of each interaction are shown in blue.

TABLE 3.1: Summary of the main parameters, their baseline values and the ranges considered.

parameter	description	baseline value [range]	reference
H	human population size	10000 [4000, 25000]	(Ferguson and Read, 2004; Charlwood et al., 2003)
M	mosquito population size	10000 [4000, 25000]	
b	mosquito biting rate	0.12 [0.05, 1.0]	
p_s	intergenic (meiotic) recombination rate	0.01 [0, 0.1]	(Barry et al., 2007)
p_c	intragenic (mitotic) recombination rate	0.002 [0, 0.02]	
A_{max}	size of antigen space	50000 [60, 100000]	
S_{init}	number of initial strains	40 [1, 1000]	
A_{init}	number of initial antigens	2400 [60, 60000]	
I_{init}	number of initially infected mosquitoes	250 [250, 2000]	
R	repertoire size	60	(Rask et al., 2010)
n	number of genotype only encoding bits	7	
L	extrinsic incubation period (days)	10	(Dietz, Molineaux, and Thomas, 1974)
μ_M	mosquito life expectancy (days)	32	(Bellan, 2010)
p_{trans}	infectivity to mosquitoes	0.5	
α	base antigen infection contribution (days)	3 [0.4, 90]	(Felger et al., 2012; Bruce et al., 2000; Smith et al., 1999a; Bousema et al., 2011; Kitua et al., 1996; Bekessy, Molineaux, and Storey, 1976)
θ	recombination scaling factor	1	
γ	strength of cross-immunity	0 (0, 10)	

3.3 Results

3.3.1 Static diversity

First, we consider the basic behaviour of the model without diversity generation (no recombination).

Basic model behaviour

Strains were initialised by randomly selecting antigens without replacement to ensure that there is no overlap between repertoires. This allows for the effect of diversity to be assessed without immune interactions between repertoires. In other words, we assume a strongly strain structured population with respect to the parasite's antigen repertoires. Example timeseries for the main outputs are shown in figure 3.8. There is an initial peak in prevalence and EIR as parasites spread quickly through a naive host population, followed by rising immunity, which slows transmission until an approximate equilibrium is reached. This equilibrium is reached when the immunity gained through exposure is balanced by immunity lost through host death. We describe this state as an approximate equilibrium because, while it resembles an equilibrium state, diversity is slowly lost due to stochastic extinction, which in turn leads to a slow decrease in prevalence. This can be seen most clearly in the relative diversity timeseries in figure 3.8c, which decreases monotonically over time.

Without recombination no new antigens are generated. The ability to retain initial diversity therefore depends simply on the size of the parasite population. That is, the rate of stochastic extinction for a given antigen type depends on its abundance and the size of the parasite population, with lower abundance putting antigen variants at higher risk of extinction. Larger populations can therefore maintain greater
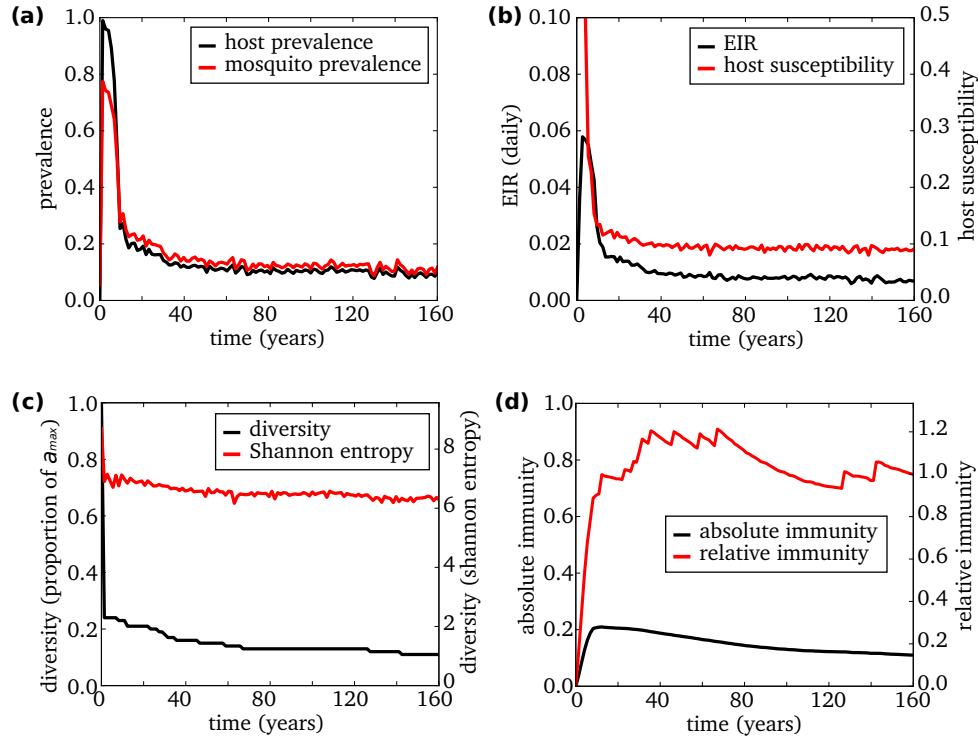
FIGURE 3.8: Example time series output from a single simulation. (a) Host (black) and mosquito (red) prevalence. (b) Daily entomological inoculation rate (EIR) (black) and host susceptibility (red). (c) Number of circulating antigens (black) and Shannon entropy of antigen composition of the parasite population (red). (d) Absolute (black) and relative immunity (red), see section 3.2.9 for an explanation of the difference between these measures. Parameters are as follows: $H = 16000$, $M = 16000$, $b = 0.12$, $A_{init} = 6000$, $S_{init} = 100$.

diversity at a given abundance. Similarly, larger host and mosquito populations can support larger parasite populations for a given level of prevalence. The effect of host and mosquito population size on the number of antigens retained after 82 simulated years is shown in figure 3.9.
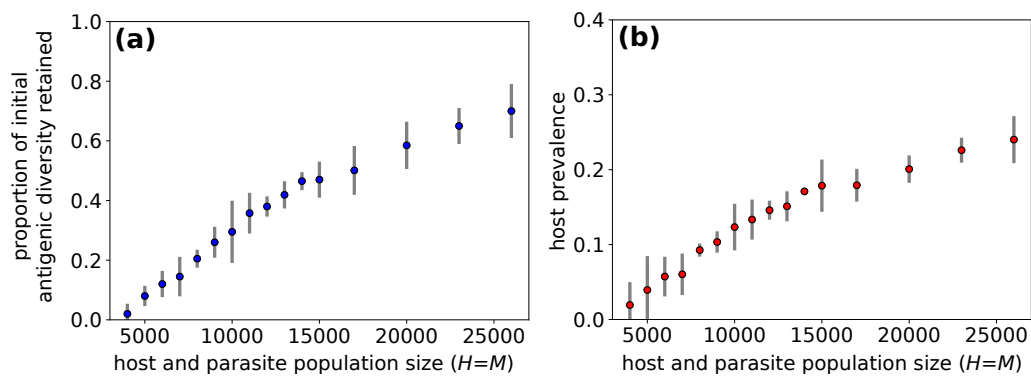
FIGURE 3.9: **(a)** Large host and mosquito population sizes can support large parasite populations, which reduce effect of stochastic extinction by maintaining larger proportion of initial antigenic diversity ($S_{init} = 50, A_{init} = 3000$). **(b)** High antigenic diversity in turn facilitates immune evasion and leads to higher proportion of infected hosts. Host and mosquito population size was increases proportionally to maintain a constant $M : H = 1$ ratio. The x axis therefore refers to both the mosquito and host population size. Each point represents the mean calculated from 10 repeat simulations. Error bars indicate the 95% confidence interval. Simulations were ran for 82 years. Other parameters which differ from the baseline values (see table 3.1): $b = 0.12$.

Diversity and transmission saturation

As described previously in Chapter 2, prevalence can be limited in two ways: by transmission saturation or diversity saturation. Both of these states are observed in our model. When in a diversity saturated state, the parasite population is well adapted and can evade host immunity with relative ease; prevalence, on the other hand, is limited by a bottleneck in transmission. Increasing antigenic diversity does not increase prevalence further, because it does not significantly increase the parasite population's ability to evade immunity. This insensitivity to diversity is shown by the shaded region in figure 3.10a. Under these conditions we also find that surplus diversity tends to be lost over time, because immune selection is not sufficiently strong to induce the negative frequency selection required to maintain it.

Transmission saturation is characterised by high relative immunity and a poorly adapted parasite population. Greater transmissibility does not lead to appreciable increases in prevalence because there are few susceptible hosts for parasites to infect (shaded region in figure 3.10b); however, increases in diversity increase the ability of the parasite population to evade immunity and cause prevalence to plateau at higher values.
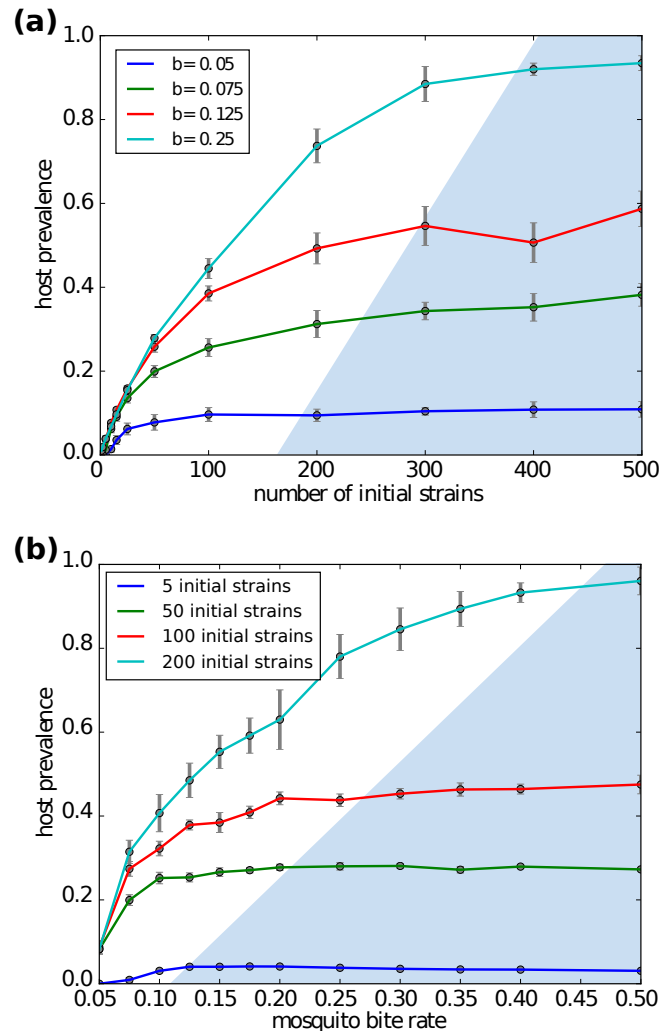
FIGURE 3.10: Prevalence can be limited by diversity or transmissibility. To demonstrate this the number of strains used to initialise the model (S_{init}) and the mosquito biting rate (*b*) were varied. **(a)** In a diversity saturated setting higher diversity offer little additional advantage for the parasite population, and fail to induce large increases in prevalence (shaded region). Parameters: $H = 10,000$, $M = 10,000$, $A_{init} = RS_{init}$. **(b)** In a transmission saturated setting increases in the transmissibility do not lead to higher prevalence because there is a relatively low proportion of susceptible hosts. Diversity is the limiting factor and prevalence plateaus at higher values for higher prevalence (shaded region). Means and 95% confidence intervals are calculated from 10 repeats. Parameters: $H = 10,000$, $M = 10,000$, $A_{init} = RS_{init}$.

Diversity and the transmission-prevalence relationship

Unlike the cumulative exposure model in the previous chapter, transmission intensity in the dynamic diversity model emerges out of the model from the underlying interactions between hosts, mosquitoes and parasites. Hence, transmission intensity cannot be explicitly set. We therefore explored the effect of transmission intensity on prevalence by varying mosquito biting rate, b. We choose to vary b because it is a computationally convenient and easily implemented way to modulate transmission rate. Figure 3.11a shows a strong linear relationship between b and EIR for the parameter range considered. When EIR is plotted against prevalence, the characteristic transmission-prevalence relationship is observed, with a steep initial rise in prevalence, which plateaus at moderate to high transmission intensity (figure 3.11b). As expected, this relationship is highly sensitive to diversity. Comparing this individual-based model (b) to the cumulative exposure model (c) we find that the absolute values differ (at least in part due to differences in infection length) but that the qualitative effect that diversity has on this relationship is the same. Higher antigenic diversity delays the onset of transmission saturation, causing prevalence to plateau at higher levels.

FIGURE 3.11: **(a)** Entomological inoculation rate (EIR) emerges from the interactions between hosts, mosquitoes and parasites, but exhibits a strong linear relationship with transmission intensity (here modulated by changing b). **(b)** The degree of antigenic diversity (initial diversity) strongly affects the transmission-prevalence relationship. Higher diversity delays the onset of transmission saturation, allowing prevalence to plateau to higher levels. For comparison **(c)** 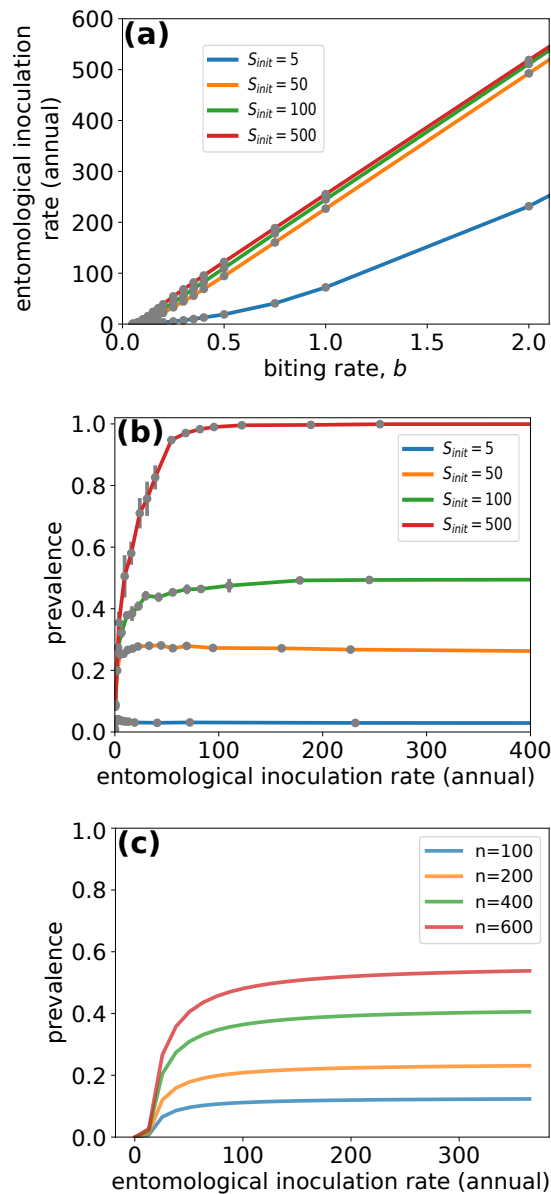shows the effect of diversity on the transmission-prevalence relationship in the deterministic cumulative exposure model from Chapter 2. Error bars in **(a)** and **(b)** show the 95% confidence interval calculated from 10 repeats.

3.3.2 Dynamic diversity

When recombination is enabled, diversity emerges as a property of the underlying transmission dynamics. A stable equilibrium is reached when the rate of diversity generation from recombination equals the rate of diversity loss through stochastic extinction. This is shown in figure 3.12, which shows time series of prevalence, diversity, immunity and EIR for a single run with recombination. When recombination is turned on (indicated by the vertical dotted line), the parasite population undergoes a phase of rapid diversification in response to immune selection (figure 3.12c). New regions of antigenic space are explored and negative frequency selection favours novel antigen variants. As diversity increases, new antigen variants are able to evade pre-existing immunity, causing an increase in EIR and prevalence (figures 3.12b and d).

A new equilibrium is reached as diversity, transmission, prevalence and immunity plateau. In figure 3.12, the level of diversity is large and immunity takes a long time to equilibrate, requiring the host population to be completely renewed so that all hosts have grown up under the new equilibrium conditions. New antigen variants continue to be generated and are lost at an equal rate to maintain a dynamic equilibrium.
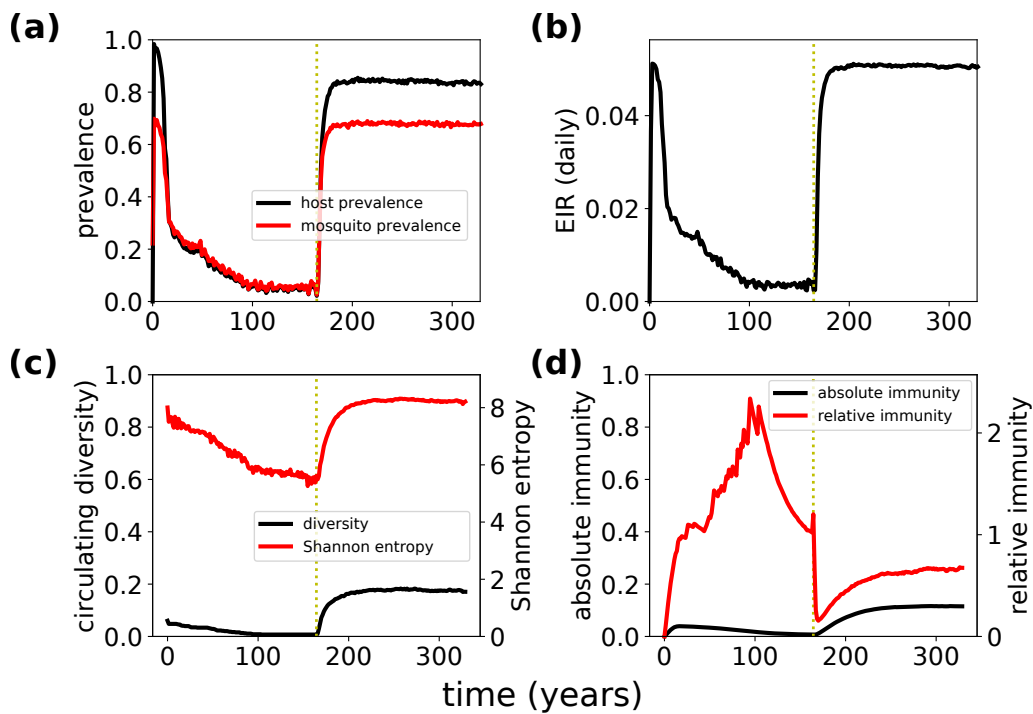
FIGURE 3.12: Model output from a typical run with recombination occurring after a burn-in period. The end of the burn in period is indicated by the dotted vertical line. Timeseries shown for **(a)** host and mosquito prevalence, **(b)** transmission intensity (entomological inoculation rate, EIR), **(c)** antigenic diversity, **(d)** relative and absolute immunity (see section 3.2.9 for how relative and absolute immunity are calculated). Parameters as follows: $H = M = 10000$, $S_{init} = A_{max} = 50000$, $A_{init} = 50$, $b = 0.12$, $P_c = 0.002$, $P_s = 0.01$, $\gamma = 0$.

Determinants of diversity

New antigen variants are generated in the model through intragenic recombination, controlled by the per-gene probability of intragenic recombination P_c. Larger P_c increases both the rate of diversity generation and the ability of the system to maintain diversity (the equilibrium level of diversity), illustrated in figure 3.13 for different values of P_c.
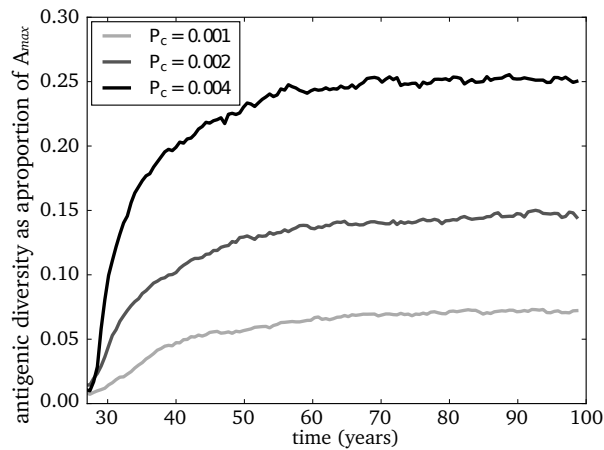
FIGURE 3.13: The per-gene probability of intragenic recombination (P_c) strongly influences equilibrium antigenic diversity by driving the generation of new antigen variants. Three time series from individual simulations show the number of antigen variants in circulation for different values of P_c. Both the rate of diversification and equilibrium point for diversity increase at higher P_c.

Diversity is lost through stochastic extinction which, as discussed previously, is strongly affected by population size. The risk of stochastic extinction for a given antigen variant depends on the abundance of that antigen in the parasite population. Larger parasite populations are able to accommodate greater abundance with the same degree of diversity, or greater diversity with no change in the abundance of each antigen variant.

Diversity generation is also driven by transmission as each transmission event provides the opportunity for recombination. The parameters that affect transmission intensity therefore also govern both the generation and the equilibrium level of

diversity. This includes host and mosquito population size, transmissibility, and the strength/mode of immunity. Interestingly, since diversity facilitates immune evasion, diversity is itself an indirect factor governing the generation of diversity. This introduces a key concept - that transmission and diversity are inextricably interlinked. Each time transmission is successful there is a chance that new antigens will be generated through intragenic recombination. This increases the diversity of the parasite population and hence the ability of the parasite population to evade host immunity, which in turn increases transmission and forms the basis of a positive feedback loop between diversity and transmission.

3.3.3 The transmission-prevalence relationship

Host immunity causes negative frequency selection on antigen variants in the parasite population. This can lead to novel antigen variants to emerge and the parasite population to adapt to host immunity. This is the fundamental difference between treating diversity as a dynamic versus static property. When diversity responds dynamically to changes in transmission and host immunity, the parasite population can overcome transmission saturation, which would ordinarily cause prevalence to plateau under static diversity (compare figures 3.14a with static diversity to 3.14b with dynamic diversity). When the system is transmission saturated, prevalence is limited by a lack of susceptible hosts, leading to strong immune selection for the parasite population to diversify. Assuming that there are novel regions of antigen space to be explored, the parasite population diversifies to escape host immunity. This results in prevalence increasing asymptotically to close to 100% as transmission intensity increases (figure 3.14b). This results in a similar transmission-prevalence

curve produced by the SIS framework in the previous chapter, which was also unable to limit prevalence at high transmission intensity. Lower intragenic recombination rates can reduce the equilibrium diversity for a given transmission intensity, therefore slowing the rate of increase in prevalence with transmission, but this does not prevent the monotonic increases toward full prevalence as transmission intensifies. Furthermore, slowing the rate of prevalence growth with transmission intensity leads to a poor approximation of the initial steep region of the transmission-prevalence curve. This is shown in figure 3.14b, in which prevalence approaches 100% even for very low mitotic recombination rate (e.g. $P_c = 5e-5$).

As already discussed, equilibrium diversity is sensitive to factors which contribute to the rate of diversity loss and generation. To test whether these could limit prevalence at high transmission intensity we simulated transmission-prevalence curves using different strengths of cross-immunity, population sizes and antigen space (shown in figure 3.14c, d and e respectively). We find that population size behaves similarly to recombination rate, with smaller populations slowing the rate at which prevalence increases with transmission intensity (figure 3.14c). This makes intuitive sense because population size, when controlled to keep the ratio of hosts to mosquitoes ($M : H$) constant, scales the absolute number of transmission events and therefore the rate of diversity generation in much the same way as P_c does. Smaller parasite populations can maintain less antigenic diversity at equilibrium, just as parasites with lower recombination rates do. Hence, hosts can gain immunity to a large proportion of circulating diversity with fewer infections, which reduces prevalence. Both recombination rate and population size are important determinants of equilibrium antigenic diversity, but both only slow the rate of growth in prevalence without causing it to plateau at high transmission intensity.
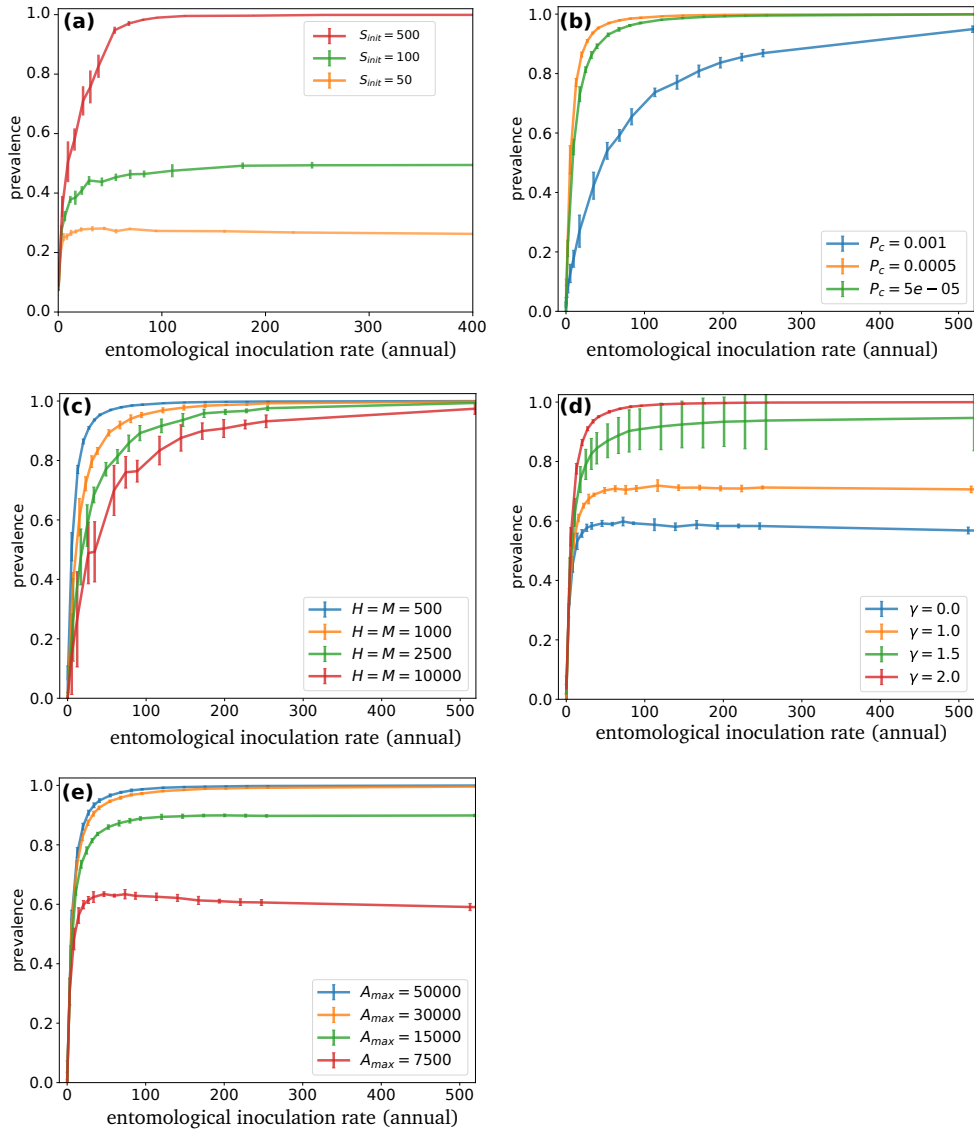
FIGURE 3.14: Simulated transmission-prevalence relationships. **(a)** Under static diversity different levels of initial diversity determine the transmission intensity at which transmission saturation occurs, and therefore also determines maximum prevalence. **(b)** When diversity is dynamically generated in response to immune selection the parasite population adapt to overcome transmission saturation. The rate of diversity generation, here manipulated by controlling the per-gene probability of mitotic recombination, P_c, affects the speed with which prevalence increases but does not determine a maximum level of prevalence. **(c)** Population size effectively scales the absolute number of transmission events and therefore also the rate of diversity generation, acting in the same way as changes to P_c. **(d)** Cross-immunity limits the ability of parasites to generate and utilise novel antigenic diversity. This prevents unconstrained diversification and immune evasion, leading to the re-establishment of transmission saturation at high transmission intensity. Stronger cross-immunity (higher λ) causes prevalence to plateau at lower values. **(e)** The size of antigenic space (A_{max}) acts through a similar mechanism as cross-immunity by limiting the ability of the parasite population to diversify and preventing diversity saturation at high transmission intensities. This brings about transmission saturation and causes prevalence to plateau at lower values for smaller antigen spaces. All sub-figures show the mean of 10 repeats, with error bars indicating the 95% confidence interval.

The strength of cross-immunity and the size of antigenic space affect prevalence in a qualitatively differently way. Increasing cross-immunity (γ) or reducing the size of antigenic space (A_{max}) limit the number of antigenically distinct parasite strains possible. This constrains the ability of the parasite population to evade immunity and can result in transmission saturation at high transmission intensity. Higher γ and/or lower A_{max} induces transmission saturation at lower transmission intensities, causing prevalence to plateau at a lower level. This suggests that limits on the ability of the parasite population to generate diversity in key antigens, for example due to structural constraints that arise from the necessity to retain functional binding, could limit parasite prevalence at higher transmission intensities and in larger parasite populations. Additionally, the existence of some degree of cross-reactive immunity, which can recognise a broad range of antigen variants, e.g. due to shared epitopes, could equally limit prevalence in high transmission settings.

3.3.4 Changing transmission intensity

To better understand the effect of dynamic diversity in conditions of variable transmission intensity, either due to natural fluctuations or human induced interventions, we varied transmission intensity by changing mosquito biting rate (b) over a three year period. Simulations were allowed to reach an equilibrium state before perturbation, and the timeseries outputs for diversity and prevalence were averaged over 30 repeat simulations (shown in figure 3.15). We found that the feedback between diversity and transmission caused a temporal lag in the response to perturbations in transmission intensity. This can be seen most clearly in figure 3.15a and b, which shows diversity continuing to increase (3.15a) or decrease (3.15b) many years after the period of perturbation. Since diversity is a key factor in driving prevalence, we

predicted a similar effect in prevalence. The temporal lag in prevalence was gener-

ally less pronounced but strongly influenced by the rate of intragenic recombination,

such that lower recombination rates increase the time taken to reach an equilibrium

after the perturbation period (figure 3.15c). Interestingly, relatively little lag was

seen in prevalence when transmission intensity was decreased (figure 3.15d). When

the rate of intragenic recombination is low, there is a slight rise in prevalence shortly

after the perturbation period. We hypothesised that this might be caused by the per-

turbation disrupting strain structure and the subsequent re-establishment of strain

structure. This is examined in more detail in Chapter 5, which focuses on the inter-

action between parasite population structure and immune selection.

New diversity is generated in response to increases in transmission intensity,

which in turn drive non-linear changes in prevalence by increasing the ability of

parasites to evade host immunity. The temporal lag introduces an inertia in the

responsiveness of the system and leads to hysteresis in diversity and prevalence.

This is illustrated in figures 3.16a and b. We found that factors which reduce the

strength of transmission-diversity feedback (e.g. lower P_c) increase this inertia and

make the system slower to reach an equilibrium after perturbation. This effect oc-

curs when transmission intensity is perturbed over regions of the transmission-

prevalence curve, which are not transmission saturated, highlighted in red in figure

3.16c. This could therefore have particular importance for the planning and moni-

toring of control efforts in that the effects of control could take a considerable time to

play out and that historic changes in transmission intensity could have long reaching

implications.
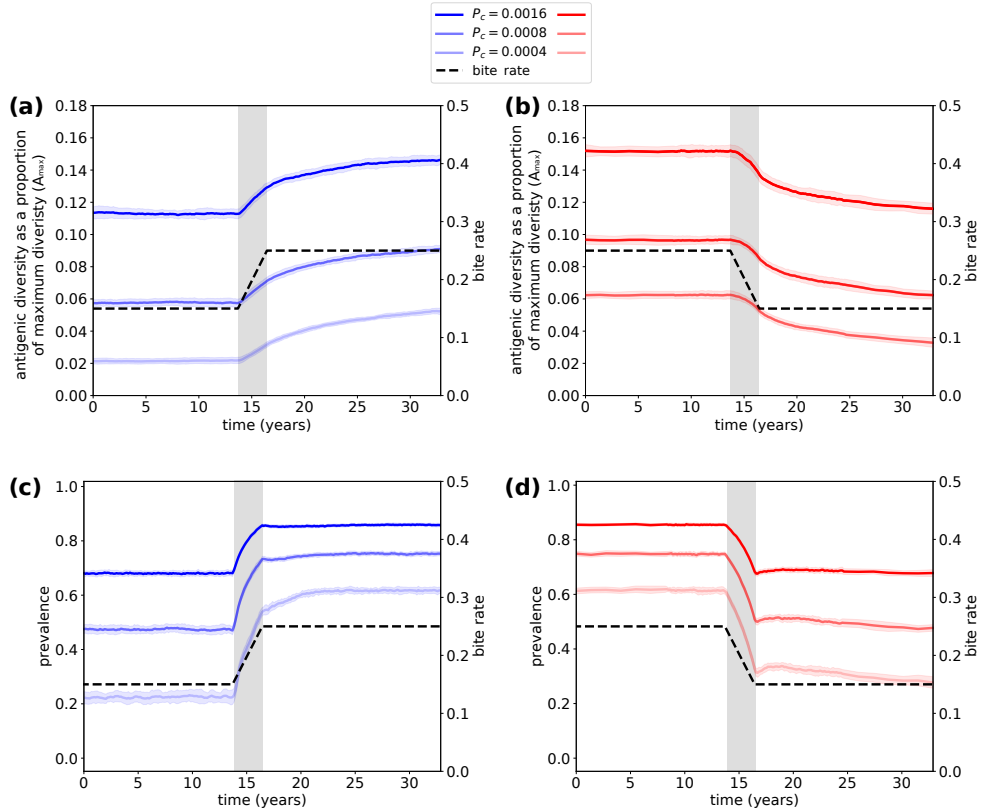
FIGURE 3.15: Perturbations in transmission intensity over time leads to hysteresis in diversity and prevalence. Transmission intensity was manipulated by altering mosquito biting rate (black dashed line) over a short time period (indicated by the shaded region). Transmission-diversity feedback introduces a temporal lag in the response to perturbations. **(a)** Response of antigenic diversity over time resulting from an increase in transmission intensity. **(b)** Response of antigenic diversity over time resulting from a decrease in transmission intensity. **(c)** Response of prevalence over time resulting from an increase in transmission intensity. **(d)** Response of prevalence over time resulting from a decrease in transmission intensity. Bold lines show the mean of 30 repeat and shaded area indicates the 95% confidence interval. Parameters: $H = 8000$, $M = 8000$, $P_s = 0.002$, $P_c = 0.0008$, $A_{max} = 50000$, $\gamma = 0.0$, $A_{init} = 2400$, $S_{init} = 40$.
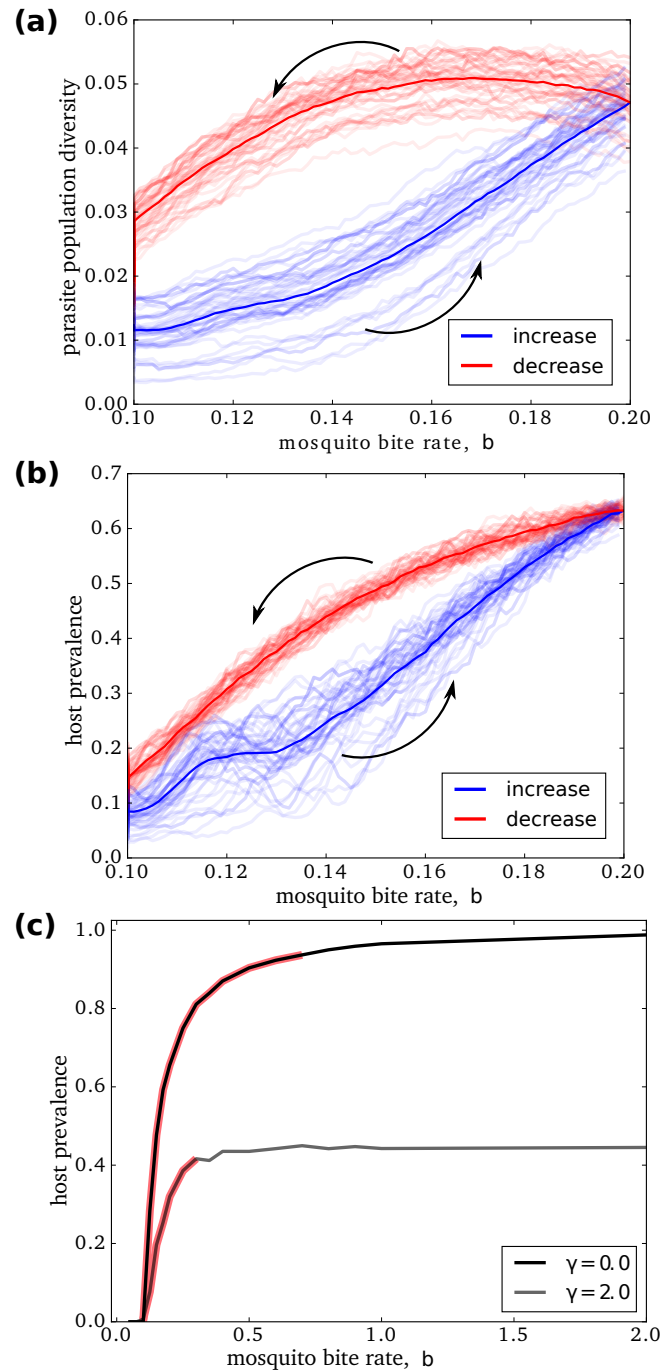
FIGURE 3.16: Phase diagrams showing the hysteresis in the response of parasite diversity (**(a)**) and prevalence (**(b)**), which arises as due to transmission-diversity feedback. Transmission intensity was increased (blue lines) then decreased (red lines) over a short period of time. Bold lines show the mean of 30 repeats while faint lines show individual simulations. The black arrows indicate the direction of time, while mosquito biting rate and parasite diversity/prevalence are shown on the x and y axes respectively. **(c)** Shows two simulated transmission-prevalence relationships with and without cross-immunity, with the regions in which hysteresis is most apparent are highlighted in red. Hysteresis was not as pronounced when the system was in a transmission-saturated state.. Parameters: $M = 8000$, $H = 8000$, $P_s = 0.002$, $A_{init} = 2400$, $S_{init} = 40$, $A_{max} = 50000$.

3.4 Discussion

When diversity is treated as a static property, prevalence is limited by transmission saturation at high transmission intensity, and the level at which prevalence plateaus is strongly influenced by the antigenic diversity in the parasite population. We argued that transmission and diversity generation are inextricably interlinked through the processes of meiotic and mitotic recombination. PfEMP1 diversity is therefore more accurately portrayed as a dynamic property which emerges from the underlying ecology and transmission setting rather than as a static quantity.

To investigate the effects that this has on transmission and prevalence we used a stochastic individual-based model in which new antigen variants could arise through mitotic recombination. We explored the effect that diversity generation has in transmission saturated and diversity saturated systems by developing a model which treats diversity as a dynamic emergent property. This model was used to elucidate the factors that drive prevalence under different transmission and ecological settings. We identified a potentially important mechanism, which could suggest that historic changes in transmission intensity can influence *P. falciparum* epidemiology over much longer periods than previously thought.

Little is currently known about the extent to which individuals with clinical immunity contribute to transmission. A central assumption in our model is that infection length is proportional to the parasite's ability to express novel PfEMP1 variants. This means that duration of infection and host immunity arises naturally from the transmission dynamics of the system. Hosts with greater exposure, on average, experience shorter infections which reduced the potential for onward transmission, and in the context of our analysis, is equivalent to reducing the probability of transmission or the infectivity to vectors. The parasite population was therefore under

selection to diversify in order to evade host immunity. Importantly, diversity did not increase indefinitely but reached an equilibrium in which stochastic extinction balances the rate of generation of new diversity. The level of diversity at equilibrium was determined by the specific disease ecology, including transmission intensity and host population size, and responded to temporal perturbations in these factors accordingly. While we do not expect that adding other non-VSA antigens to our model would effect our results, this could open additional routes by which cross-immunity could lead to transmission saturation and limit prevalence in high transmission regions.

The principal effect of dynamically generating/maintaining diversity was to allow the parasite population to respond to immune selection through diversification. Provided that the ability of the parasite population to find new regions of antigenic space was not constrained, transmission saturation was easily overcome and prevalence increased toward 100% as transmission intensity increased. We therefore hypothesised that another mechanism must be responsible for limiting prevalence at high transmission intensity.

Population size and recombination rate were unable to explain both the rapid initial increase in prevalence at low-moderate transmission intensity and the plateau in prevalence at high transmission intensity. We tested the ability of two mechanisms to limit prevalence at high transmission intensity. Firstly we considered cross-immunity, which could arise through the presence of shared epitopes in the highly mosaic PfEMP1 protein family (Ward et al., 1999; Bull et al., 2008; Nielsen et al., 2004). Secondly, we considered the effect of functional constraints on the ability of antigens to diversify while retaining functional binding to their target receptors. Both of these constraints effectively reduce the size of the antigenic space and can be considered as acting through the same mechanism by limiting the number

of immunogenically distinct antigen variants that can be generated. This in turn limited the ability of parasites to diversify in response to immune selection and re-established transmission saturation at high transmission intensity. While our results should be interpreted qualitatively, we still found that very high numbers of immunogenically distinct antigen variants were required to resolve empirically observed transmission-prevalence curves, which plateau at around the 80-90% prevalence, agreeing well with estimates of large PfEMP1 diversity reported from genetic studies (Barry et al., 2007; Chen et al., 2011; Day et al., 2017). We also found that the number of antigen variants in circulation was closely associated with both transmission intensity and population size, and may explain the relatively low diversity observed in some low-transmission regions (e.g. (Albrecht et al., 2006)). Differences in local disease ecology, such as effective host population size or the ability of parasite populations to maintain antigenic diversity, could be one contributing factor in explaining variation in prevalence across regions with similar entomological inoculation rates.

Diversification facilitates greater immune evasion, which in turn promotes transmission. This forms the basis of a positive feedback loop, which drives rapid non-linear increases in prevalence with increasing transmission intensity. We analysed the effect of this feedback on the response to perturbations in transmission intensity and found that it leads to temporal lags in diversity and prevalence. One consequence of this is that it may take many years for the effects of natural or human-induced changes in transmission to become fully apparent. We would therefore recommend that caution should be used when trying to infer the epidemiological basis of recent changes in prevalence or transmission, and that historic transmission patterns may have long lasting effects. This could have implications for the monitoring of control efficacy particularly in low to moderate transmission intensities,

where the transmission-prevalence curve is steep and the epidemiology is predicted

to be predominantly limited by diversity saturation, or where transmission intensity

is historically variable.

Chapter 4

The maintenance of phenotypically diverse *var* gene repertoires

4.1 Introduction

In the previous chapters we focused on the role of population level diversity in mediating immune evasion and transmission. This chapter is primarily concerned with the diversity within the *var* gene repertoires of individual parasites.

var genes can be categorised into groups based on their chromosomal position, transcription direction and upstream promoter sequence (Ups) (Gardner et al., 2002; Lavstsen et al., 2003). The *var* gene expression products of different groups have been associated with different binding phenotypes and disease outcomes (e.g. (Robinson, Welch, and Smith, 2003; Howell et al., 2008; Vigan-Womas et al., 2008), and for an in-depth review see (Rowe et al., 2009). The clearest example of this is *var2csa*, an UpsE gene encoding VAR2CSA which binds placental CSA and has been implicated in pregnancy-associated malaria (Beeson et al., 2001; Salanti et al., 2003; Salanti et al., 2004). Furthermore, expression of UpsA *var* genes is associated with binding uninfected erythrocytes (rosetting) (Rowe et al., 1995) and severe infection outcomes, and appears to be expressed predominantly in hosts with little prior exposure (Bull et al.,

2005).

Individual *var* repertoires exhibit large diversity both within and between parasite genomes. However, genetic studies have revealed the existence of a highly conserved partitioning structure in which repertoires encode unequal numbers of different *var* groups (figure 4.1). Theory predicts that immune selection acting on populations of pathogens leads to the emergence of structured parasite populations, which minimise the adverse effect of immune interference between parasites (Gupta et al., 1996). An unequally partitioned repertoire is therefore unexpected because it places restrictions on the number of possible non-overlapping repertoire configurations and appears to result in the inefficient use of global diversity (Recker, Arinaminpathy, and Buckee, 2008).
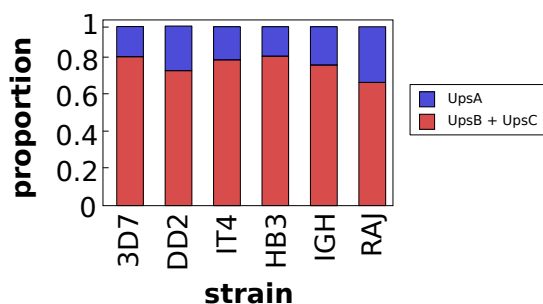
FIGURE 4.1: Approximate partitioning structure of the *var* gene repertoires in six *P. falciparum* strains. Adapted from (Buckee and Recker, 2012). Figure published previously in (Holding and Recker, 2015) (see Appendix A).

For this partitioning structure to be maintained it must confer an advantage to the parasite. However, the evolutionary forces which, select for and maintain repertoire structure, are not well understood. Given the important role that PfEMP1 plays in both immunity and immune evasion, immune selection is a strong candidate. One hypothesis, which arises from the observation that different *var* subgroups are associated with different cytoadhesion and virulence phenotypes, posits that partitions may have evolved to exploit different host niches. According to this hypothesis, the

differential expression of phenotypes enables the parasite to utilise flexible strategies depending on the immune status of the host in which it finds itself (Kraemer and Smith, 2003). It has been hypothesised that anti-disease immunity is acquired relatively quickly because the subset of highly virulent *var* genes associated with severe outcomes exhibits limited diversity. This group may therefore be specialised to naive hosts with limited exposure and experience weaker selection for immune evasion compared to other groups. In contrast, antigens which are adapted to a semi-immune hosts should prioritise immune evasion and exhibit higher sequence variation.

Parasites express PfEMP1 variants in a mutually exclusive fashion, which has presumably evolved in response to strong selection by the host's immune response during infection. Variant specific immunity is thought to select for strain structure, in which strains minimise their immune interference with one another by minimising overlaps in their repertoires. Theoretical work has demonstrated the effect of these processes on the co-circulation and emergence of distinct strains (Gupta, Swinton, and Anderson, 1994; Gupta et al., 1996). Recent work has helped to uncover how a combination of immune selection acting to inhibit individual binding domains (assumed to reduce fitness) and whole antigens can select for antigens which specialise to different host niches (Buckee and Recker, 2012). This offered an explanation for the existence of phenotypes specialised to different host niches, but did not address why parasites should contain antigen repertoires that are phenotypically diverse, and what evolutionary forces could maintain the partitioning structure of these phenotypes. There is currently little theoretical understanding of the evolutionary forces which shape and maintain structure within antigen repertoires.

In the present study we investigate the interaction of host immunity and antigenic diversity to elucidate how immune selection shapes antigen repertoires. Specifically, we consider the effects of variant-specific and cross-reactive immunity on the within- and between-host fitness of *var* gene repertoires using a stochastic individual-based model (the 'repertoire structure model'). Throughout this chapter, we use the term 'phenotype' to refer to a group of PfEMP1 variants with similar binding properties. We found that between-host fitness is maximised by phenotypically homogeneous repertoires, whereas a combination of variant-specific and cross-protective immunity selected for mixed phenotype repertoires. Finally, we show that constraints on diversity in different phenotype groups could select for partitioning structures similar to those observed empirically.

4.2 Methods

4.2.1 The repertoire structure model

Parasite strains are defined by a repertoire of R antigens belonging to either phenotype A or B. We define the repertoire structure of a parasite as the number of variants belonging to each phenotype, i.e. the repertoire structure $n_A : n_B$ refers to any repertoire with n_A antigens of phenotype A and n_B antigens of phenotype B, such that $n_A + n_B = R$. Importantly, we assume no inherent fitness difference between antigen variants or phenotype groups. Each phenotype is assumed to have a global pool of N_A and N_B of antigen variants in circulation.

Hosts are modelled individually and die of natural causes with an age-dependent probability given by

$$p_{death}(T) = \frac{0.002e^{0.07T}}{365},$$

(4.1)

where T is the host age in years. The probability of death is assumed to be constant for each day in a given year. Dead hosts are immediately replaced by a newborn naive host. For simplicity we ignored maternal protection and any risk of mortality associated with infection. Hosts are initialised to be at demographic equilibrium.

Hosts can be infected by one strain at a time and, upon infection, parasites are assumed to attempt to express their entire antigen repertoire. Each successfully expressed antigen makes an incremental contribution to the duration of infection. Hosts acquire variant-specific immunity to any antigen variant they are exposed to, which blocks antigens from contributing to future infections in a variant-specific manner. We also considered the scenario in which the acquisition of immunity was probabilistic, with immunity being acquired with a probability proportional to the infection contribution induced by that variant.

Where considered, cross-immunity was assumed to act such that exposure to an antigen induces partial immunity to other antigens, the net effect of which is to reduce the benefit that the parasite gains for each subsequently expressed antigen. This was implemented by reducing the contribution to infection length that antigens make. Cross-immunity was considered either as a transient and permanent change to the immune status of the host. In the case where cross-immunity was transient, any cross-reactive aspect of immunity was reset upon clearing the infection. Variant specific immunity is always permanent and can only be lost through death. Additionally, cross-immunity could act in a phenotype-specific manner (i.e. that exposure to group *A* antigens does not affect the fitness of group *B* antigens and *vice versa*), as well in a phenotype-transcending manner.

The contribution a single antigen makes to the length of infection is given by

$$l_k = l_0 e^{-k\gamma}, \tag{4.2}$$

where γ is controls the strength of cross-immunity. l_0 is the contribution to the duration of infection for a novel antigen variant. Variant specific immunity is realised through the special case where $\gamma = 0$. The interpretation of k depends on the type of cross-immunity considered. When cross-immunity is transient and phenotype-specific, k is the number of antigen variants previously expressed of the same phenotype group during an infection. When cross-immunity is permanent and phenotype-specific, k is the total number of antigen variants from the given phenotype group to which the host has been exposed. When cross-immunity is phenotype-transcending and transient, k is the sum of the number of previously expressed antigens in the current infection. When cross-immunity is phenotype-transcending and permanent the k is the total number of antigens to which the host has been exposed over its life.

The total duration of infection for a host infected by a parasite expressing c novel antigens is thus given as

$$L = \sum_{k=1}^{c} l_k = \sum_{k=1}^{c} l_0 e^{-c\gamma} = l_0 \frac{1 - e^{-c\gamma}}{1 - e^{-\gamma}}, \ \ \gamma \geq 0. \tag{4.3}$$

The rate of transmission was kept constant for simplicity, and the probability of a host being infected by strain i was frequency-dependent, i.e.

$$p_{inf} = \lambda_0 \frac{s_i}{\sum_j s_j}, \tag{4.4}$$

where $\lambda_0 = 0.033$ is the daily transmission rate and s_i is the proportion of hosts currently infected with strain i. For simplicity, we did not consider mosquitoes, and hence strains are transmitted directly between hosts.

A strain's fitness is primarily determined by its ability to induce long infections

and thereby to increase its chance of onward transmission. Parasite fitness is dependent upon the immune status of the infected host, because host immunity modulates the number of antigens expressed as well as the contribution that each antigen makes to the duration of infection.

A very large number of strains are possible through combining the global pool of antigens. We therefore initialised the parasite population by generating one strain for each possible repertoire structure (i.e. one of each $R : 0, R - 1 : 1, ... , 0 : R$). Strains were generated by selecting antigen variants without replacement from the antigen pools as appropriate for each repertoire structure. Each simulation was run until it reached a dynamic equilibrium (around 10,000 simulated days), by which point most strains would be driven to extinction. The remaining strains were considered to be 'dominant' regardless of their abundance in the population. Dominance is therefore taken as a qualitative measure of the ability of a set of strains to competitively exclude all other strains. We ran repeated simulations and calculated the frequency with which each repertoire structure is represented by a dominant strain. After each simulation, the model was initialised again with a new set of strains. The aggregated dominance frequency therefore provides a measure of the relative fitness of each repertoire partitioning structure. Simulations were repeated until a clear structure emerged, which typically took between 5,000 and 10,000 simulations.

4.2.2 Extension to the dynamic diversity model

We extended our individual-based 'dynamic diversity' model from Chapter 3 to utilise antigens from more than one phenotypic group and to test whether our results represent evolutionary stable strategies. As before, we focus on two phenotype groups, A and B. Antigen space, parasite repertoires and host immune status are defined by two vectors (one for each phenotype group). Parasite repertoires are

therefore defined by a total of R genes, which can come from either of the antigen groups, i.e.

$$s = \{\mathbf{g_A}, \mathbf{g_B}\}, \tag{4.5}$$

where strain s is a set of two column vectors of antigens, g_A and g_B, with lengths N_A and N_B for antigen groups A and B, respectively. The presence (or absence) of a particular antigen type i is indicated by a 1 (or 0) in row i of the strain vector for the relevant group. These group vectors correspond to the strain vector described in Chapter 3.

The immune state of a host is similarly defined by a set of vectors comprising an immune status vector for each phenotype group:

$$h = \{\mathbf{h_A}, \mathbf{h_B}\}, \tag{4.6}$$

where each immune vector is a column vector of size N_A and N_B for antigen groups A and B, respectively. The previous definitions of infection length, immune acquisition and recombination are similarly adapted to apply over each group. The duration of an infection is thus given by summing the contribution to infection made by each antigen in each group,

$$\omega(h, s) = \alpha((\mathbf{1} - \mathbf{h_A})^\mathsf{T} \mathbf{s_A} + (\mathbf{1} - \mathbf{h_B})^\mathsf{T} \mathbf{s_B}), \tag{4.7}$$

or in the general case

$$\omega(h, s) = l_0 \left(\sum_{i=1}^{G} (\mathbf{1} - \mathbf{h_i})^\mathsf{T} \mathbf{s_i} \right), \tag{4.8}$$

where l_0 is the baseline contribution to the infection length that a single antigen makes. G is the number of groups, and $\mathbf{h_i}$ and $\mathbf{s_i}$ are the immune vector and strain vector for antigen group i, respectively.

Immunity is acquired independently for each antigen group (whether variant-specific or cross-immunity), i.e. for variant specific-immunity the host's immune state changes with respect to antigen group i as

$$\mathbf{h}_{i,t_2} = \phi(\mathbf{h}_{i,t_1} + \mathbf{s_i}), \tag{4.9}$$

where t_1 is a time point immediately before infection, t_2 is the time point immediately after infection and s_i is the strain vector of the infecting parasite for antigen group i. ϕ is an element-wise function which ensures $0 \geq h_i \geq 1$ for all elements, i.e.

$$\phi(\mathbf{x}) = \begin{cases} 1 & \text{if } x_i \geq 1 \\ x_i & \text{if } 0 < x_i < 1 \qquad \text{for all } x_i \in \mathbf{x}. \\ 0 & \text{if } x_i \leq 0 \end{cases}$$

4.2.2.1 Recombination

There is evidence that recombination is more likely between *var* genes of the same Ups group (Kraemer et al., 2007). Differences in transcription direction, domain level structure and associated phenotype between Ups groups mean that recombination between similar grouped antigens is more likely to result in a functional and viable recombinant. Furthermore, the fact that *var* groups with different attributes exist suggests that inter-group recombination is not common, otherwise one would expect that the distinctions between groups become less pronounced, and eventually disappear. In order to understand the robustness of our model to the specific

assumptions about recombination, we tested different types of recombination. Meiotic recombination can be group-agnostic or strictly intra-group. When meiotic occurs between antigens of the same group, the repertoire partitioning structure of one of the parents (selected at random) is preserved. We refer to this as having a strict meiotic recombination hierarchy. We also considered the scenario in which there is a small probability, p_g, of intragenic recombination producing a recombinant antigen which belonged to a different group than its parent antigens.

4.2.2.2   Cross-immunity in the dynamic diversity model

The method of cross-immunity used in the dynamic diversity model described in Chapter 3 models cross-immunity as an additive Gaussian transformation of the host's immune status vector. We modified this for the present study to use the same implementation of cross-immunity defined above. This allowed for more direct comparisons between the two models.

4.2.2.3   Initialisation

As in the repertoire structure model, we initialised the dynamic diversity model with an equal number of strains for each possible $R + 1$ repertoire structures (typically one strain for each). An even number of mosquito infections were also initialised for each strain. Strains were generated as follows. Initial antigen pools N_A^* and N_B^* were first created as a subset of the total antigen space of each group (N_A and N_B). Then strains are generated from these initial antigen pools by uniformly selecting antigen variants without replacement, such that an equal number of antigen repertoires are created for each possible repertoire structure. Note that overlaps are possible between antigen repertoires but not within the initial antigen repertoires. This creates

a scenario, equivalent to the initialisation scheme of the repertoire structure model, in which each repertoire structure has equal representation in the model.

4.2.2.4 Dominance frequency

In the repertoire structure, model dominance frequency was calculated as the number of times a strain with a particular repertoire structure is part of a dominant set of strains. This was convenient because, in the absence of recombination, new strains could not emerge. We modified this for use with the dynamic diversity model and used the sum of strain abundance for each repertoire structure. This takes into account the possibility that there may be many strains with low abundance for one repertoire structure (high dominance frequency) or a single strain with high abundance (but low dominance frequency). Using repertoire structure abundance rather than dominance frequency avoids biasing the metric towards many low-abundant strains.

4.3   Results

4.3.1   Variant specific immunity

Parasite fitness depends on the ability to induce long infections in hosts, which maximise the potential for onward transmission. When hosts acquired immunity in a variant-specific manner, the duration of infection was assumed to be proportional to the number of novel antigens expressed by the parasite. Parasite fitness is therefore indirectly determined by competition for susceptible hosts with other strains with overlapping sets of antigens. We refer to the immune selection which arises from this interaction as between-host immune selection, because it causes parasite fitness to vary between hosts. To understand how this interaction affects the fitness of different repertoire structures, we initialised the model with equally sized antigen pools of A and B phenotypes. One strain was generated for each possible repertoire structure by drawing antigens from these pools in the appropriate proportions, and each simulation was run until a dynamic equilibrium was reached. Figure 4.2a shows the strain abundance over time of a typical run, in which a small set (typically a pair) of strains begins to dominate and all other strains are driven to extinction through competitive exclusion. The repertoires of dominant strains tend to exhibit low overlap with one another. The distribution of overlaps between pairs of dominant strains is shown in figure 4.2b for 50,000 repeat runs.

We find that variant specific immunity selected for phenotypically homogeneous repertoires. The frequency with which each repertoire structure was included in the dominant set was calculated using 5,000 repeat simulations, and shown in figure 4.2c. Homogeneous repertoire structures (that is, all phenotype A or all phenotype B) clearly emerge as the most frequently dominating structures. Dominance frequency can be interpreted as a measure of relative fitness, and these results therefore

suggest that variant specific immunity alone favours the evolution of phenotypically
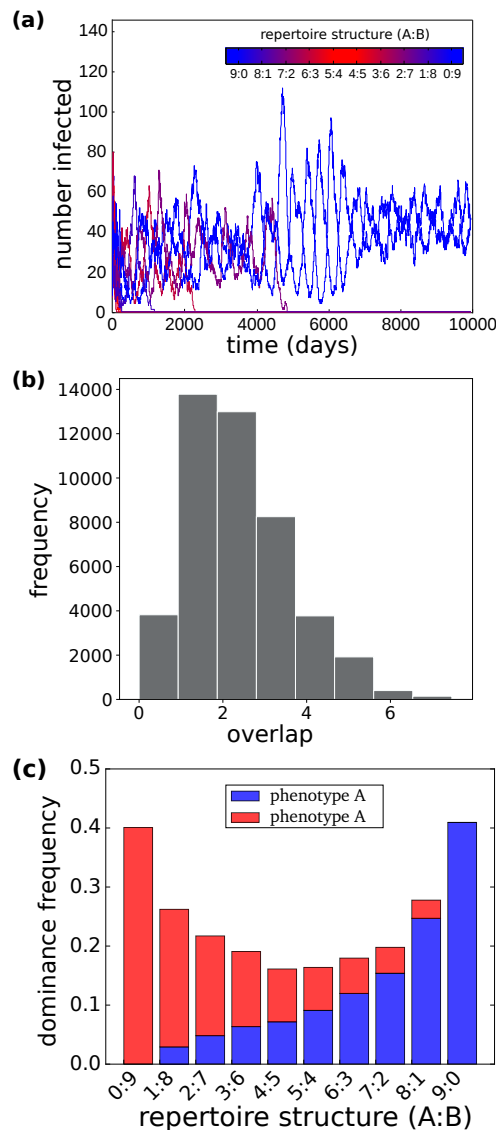
homogeneous repertoires.

FIGURE 4.2: Variant specific-immunity. **(a)** Time series of a single simulation, showing the emergence of a pair of dominant homogeneous strains towards the end of the simulation. **(b)** Distribution of mean overlaps between pairs of dominant strains. The overlap between dominant strains is typically low. **(c)** The distribution of dominant repertoire structure frequency for 5000 repeated simulations shows that homogeneous repertoire structures are most commonly selected for. Parameters: $R = 9$, $N_A = N_B = 13$, $S_{init} = 10$. Subfigure **(a)** is adapted from previously published work (Holding and Recker, 2015).

Theoretical work demonstrates that immune selection acts to structure parasite

populations into dominant sets of strains which exhibit minimal antigenic overlap

to one another (Gupta et al., 1996). Pairs of homogeneous repertoires that have different phenotypes (i.e. $0 : R$ and $R : 0$) minimise the probability of containing within-pair overlaps, while pairs of similarly structured repertoires maximise this. More homogeneous pairs therefore have an advantage because exposure to one is less likely to induce partial immunity to the other. Hence, dominant sets tend to form between repertoire structures at opposite ends of the spectrum. This is highlighted in figure 4.3a, which shows the mean overlap between pairs of strains with different repertoire structures.

Pairs of strains are more likely to overlap one another if they have similar repertoire structures and are therefore less likely to form a dominant pair. Pairs of repertoires which are least similar to one another, i.e. $0 : R$ and $R : 0$, are guaranteed to contain no overlaps. In dominant sets with more than two strains, this would not necessarily be the case and even homogeneous strains would be likely to overlap with a third strain. However we would still expect the same overall combinatoric logic to hold. To test this, the frequency with which different repertoire structures dominate was calculated for each dominant set size, using the same data as used to produce figure 4.2b. The largest dominant set size was four, with most dominant sets being size two (2652) or three (2192), leaving just 156 dominant sets of zero, one or four strains. The distribution of dominant repertoire structures for each dominant set size are shown in figure 4.4. We observed no difference in these distributions with differently sized dominant sets, confirming that the same processes are likely guiding repertoire fitness regardless of the size of the dominant sets.

That homogeneous repertoire structures appeared to be strongly selected for is at odds with the observed structure in sequenced *var* gene repertoires, which exhibit phenotypically diverse repertoires with a strongly conserved partitioning structure (Buckee and Recker, 2012).
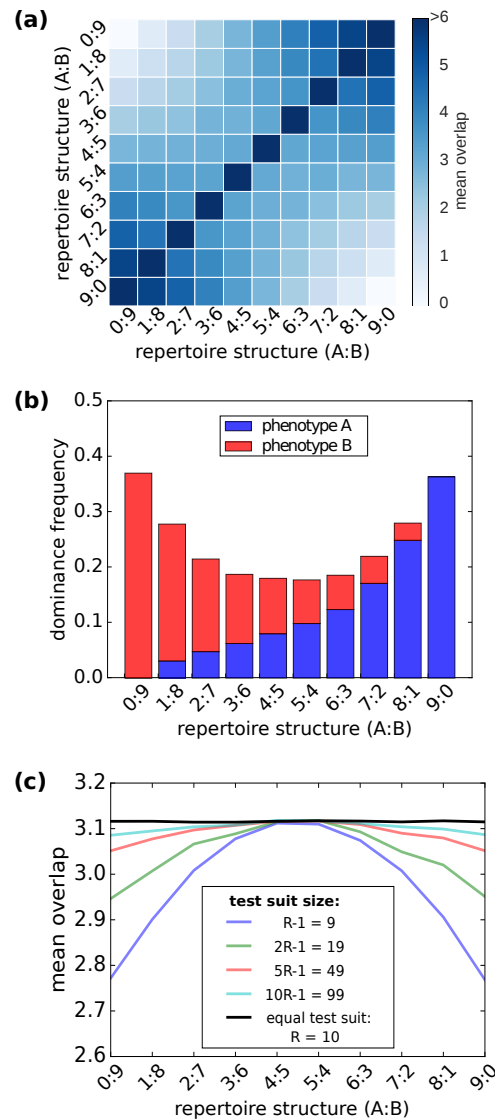
FIGURE 4.3: **(a)** The expected overlap between pairs of randomly generated strains. Overlaps are most common between similarly structured strains, while mixed phenotype repertoires are less likely to overlap with similar structures, and homogeneous repertoires have a much lower probability of overlapping with dissimilar structures. **(b)** When $S_{init} = 100$ (compared to $S_{init} = 10$ in figure 4.2b) there was no change in the relative fitness of repertoire structures. This is despite each strain competing with multiple strains with self-structured repertoires. **(c)** The fitness of each repertoire structure was tested independently by using mean overlap against a test suite of strains as an negatively correlated measure of fitness. When S_{init} increases we find that the number of overlaps becomes more uniform, but the relative difference between repertoire structures remains the same. When test suites utilised an equal number of strains for each repertoire structure these differences disappear, suggesting that repertoire structures are equally fit. Other parameters: $R = 9$, $N_A = N_B = 13$.

The simulations used to produce figure 4.2 were each initialised with one strain
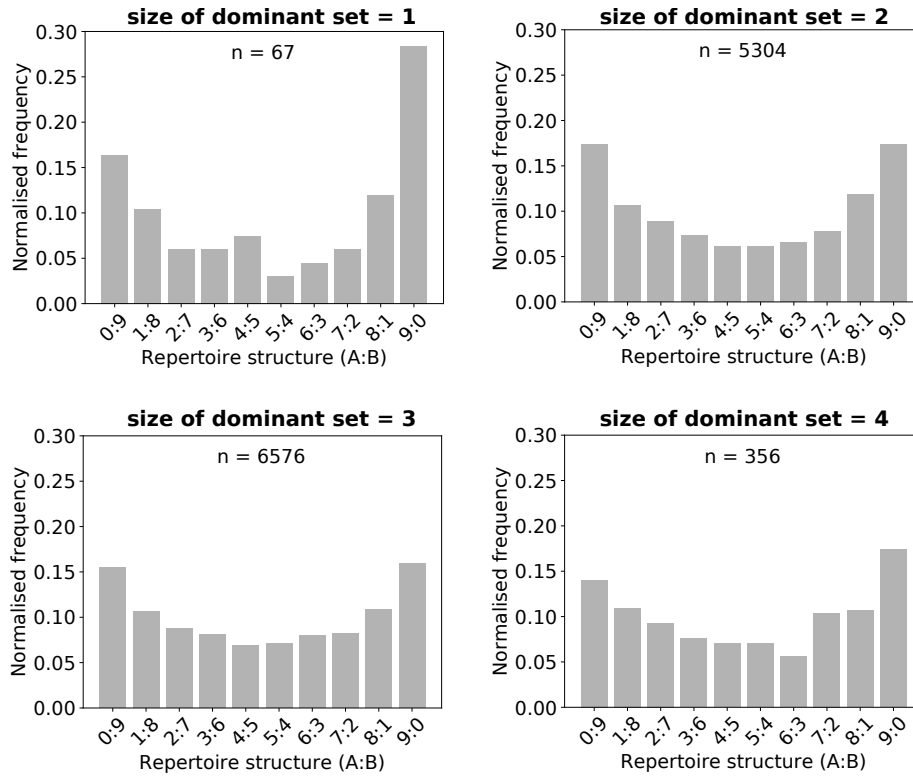
FIGURE 4.4: The distribution of dominant repertoire structures for 5000 repeated simulations plotted according to the size of the dominant set. All frequencies are normalised by the number of strains used (*n* in the figure). No difference was observed in the relative frequency of repertoire structures for different sized dominant sets. Parameters: $R = 9, N_A = N_B = 13, S_{init} = 10$.

to represent every possible repertoire structure. This is a somewhat artificial situation and does not necessarily reflect the dynamic setting in which strains would be competing or evolving. Strains would presumably not only be competing with strains of different repertoire structures but also with strains of the same structure, especially if there is a selective advantage associated with particular structures. In other words, while these results demonstrate that variant specific immunity selects for phenotypically homogeneous repertoire structures, we cannot directly comment on whether or not this constitutes a stable evolutionary strategy. To further understand this, we repeated the above analysis by initialising the model with 2, 5 and 10 strains per repertoire structure. When higher numbers of strains are used, each strain

is competing with a larger number of self-structured strains as well as differently-structured strains. We observed no difference in the distribution of dominant repertoire structures (for comparison, the distribution of dominant strain structures for a total of 100 strains (10 per repertoire structure) is shown in figure 4.3b), suggesting that evolution toward homogeneous repertoire structures would occur even in the presence self-structured competitors.

It is a constraint of our simulation framework that there is necessarily one less self-structured competitor than there is of every other strain structure. Since mixed-phenotype repertoires effectively draw their antigens from a larger global pool, overlaps between two self-structured strains with mixed-phenotype repertoires are less common than between two self-structured strains with more homogeneous structures. Differences in the relative fitness against self-structured strains could therefore be the cause of these differences. To mitigate this, we tested the fitness of each repertoire structure independently by randomly generating strains and measuring the number of antigen overlaps present in each competitor. For each structure a single test strain was generated alongside a suite of competitor strains, and the mean overlap between the test strain and each test suite strain was calculated. This process was repeated until a clear pattern emerged; and is shown in figure 4.3c. The number of overlaps is taken to be negatively correlated with repertoire structure fitness. When we used test suites which emulate the initial conditions of the model, we find that the overlap distribution correctly predicts the fitness distributions shown in figure 4.2c and 4.3b. As expected, the mean number of overlaps increases as strains compete with larger numbers of self-structured competitors and the absolute difference in overlaps between strain structures was larger for smaller test suite sizes. However, the relative difference was constant (i.e. normalising the mean overlaps between 0 and 1 produced identical distributions regardless of the number of strain

used), suggesting that the strength of between-host selection may not change. Interestingly, when we used test suites with an equal number of strains of each repertoire structure (including the self-structure), the differences in mean overlaps between strain structures disappeared. This is shown by the black line in figure 4.3c, and suggests again that there may be no differences in the average fitness contributed by different repertoire structures when immunity is variant specific.

Taken together, these results suggest that homogeneous repertoires may be favoured in the repertoire structure model as an artifact of the initialisation scheme, and not because they minimise overlaps with other repertoire structures. Homogeneous repertoires are more likely to contain overlaps with self-structured strains compared to more heterogeneous repertoire structures and therefore receive a greater advantage when competing against one fewer self-structured strain than do strains with phenotypically mixed repertoires. However, it is difficult to extrapolate this to the complex dynamics of freely evolving parasite strains, and it is difficult to draw a firm conclusion based on this analysis alone. In either case, we do not anticipate this to affect our results because the aim of this study is to explore the conditions which result in selection for phenotypically mixed repertoires. The case of variant specific immunity serves as a convenient baseline from which to investigate this regardless as to whether the distribution of repertoire structure fitness is uniform or favours homogeneous repertoires. We return to this question of whether homogeneous repertoires represent a stable evolutionary strategy towards the end of this chapter.

4.3.2 Cross-immunity

Next, we investigated the effect of cross-immunity on repertoire structure fitness. We modelled this by adding phenotype-specific cross-immunity, which leads to diminishing returns in the benefit that each subsequently expressed antigen confers to the parasite by attenuating the contribution to infection length. The rate at which the contribution of expressed antigens diminishes determines the strength of cross-immunity, controlled by σ. The effect of cross-immunity is shown in figure 4.5a for different values of σ. Because we assume that cross-immunity is facilitated by shared conserved epitopes between functionally similar antigens, the effect of cross-immunity was initially assumed to act in a phenotype-specific way, such that exposure to type *A* antigens has no bearing on the fitness of type *B* antigens and *vice versa*.

We used the same strain initialisation scheme as described previously and repeatedly simulated the transmission dynamics to determine the distribution of dominant repertoire structures (5,000 repeats). Figure 4.5b shows a typical time series for strain abundance, in which a dynamic equilibrium is reached when a pair of mixed phenotype repertoires dominate the transmission dynamics by competitively excluding all other strains.

We found that cross-immunity selects for antigen repertoires that exhibit maximum phenotypic diversity, i.e. with an equal number of type *A* and type *B* antigens. Importantly, this result held both when cross-immunity was either transient (reset between infections) or permanent. The effect of cross-immunity is to decrease the contribution that each antigen makes to infection length; hence, parasites which express many antigens of a single phenotype incur a greater cost due to cross-immunity. As cross-immunity becomes stronger, it is possible for parasites to reach a

point in which there is essentially no advantage in expressing additional antigens of a particular phenotype (e.g. $\sigma = 0.7$ in figure 4.5a). This puts an upper limit on the number of antigens which can be usefully expressed. Crucially, this limit can only be overcome by expressing antigens of a different phenotype. Even without this limit, diminishing returns in the advantage gained from one phenotype can mean it is more advantageous to express antigens of another phenotype, despite an increased risk of there being overlaps with other strains. Mixed phenotype repertoires minimise the detrimental effect of cross-immunity by maximising within-host fitness. The net result is that strains with phenotypically-mixed antigen repertoires can, on average, induce longer infections. Given sufficiently strong cross-immunity, selection to maximise within-host fitness overcomes the influence of between-host immune selection and phenotypically-mixed repertoires are predominantly favoured. This is illustrated in figures 4.5c and d, which show the distribution of dominant repertoire structures using 5,000 repeat simulations for transient and permanent cross-immunity, respectively. There was no qualitative difference between the effect of transient and permanent cross-immunity because the same evolutionary processes shape fitness in both cases. When cross-immunity is permanent, however, narrower distributions were favoured because the effect of within-host immune selection is amplified compared to when cross-immunity resets between infections.

Immune selection acts on repertoire structure at two scales. Firstly, variant specific immunity leads to between-host immune selection for repertoires which minimise the probability of immune interference with competing strains. This favours homogeneous repertoires. On the other hand, within-host immune selection selects for repertoire structures which minimise the detrimental effect of cross-immunity, and favours mixed repertories. The balance between these opposing evolutionary forces determines the relative fitness of repertoire structures and, as the strength of
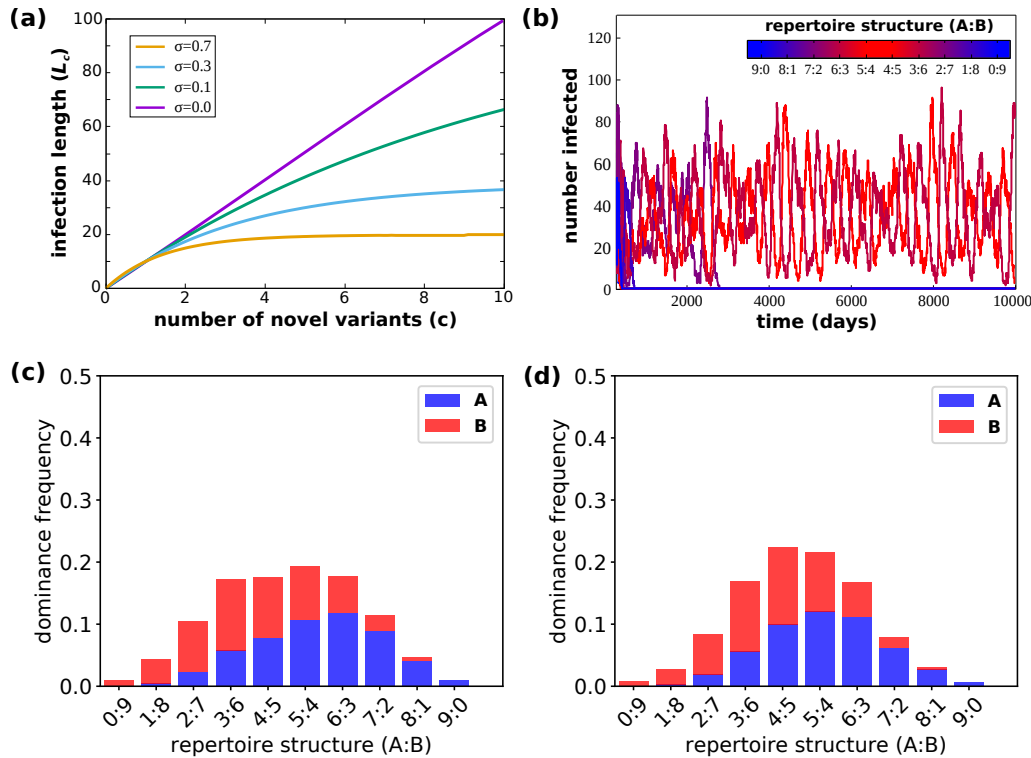
FIGURE 4.5: **(a)** When hosts develop cross-immunity there is a reduction in the contribution to fitness (infection length) associated with the expression of each additional antigen. This limits the maximum benefit associated with expressing additional antigens of the same phenotype. Stronger cross-immunity (higher σ) lead to sharper decreases in fitness. **(b)** Heterogeneous repertoire structures have a selective advantage because they allow parasites to express different phenotypes and minimise the detrimental effect of cross-immunity, thereby inducing longer infections and increasing their potential for onward transmission. **(c)** and **(d)** Simulated transmission dynamics are therefore most commonly dominated by maximally heterogeneous repertoire structures when cross-immunity is either **(c)** transient (reset between infections) or **(d)** permanent. Other parameters: $R = 9$, $N_A = N_B = 13$, $S_{init} = 10$. Subfigures **(a)** and **(b)** are adapted from previously published work (Holding and Recker, 2015).

cross-immunity (σ) increases, the within-host immune selection becomes a more important influence shaping repertoire evolution. This is shown in figure 4.6, which shows the distribution of dominant repertoire structures for different strengths of cross-immunity.

For computational tractability we ran simulations with a relatively small host population ($H = 5000$) and used a smaller repertoire size ($R = 9$) than that of real *var* gene repertoires. We found no qualitative difference in our results when using larger populations (figure 4.7a) or more realistic repertoire sizes (e.g. $R = 60$
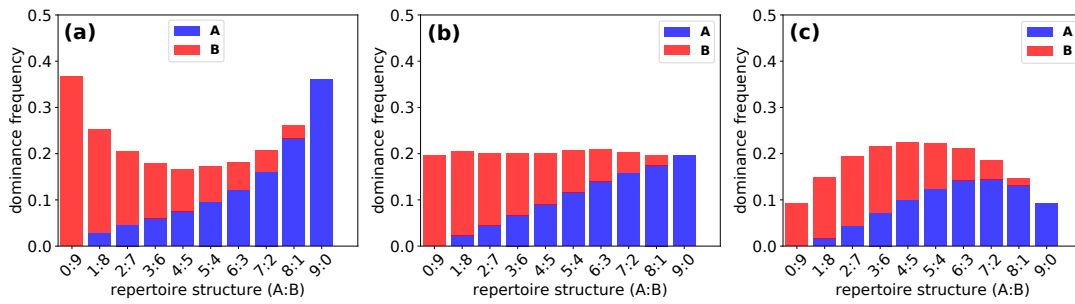
FIGURE 4.6: The strength of cross-immunity determines the balance between within-host immune selection and between-host immune selection. **(a)** $\sigma = 0.01$, **(b)** $\sigma = 0.05$, **(c)** $\sigma = 0.1$. Other parameters: $R = 9$, $N_A = N_B = 13$, $S_{init} = 10$.

in figure 4.7b). Similarly, our results are robust to differences in transmission intensity (figure 4.7c and d), although higher transmission rates scaled up selection pressure because there were more host-parasite interactions and resulted in a more pronounced distribution in dominance frequency.

When we tested the sensitivity of our results to the degree of global diversity we found that it exerted a strong influence on the strength of between-host selection. Increases in global diversity (assuming $N_A = N_B$) decreases the probability of overlaps between repertoires and has a net effect of reducing between-host immune selection. Figure 4.8 illustrates that larger diversity increases the relative strength of selection for mixed repertoires leading, to a more pronounced pattern in the dominance frequency distribution. It should be noted that increasing global diversity is also linked to the number of strains in circulation, with larger antigen diversity facilitating the co-existence of a larger number of strains (i.e. larger dominant sets are possible). In freely evolving parasite populations we would expect this to manifest such that the increase in the number of co-circulating strains largely compensates for the reduced probability of repertoire overlap, resulting in little change in between-host immune selection.

host population size
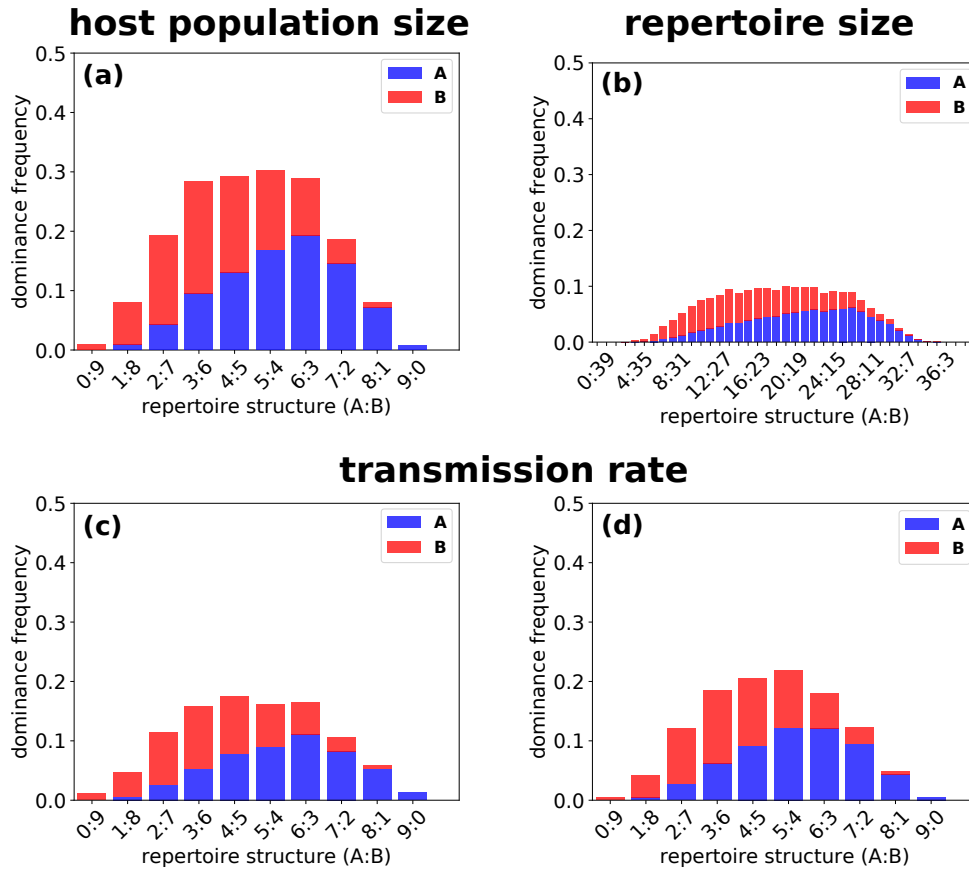
repertoire size

transmission rate

FIGURE 4.7: Cross-immunity robustly selects for phenotypically diverse repertoire structures over a wide range of parameters. **(a)** Large host population size, $H = 20000$. **(b)** Antigen repertoire size $R = 39$, $S_{init} = 40$ and $H = 10000$. **(c)** and **(d)** Transmission rate; $b = 0.25$ and $b = 0.45$, respectively. Unless otherwise specified other parameters are: $H = 5000$, $R = 9$, $S_{init} = 10$, $N_B = N_B = 13$, $\sigma = 0.3$.
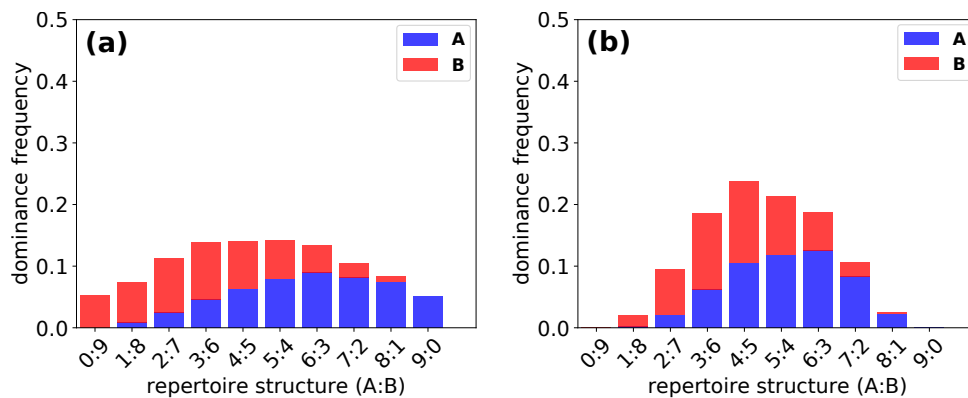
FIGURE 4.8: The degree of global antigenic diversity (i.e. the number of antigens used to initialise strains) decreases the strength of between-host immune selection by decreasing the probability of overlapping repertoires between strains. In relative terms, this results in a stronger effect of cross-immunity, and heterogeneous strains are more likely to dominate. **(a)** $N_A = N_B = 9$, **(b)** $N_A = N_B = 36$. Other parameters: $R = 9$, $S_{init} = 10$, $\sigma = 0.3$.

When we assumed that cross-immunity is phenotype-specific, exposure to antigens in one phenotype group had no bearing on the host's immunity to antigens in the other. To test the effect of this assumption we compared this scenario to one in which cross-immunity transcends phenotype group. As expected, we found that phenotype-transcending cross-immunity did not select for phenotypically mixed antigen repertoire. Phenotype-specific cross-immunity allows parasites with mixed repertoires to evade within-host immune pressure by expressing antigens with a different phenotype group. This is not possible when cross-immunity is acquired to all phenotype groups together (phenotype transcending cross-immunity), and hence mixed repertoires do not have a relative within-host fitness advantage compared to phenotypically homogeneous repertoires. With no evolutionary force to counteract the effect of between-host selection, phenotypically homogeneous repertoires are selected for (figure 4.9).
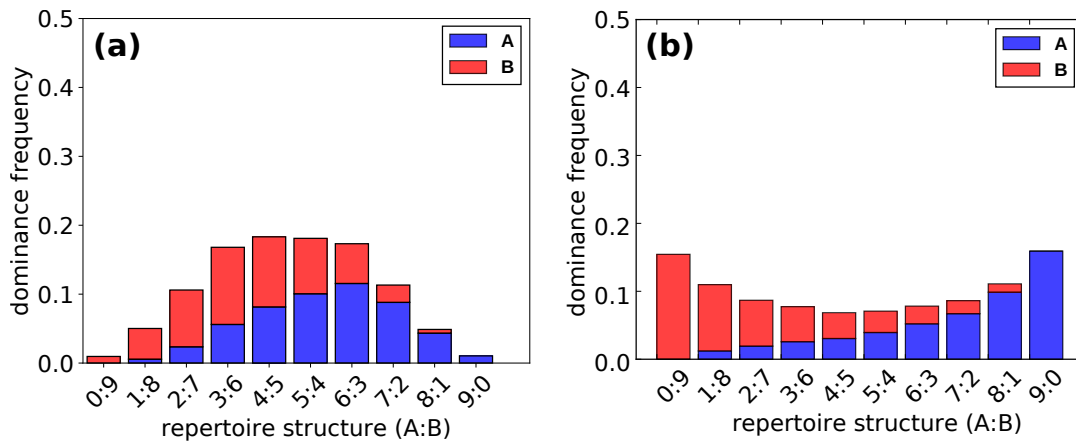
FIGURE 4.9: The distribution of dominant repertoire structures for **(a)** phenotype-specific cross-immunity and **(b)** phenotype-transcending cross-immunity. Parameters used: $R = 9$, $S_{init} = 10$, $N_A = N_B = 13$, $\sigma = 0.3$.

We also considered an alternate mode of immune acquisition in which hosts acquire variant-specific immunity in a probabilistic fashion. In this scenario the probability that hosts acquire long-lasting immunity to an antigen variant is proportional to its contribution to infection length. The rationale being that antigens which are

expressed for long periods will be exposed to the host immune system for longer. When considered alongside the transient cross-immune response this means that antigens which are expressed later in the infection are less likely to induce a long-lasting variant-specific immune response. Under this assumption the order in which antigens are expressed is now important in determining the change in immune status as a result of infection. Two expression ordering schemes were tested: random and *A-first*. In the first, antigens are expressed in a random order for each infection and is the default expression scheme used for all previous simulations. The inclusion of *A-first* expression was motivated by the observation that UpsA *var* genes appear to be preferentially expressed in naive individuals (Bull et al., 2005). We determined the dominance frequency distributions for variant-specific immunity and cross-immunity using random expression and *A-first* expression, shown in figure 4.10. As expected, there was no qualitative change in selection for repertoire structures under *A-first* expression with variant-specific immunity only, because without cross-immunity there is no variation in the time each antigen is expressed (it is either expressed for a constant period, or it is not expressed at all). Interestingly, we saw no change when we introduced cross-immunity, presumably because we assume cross-immunity to act in a phenotype-specific manner, and therefore the expression of group *A* antigens before group *B* has no impact on the acquisition of immunity. We investigate the effect of *A-first* expression on the acquisition immunity and the age-distribution of immunity to different phenotype groups in Chapter 5.

4.3.3 Constraints on phenotype diversity

Antigens are under selection pressure to diversify in order to evade existing immunity but must retain functional binding to their target host receptors. It is reasonable
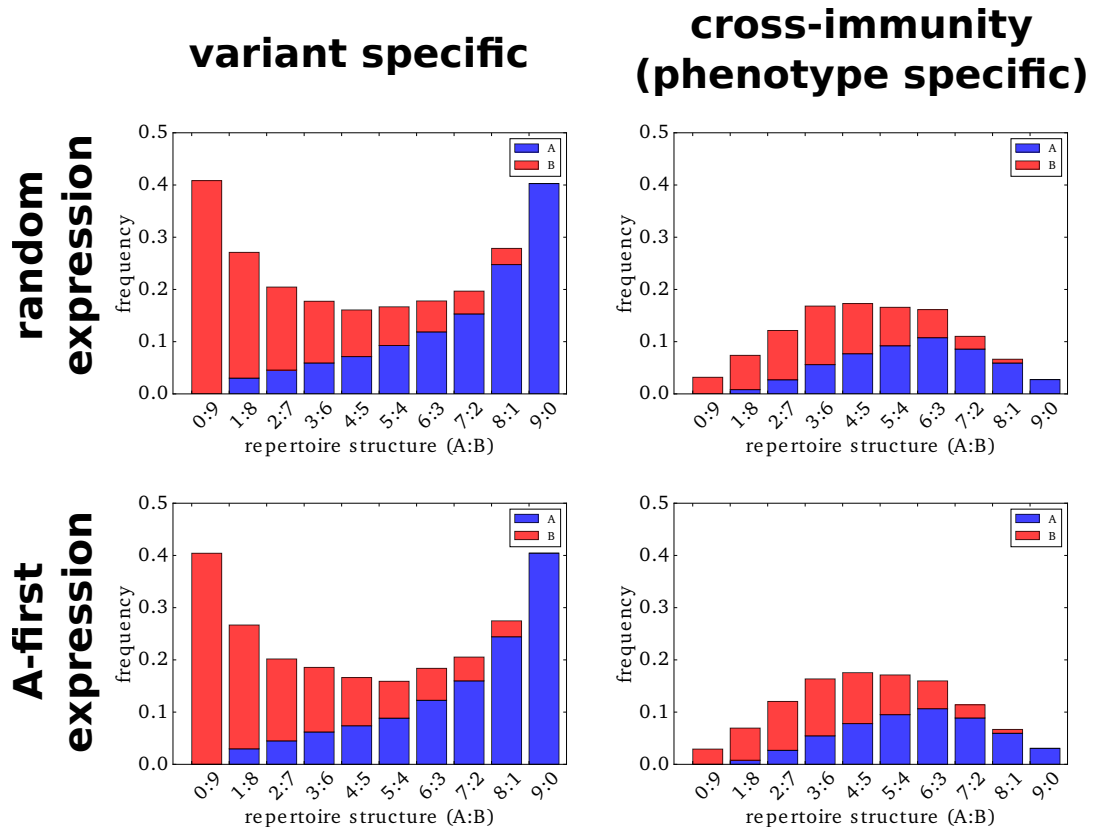
FIGURE 4.10: Expression order had no effect on the distribution of dominant repertoire structures. Variant-specific immunity was acquired with a probability proportional to the duration that each antigen is expressed, such that one effect of cross-immunity is to reduce the likelihood of hosts gaining long term immunity to antigens expressed later in the sequence. $R = 9$, $S_{init} = 10$, $N_A = 13$, $N_B = 13$ where used $\sigma = 0.3$.

to assume that structural differences in the binding targets of different PfEMP1 cytoadhesion phenotypes mean that different phenotype groups are constrained to different degrees. We hypothesised that this could lead to differences in the number of immunologically distinct variants that are possible for different phenotype groups. To test the effect this has on the selection for different repertoire structures we considered a scenario in which differently sized antigen pools are used for different *A* and *B* antigen groups.

The size of the global antigen pools determines the number of antigen variants that are available from which to generate strains during the initialisation of

the model. Larger antigen pools lower the probability of there being overlaps between strains. When one phenotype group has a larger pool than the other, the between-host selection pressure for that group is weaker. When we simulated this scenario using $N_A = 9$ and $N_B = 36$ with variant specific immunity only, this resulted in strong bias for antigen repertoires which contain large numbers of B group antigens (figure 4.11a). As before, we found that adding phenotype specific cross-immunity selected for mixed repertoires, but the relative differences in the strength of between-host immune selection on the phenotype groups meant that the dominance frequency distribution was skewed toward repertoires that contained more group B antigens (figure 4.11b). This suggests that uneven repertoire partitioning structures can be selected for by differences in the global diversity of different phenotype groups.

We next considered the scenario in which differences in the structural similarities of antigens in different phenotype groups led to different degrees of cross-reactivity. The modular structure of PfEMP1 implies that different variants may share sequence regions or whole domains. We hypothesised that this may lead to differences in the extent to which cross-immunity is important in different phenotype groups. The strength of transient cross-immunity in group A and B phenotypes was set to $\sigma_A = 0.6$ and $\sigma_B = 0.1$ respectively. This resulted in a situation in which the advantage associated with expressing multiple group A antigens decreased more rapidly than it did for group B. This resulted in the selection for mixed antigen repertoires which are skewed toward the phenotype group with the lowest level of cross-immunity (figure 4.11c). The extent of the skew could again be determined by the relative difference between phenotype groups, in this case by the relative strength of cross-immunity. The underlying mechanism is essentially the

same as for differences in phenotype group diversity and comes down to the trade-off between within- and between-host immune selection, leading to an optimal balance which favours utilisation of one phenotype group over another. The difference between this scenario and the previous one is that changes to the antigen pool size affects the relative strength of between-host selection while cross-immunity affects the relative strength of within-host selection.

Finally, we considered a third scenario in which we assume that cross-immunity is not gained gradually but that there is a small probability ($p_{\gamma A}$ and $p_{\gamma B}$ for A and B groups respectively) that a host will develop an antibody response which binds a highly conserved but weakly immunogenic epitope. Once acquired, this was assumed to provide group-wide immunity, which prevents antigens from that group contributing to infection length. When we used uneven recognition probabilities for A and B groups ($p_{\gamma A} = 0.01$, $p_{\gamma B} = 0.06$) we find that the repertoire structures selected for qualitatively match those in the other scenarios. While this last scenario may not be the most realistic model for anti-PfEMP1 immunity, it demonstrates that the underlying evolutionary mechanism we have described is robust to the way in which immunity is gained and to the specific implementation used.
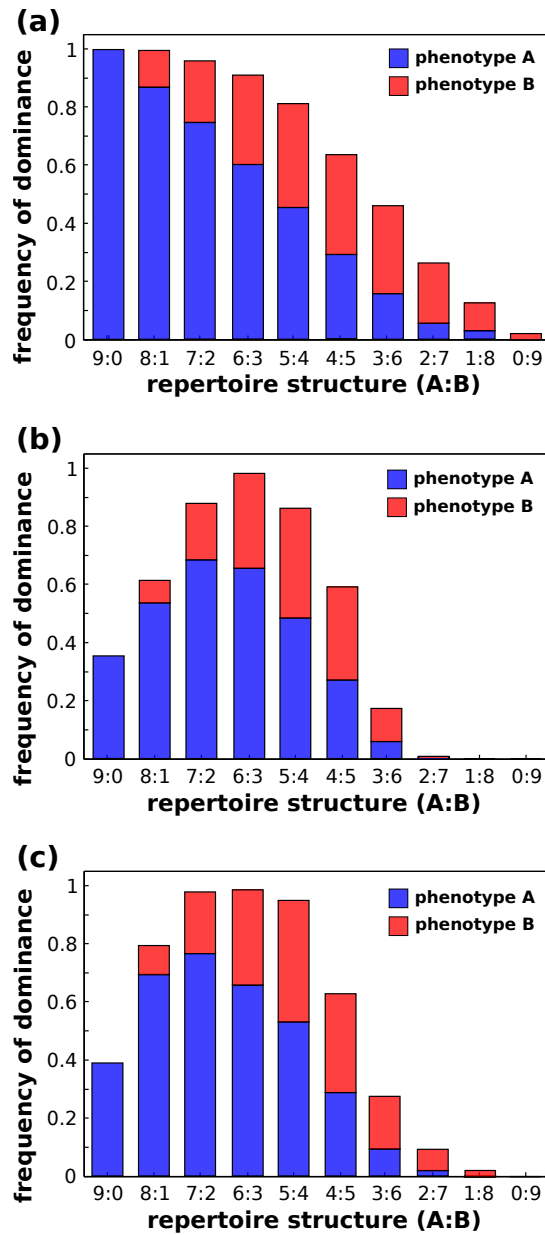
FIGURE 4.11: Differences in the global diversity of phenotype groups can alter the balance between within- and between-host immune selection in a way which selects for repertoires skewed toward one phenotype of the other. **(a)** Differences in global diversity between phenotypes ($N_A = 36, N_B = 9$) with variant-specific immunity. **(b)** Cross-immunity and difference in global diversity ($N_a = 36, N_b = 9, \sigma = 0.1$). **(c)** Different phenotype groups exhibit different degrees of cross-immunity ($\sigma_A = 0.6, \sigma_B = 0.1$). Figure adapted from previously published work (Holding and Recker, 2015).

4.3.4   Validation of results in an evolutionary context

So far we have investigated the evolutionary processes that select for different repertoire structures using a rigid, and arguably artificial, scheme in which each strain competes with one strain from every possible non-self repertoire structure. While this is good evidence for the initial direction of selection, diversity is treated as a static property and parasite populations cannot dynamically adapt to the changing co-evolutionary environment. In this section we test this assumption by using the dynamic diversity model introduced in Chapter 3. We explore the effect of recombination and test whether the results described so far could represent evolutionary stable strategies.

In the repertoire structure model we initialised a single strain for each possible repertoire partitioning structure and measured fitness using dominance frequency (i.e. the number of times a strain of a particular repertoire structure was part of the dominant set). For use with the dynamic diversity model it was necessary to modify this metric (see Methods). In practice we found very little difference between these two metrics, probably because abundant strains have a greater opportunity for recombination and therefore spawn a greater number of strains. Note, we also normalise repertoire structure abundance to between 0 and 1 in order to remove the effect of total prevalence and to facilitate comparison between scenarios. This differs from the normalisation used for dominance frequency, which was normalised by the total number of repeat runs (the maximum possible value).

4.3.4.1 Equivalence to previous model

Firstly, to demonstrate the equivalence of the dynamic diversity model we simulated the key scenarios described so far using equivalent parameters and with recombination turned off. Figure 4.12 shows that the model can reproduce the main results of the repertoire structure model. Variant-specific immunity selected for parasites with homogeneous antigen repertoires (figure 4.12a), while cross-immunity selected for mixed phenotype antigen repertoires. If the global antigen diversity (pool size) contains uneven proportions of *A* and *B* antigens, then between-host immune selection is lower for the more diverse phenotype group, resulting in a skewed dominance frequency distribution (figure 4.12c). Mixed phenotype antigen repertoires with uneven proportions of antigen groups can therefore be selected for when there is uneven within- or between- host immune selection pressure on the different phenotype groups (figure 4.12d).

4.3.4.2 Effect of recombination

It is important to distinguish between the effect of recombination and the effect of selection on repertoire structure. The way that variation is generated can introduce a mechanistic bias that can shape repertoire structure but does not in itself lead to the evolution of fitter parasites. In particular, recombination tends to break up associations between genes and may therefore disrupt the emergence of repertoire structure. To characterise the effect of recombination on the evolution of antigen repertoires under neutral selection we turned off host immunity. Three types of recombination were considered: (1) meiotic recombination, which exchanges whole antigens between repertoires, (2) intra-group mitotic recombination, which takes place between two variants of the same group and produces a new variant of the
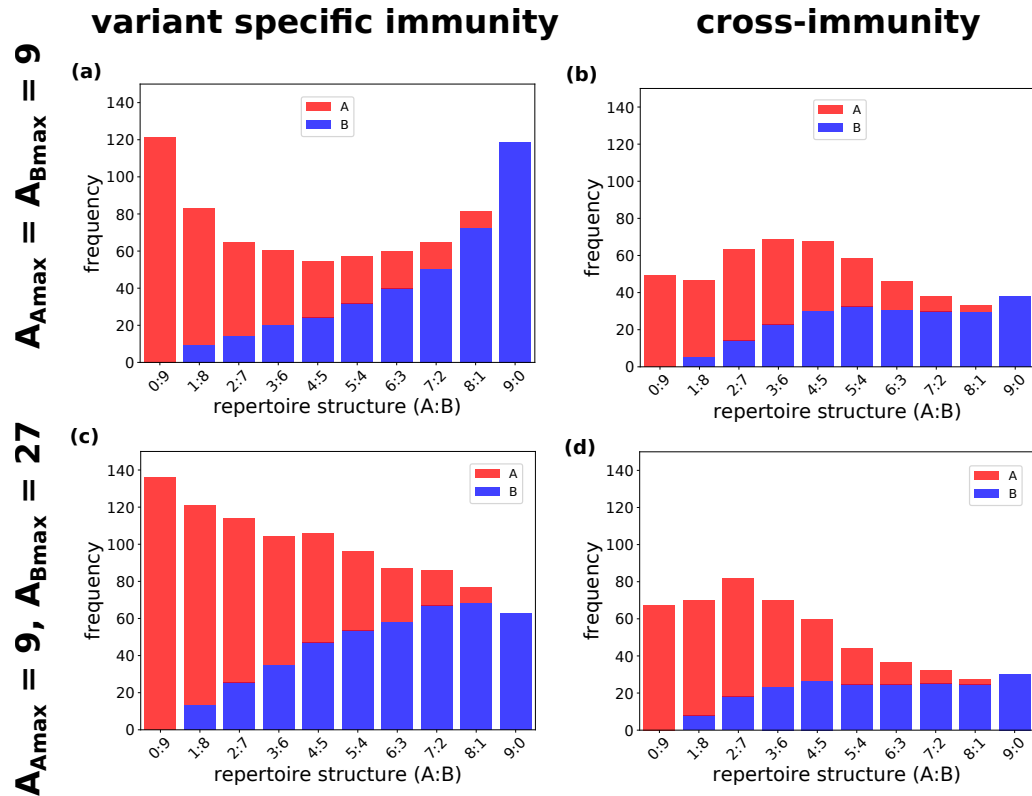
FIGURE 4.12: The distribution of dominant repertoire structures calculated over 6000 simulations using the dynamic diversity model. Immunity is either variant specific (**(a)** and **(c)**) or cross-reactive (**(b)** and **(d)**), and the global antigen diversity is either equal (**(a)** and **(b)**) or unequal (**(c)** and **(d)**). Uneven repertoire partitions maximise parasite fitness when there is within-group cross-immunity and uneven group diversity and are stably maintained. Parameters $M = H = 10000$, $R = 9$, $S_{init} = R + 1$, $p_c = 0$, $p_s = 0$. Where cross-immunity is used $\sigma = 3.0$.

same group, and (3) inter-group mitotic recombination, which takes place between any two variants and can generate a new variant of a novel type. Inter-group mitotic recombination occurs in conjunction intra-group recombination, with a probability of $p_g = 0.05$ per recombination event.

As expected, meiotic recombination tends to generate larger proportions of phenotypically mixed repertoires under neutral selection (shown in figure 4.13a). Intra-group mitotic recombination had no effect on the distribution of repertoire structures, which retained their initial distribution (compare figure 4.14b and d). When inter-group mitotic recombination was allowed, we found that it favoured mixed phenotype repertoires (figure 4.13c). Neutral selection with no recombination, as

expected, showed no clear pattern of change from the uniform distribution (figure 4.13c).
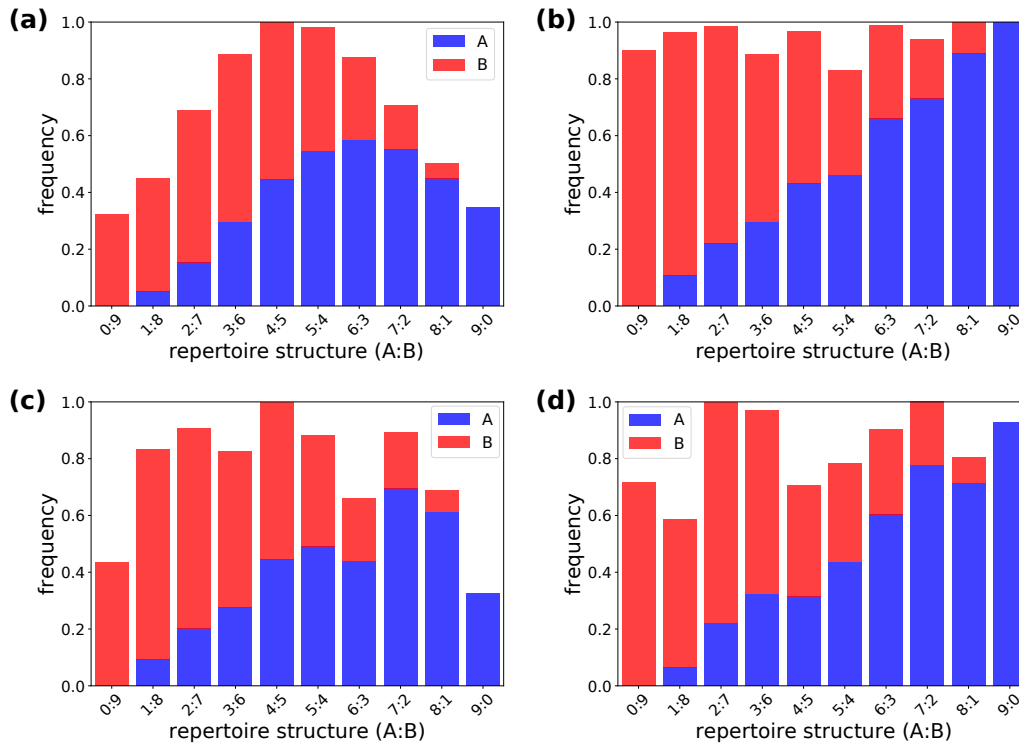
FIGURE 4.13: Mitotic and meiotic recombination can both lead to a bias toward phenotypically mixed repertoires under neutral selection. Dominance frequencies were generated from 10000 simulations for **(a)** meiotic recombination only, **(b)** mitotic recombination with no probability of group-switching ($p_g = 0$), **(c)** mitotic recombination with a small probability of group-switching ($p_g = 0.05$) and **(d)** no recombination. Other parameters: $M = H = 10000$, $R = 9$, $S_{init} = R + 1$, $N_A = N_B = R$.

Since recombination can introduce biases that favour phenotypically mixed antigen repertoires, we tested the robustness of immune selection to select for specific repertoire structures under different types of recombination. We considered variant specific immunity only because we know from the previous section that cross-immunity selects for phenotypically mixed antigen repertoires, which would make it impossible to distinguish between the mechanistic effect of recombination and the effect of immune selection itself. We tested each type of using three rates of recombination. For meiotic recombination the rates were as follows: $p_s = 0.0001$ (low),

$p_s = 0.01$ (moderate), $p_s = 0.05$ (high). For both inter- and intra-group mitotic recombination the rates were as follows: $p_c = 0.00001$ (low), $p_c = 0.005$ (moderate), $p_c = 0.05$ (high). These values were chosen to give examples of the distribution of repertoire structures when immune selection dominates (low recombination), when recombination bias dominates (high recombination) and when both factors influence the equilibrium distribution of repertoire structures (moderate recombination). We found that recombination can strongly influence the distribution of repertoires structure abundance. An exception to this is intra-group mitotic recombination, which did not disrupt selection for repertoire structure. This is shown in figure 4.14 and suggests that recombination hierarchies are an important consideration here. This agrees with other studies which suggest that recombination between antigens of different groups (e.g. UpsA and UpsB) is probably rare (Bull et al., 2008; Kraemer et al., 2007). It is worth noting that theoretical work has demonstrated that immune selection acting at the level of domains can result in divergent selection for distinct groups which exhibit different traits from one another (Buckee and Recker, 2012). This work is reviewed in section 1.3.3.3 and may explain how separate antigen groups are maintained in the population despite occasional recombination between groups. However, since the dynamic diversity model does not consider immunity explicitly acting at this level, this could not be tested. For the remainder of the analysis we restrict meiotic and mitotic recombination to be exclusively within-group.

4.3.4.3 Dynamically evolving repertoires

We re-ran the scenarios in figure 4.12 but this time used recombination and a recombination hierarchy which prevented recombination between antigens of different phenotype groups. This still allows the parasite population to evolve in response to immune selection and discrete strain structure can emerge. Parasite lineages are
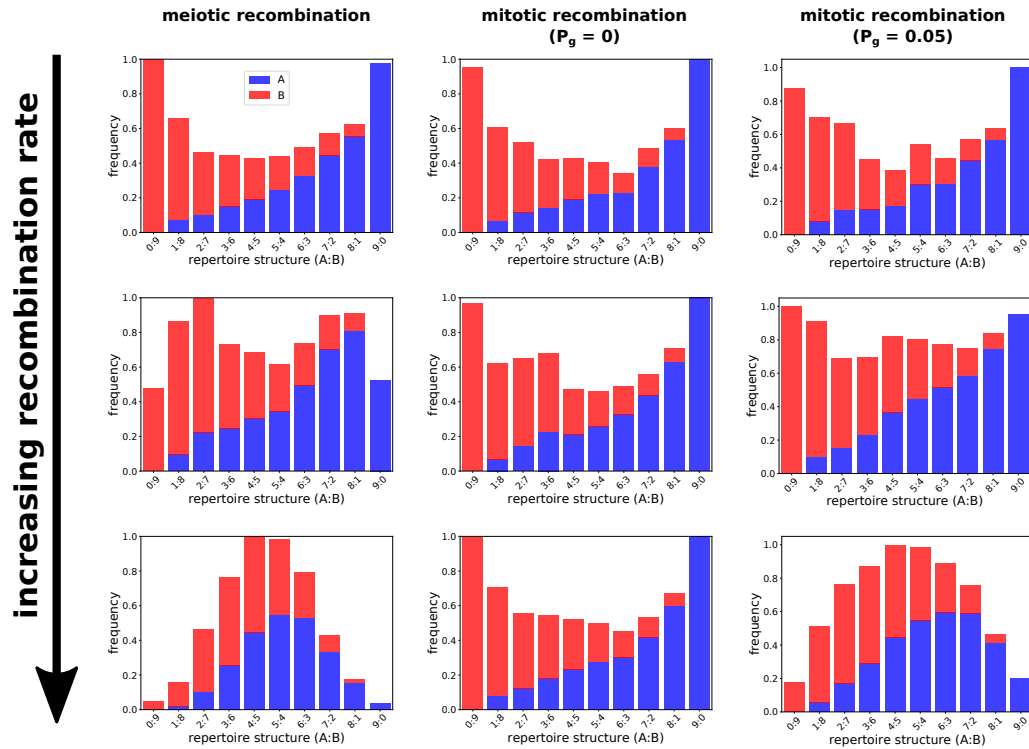
FIGURE 4.14: Meiotic and mitotic recombination can disrupt immune selection for specific repertoire partitioning structures. Meiotic and mitotic (with switching between antigen groups, i.e. $p_g > 0$) overcomes immune selection at relatively low rates of recombination, but mitotic recombination without switching between phenotype groups did not interfere with selection for structured antigen repertoires even when recombination was very high. Meiotic recombination the rates are as follows: $p_s = 0.0001$ (low), $p_s = 0.01$ (moderate), $p_s = 0.05$ (high). Mitotic recombination rates are: $p_c = 0.00001$ (low), $p_c = 0.005$ (moderate), $p_c = 0.05$ (high). These values were selected to illustrate systems dominated by immune selection (low recombination), dominated by recombination bias (high recombination), and systems in which both factors influence the distribution of repertoire structures (moderate recombination).

constrained by their repertoire partitioning structure, but a single genome can diversify into multiple strains. Successful repertoire structures may therefore be expected to generate larger numbers of decedent strains that have the same repertoire structure.

We found that our previous results broadly hold. Variant-specific immunity selected for minimal overlaps between repertoires and favoured phenotypically homogeneous repertoires. This occurred despite the ability of repertoires to change the specific antigen variants which they contain and despite there being many more

possible mixed phenotype repertoires than homogeneous ones (figure 4.15a). When we added cross-immunity we again found that the evolved repertoire structure was determined for by a trade-off between within- and between-host immune selection. Modest levels of cross-immunity were able to overcome between-host selection and favour mixed phenotype antigen repertoires (figure 4.15b). We also found that different degrees of cross-immunity associated with A and B phenotype groups selected for phenotypically mixed repertoires. The distribution of repertoire structure abundance depended on the relative strength of cross-immunity between the two phenotype groups (figure 4.15d). Interestingly, however, differences in the pool size (N_A and N_B) between phenotype groups did not appear to influence the distribution of repertoire frequencies as strongly as in the repertoire structure model (figure 4.15c). While there is a clear skew towards the utilisation of B group antigens (the more diverse group), the most abundant repertoires still exhibit approximately equal partitioning structure. This pattern held even when much higher differences in antigen diversity was used (e.g. a 20-fold difference, $N_A = 9, N_B = 180$), with greater diversity increasing the skew but only marginally change the structure at which abundance peaks. This could reflect the ability of parasites to adapt in response to between-host immune selection by exchanging specific antigen variants, which are at high frequency in the population, for lower frequency variants.
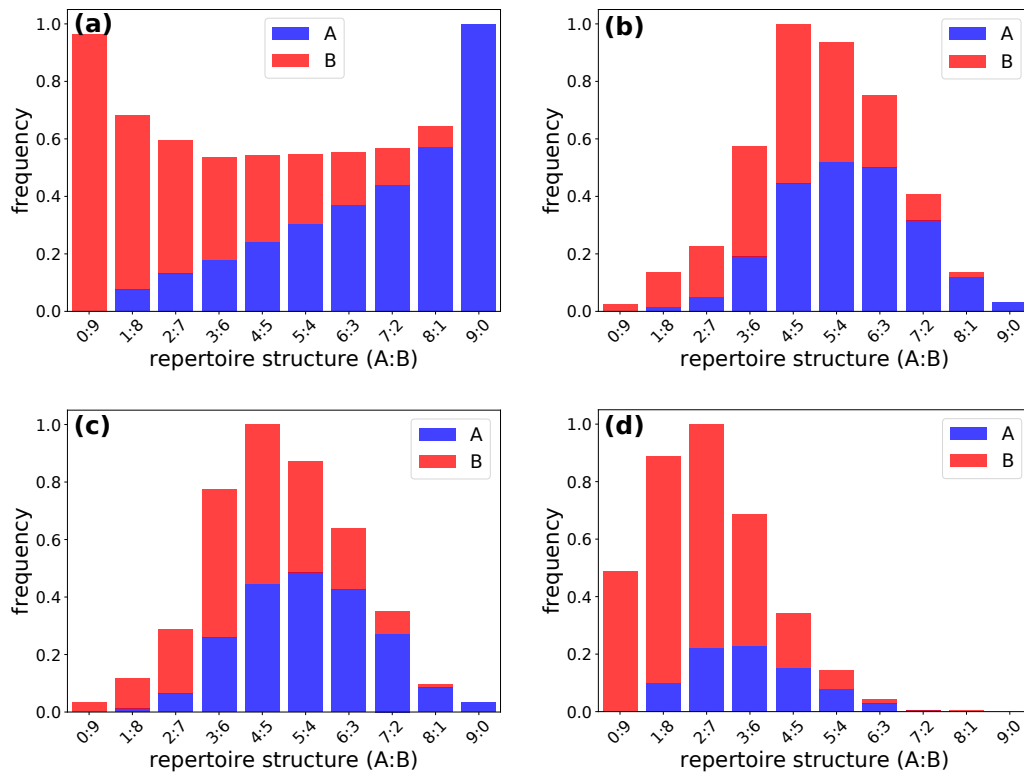
FIGURE 4.15: The abundance of strains utilising different repertoire structures under different immune assumptions: **(a)** Variant specific immunity and no cross-immunity ($\gamma_A = \gamma_B = 0$). **(b)** Variant specific immunity with phenotype specific transient cross-immunity. **(c)** Variant specific immunity with phenotype specific transient cross-immunity, and different sizes antigen pools for different phenotype groups ($N_A = 9$, $N_B = 36$). **(d)** Variant specific immunity with phenotype specific transient cross-immunity and different strength cross-immunity for different phenotype groups ($\gamma_A = 0.6$, $\gamma_B = 0.2$). Based on 10000 repeat simulations. Other parameters: $H = M = 10000$, $R = 9$, $S_{init} = 10$, $b = 16$, $p_s = 0.001$, $p_c = 0.002$, $p_g = 0$. Unless otherwise stated $N_A = 13$, $N_B = 13$, $\gamma_A = \gamma_B = 0.3$.

4.3.4.4   Repertoire structure in diversity and transmission saturated settings

We hypothesised that between-host immune selection on repertoire structure would be highest under transmission saturation because diversity is the main factor limiting prevalence in these settings. To test this we simulated transmission-prevalence curves using variant specific-immunity (without cross-immunity) and calculated the distribution of repertoire structures at different points along the curve. Figure 4.16 shows the normalised abundance of strains utilising each repertoire structure for different points along a simulated transmission-prevalence curve. This indicates that between-host immune selection is weaker and shapes repertoire structures to a lesser extent in diversity saturated settings. When the system is transmission saturated, however, there is strong selection for the parasite population to make efficient use of the available diversity and therefore for repertoire structures which minimise antigenic overlaps.
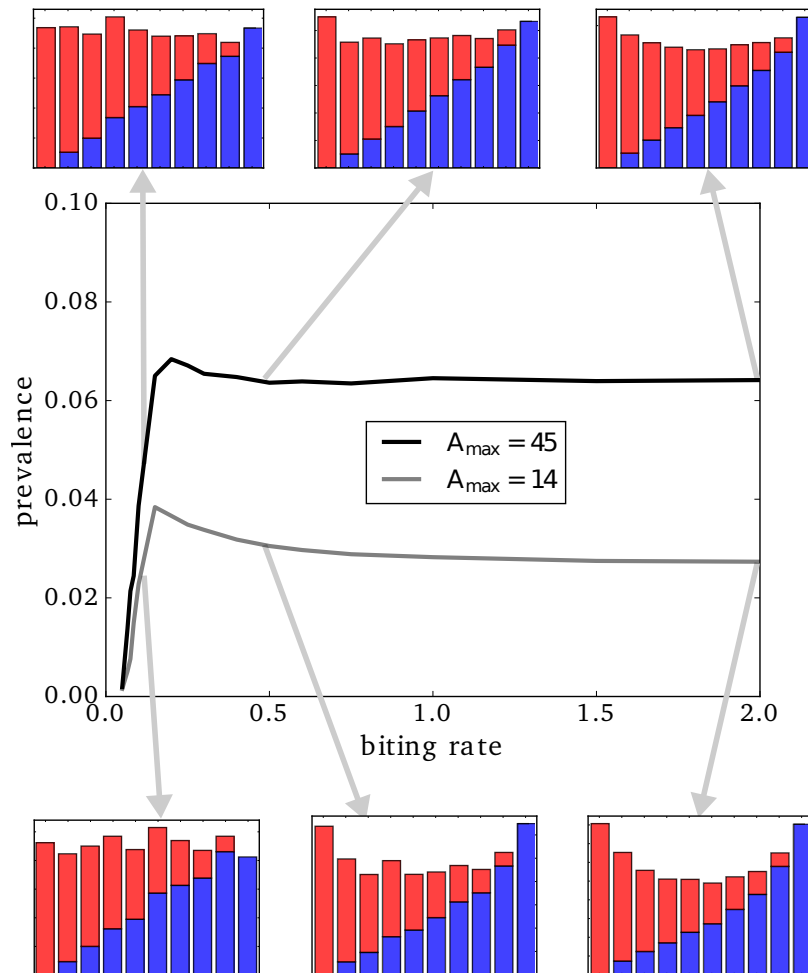
FIGURE 4.16: The effect of transmission intensity on selection for repertoire structures. Low transmission intensity and diversity saturation reduces between-host immune selection and lessens selection for homogeneous repertoire structures. As transmission intensity increases the system transitions into a transmission saturated state where prevalence is diversity limited. In this setting there is strong immune selection to utilise diversity efficiently and repertoire structures evolve to minimise the probability of overlap. Parameters: $N_A = N_B = A_{max}$, $H = M = 10000$, $R = 9$, $S_{init} = 10$, $p_s = 0.005$, $p_c = 0.0008$, $p_g = 0$, $\gamma = 0$.

4.4   Discussion

PfEMP1 is highly immunogenic and forms a crucial component of the parasites immune evasion strategy. PfEMP1 is therefore subject to strong selection pressure. Theoretical work suggests that under such circumstances immune selection will structure parasite populations into antigenically distinct strains, which minimise the immune interference with one another (Gupta et al., 1996). However, sequenced *var* gene repertoires demonstrate the existence of a highly conserved partitioning structure in which groups of genes exist in different proportions. These proportions are well conserved even between parasite genomes taken across geographically diverse regions (Rask et al., 2010; Buckee and Recker, 2012; Kyes, Kraemer, and Smith, 2007). However, the existence of this partitioning structure is counter-intuitive because it leads to inefficient utilisation of the globally available diversity (Recker, Arinaminpathy, and Buckee, 2008), suggesting that there is a selective advantage associated with this structure.

A stochastic individual-based model was used to investigate the evolutionary forces which could select for, and maintain, partitioned antigen repertoires. Under variant specific immunity we found that immune interference between strains selects for phenotypically homogeneous repertoires, which minimise the chance of overlaps between other strains. This was partly due to the initialisation setup we used, in which we randomly generated one strain for each possible repertoire structure. When we independently tested the fitness of each repertoire structure against a suite of competitor strains with an even distribution of repertoire structures, there was no overall difference in fitness across the spectrum of possible repertoire structures. However, when we simulated the evolution of repertoire structures we found

that the original result held, with variant specific immunity leading to greater abundance of phenotypically homogeneous repertoires. In either case we found that variant specific immunity is not sufficient to explain the maintenance of specifically partitioned repertoire structures.

There is evidence that exposure induces some degree of cross-reactive immunity, in which host sera can bind antigens from heterologous parasites (Bull et al., 2002; Smith et al., 1999a). Furthermore, several studies have shown the existence of partially cross-reactive PfEMP1 variants, thought to be caused by the existence of conserved epitopes (Chattopadhyay et al., 2003; Gamain, Miller, and Baruch, 2001). We considered the effect of cross-immunity, which conveyed some degree of protection against parasites that express antigens with similar cytoadhesion phenotypes. This assumed that PfEMP1 variants are constrained in their diversification or that there are structural similarities between the functional binding domains of variants which target the same host receptors. When cross-immunity acted in conjunction with variant-specific immunity we found that even modest levels of cross-immunity selected for maximally heterogeneous antigen repertoires. The balance between within- and between-host immune selection was partially determined by the strength of cross-immunity. One the other hand, phenotype transcending cross-immunity was not sufficient to explain observed repertoire structures.

We showed that uneven repertoire partitions can be selected for by altering the trade-off between within-host and between-host selection in a phenotype specific manner. We demonstrated this in two ways. We hypothesised that the degree to which antigens could diversify was related their binding target leading to different antigen pool sizes for each phenotype. This reduced the between-host selection in the more diverse phenotype group and selected for partitioned repertoires that were skewed towards this group. Secondly, we considered that differences in sequence

conservation between phenotype groups could lead to differences in the strength of cross-reactive immunity. This selected for heterogeneous repertoires, in which the proportions of each phenotype group was determined by the relative strength of within-host immune selection. Both of these scenarios can equally lead to selection for repertoire structures similar to those seen in sequenced *var* repertoires and can be considered part of the same mechanism, i.e that functional constraints on antigen diversification can limit the number of antigenically distinct antigens that can exist. These findings are in agreement with work which suggest that sub-groups of PfEMP which are associated with different phenotypes may exhibit different levels of global diversity.

The robustness of our results to recombination was investigated with respect to inter-group recombination. We found that modest levels of inter-group recombination would to disrupt repertoire structure. This occurred for both mitotic and meiotic recombination and highlights the importance of recombination hierarchies in the maintenance of repertoire structure. When recombination took place strictly between variants of the same phenotype groups, our results broadly held.

While we loosely associated binding phenotype with Ups groups, our findings are not contingent on a one-to-one mapping between the two but elucidate the conditions under which phenotypically diverse *var* gene repertoires are selected for. Provided that whichever grouping scheme is used is associated with binding phenotype, our results should broadly hold. Similarly, we see no reason why our results would not scale to more antigen groups than the two considered here.

It is becoming increasingly clear that cross-reactive immunity could be an important process in explaining *P. falciparum* evolution and epidemiology. It has been

previously proposed as a mechanism capable of synchronising the sequential expression of PfEMP1 during infections (Recker et al., 2004) and implicated in the importance of maintaining chronic infections in individuals with high exposure. In the present study, we showed that cross-immunity could be important in maintaining structured antigen repertoires. An additional consequence of this mechanism is that it limits the number of antigens which can usefully be expressed during a single infection and may therefore offer an explanation as to why *var* genes repertoires are limited in size, whereas other organisms that utilise antigenic variation can have much larger repertoires of antigens (for example *Trypanosoma brucei* (Stockdale et al., 2008)). In contrast, if immunity was solely variantspecific this would predict the evolution of increasingly large antigen repertoires.

While further investigation is needed to determine if cross-immunity acts in a phenotype-specific manner and confirm this mechanism, our results suggest that *var* gene repertoire structure could be maintained by an evolutionary trade-off between maximising within-host fitness and minimising between-host immune selection pressure. A similar evolutionary trade-off has previously been proposed to explain the differences in the domain-level structure of *var* genes (Buckee and Recker, 2012), which again highlights the importance of considering immune selection acting at different ecological scales.

Chapter 5

Structured diversity and the age-dependent burden of disease

5.1 Introduction

In previous chapters we argued that antigenic diversity is inextricable interlinked with transmission and prevalence. Immune selection not only maintains this considerable diversity but also shapes it, leading to emergent structure at multiple ecological scales. Previous theoretical work by others demonstrates how immune selection can structure parasite populations into stable non-overlapping strains (Gupta et al., 1996). In Chapter 4 we showed that a combination of between- and within-host immune selection can select for phenotypically mixed antigen repertoires with specific partitioning structures. Similarly, Buckee *et al.* (Buckee and Recker, 2012) demonstrated that immunity acting at two different scales can explain the evolution of disparate domain level structure in individual *var* genes. These studies show that the way in which diversity is structured has important evolutionary and epidemiological implications for *P. falciparum*.

In Chapter 3 we made the case for diversity being considered a dynamic property, which responds to changes in the specific ecological and transmission settings,

and developed a theoretical framework to investigate the epidemiological consequences of this. However, we have so far only used this framework to look at total antigenic diversity and generally assumed that strain structure will emerge in parasite populations. In the current chapter we examine the robustness of strain structure under different conditions. Specifically, we consider the effect of the underlying transmission and ecological setting on the degree of strain structure by investigating the effect of perturbations in transmission intensity. We also explore the temporal stability of the parasite population structure. The interaction between repertoire structure and population level structure is investigated by exploring how the factors which influence a parasite's antigen repertoire (such as constraints on diversity) influence the degree of strain structure, and we consider the consequences of this on age-dependent *P. falciparum* epidemiology.

5.2   Methods

We used the dynamic diversity model introduced in Chapter 3, and extended in Chapter 4, to investigate the factors which affect the degree of strain structure in parasite populations and their epidemiological consequences. The model is used to generate time series of antigen repertoire abundance, which we analyse for strain structure. In particular, we examine the effect that partitioned antigen repertoires have on the degree of strain structure when different constraints are placed on diversity. This allows us to test whether the evolutionary processes responsible for maintaining partitioned antigen repertoires are consistent with observations of population-level structure and epidemiology.

Throughout this chapter 'genotype' is used to refer to the entire four byte antigen encoding (including genotype only encoding bits). 'Antigen repertoire' is used

to refer to the specific set of antigen IDs within a parasite's repertoire (i.e. ignoring genotype-only encoding bits). 'Strain' is used to denote a group of one or more similar antigen repertoires, which cluster around a particular point in repertoire space. 'Group' is used to refer to the a set of antigens belonging to a particular phenotype (i.e. A or B).

5.2.1 Strain structure clustering

Parasite populations with very large numbers of distinct antigen repertoires are generated even from relatively restricted antigen spaces and small antigen repertoire sizes (R). For example, preliminary simulations showed that about 200 distinct antigen repertoires are maintained at equilibrium when there are 64 possible antigens and a repertoire size of just four ($A_{max} = 64$ and $R = 4$). Analysing and visualising parasite population structure therefore represents a significant challenge. This becomes more difficult with the use of larger antigen repertoires (e.g. $R = 60$), which increases the dimensionality of the repertoire space. In order to quantify the extent of diversity and the degree to which the parasite population is structured into strains, we used a network based method to cluster antigenically similar repertoires. Two parasites were considered to be the same strain if their repertoires shared a certain threshold number of antigens. For an antigen repertoire size of $R = 60$, we set this to $\tau = 50$.

A similarity matrix was calculated such that the similarity of one strain to another is the cardinality of the intersect of their antigen repertoires and normalised by R. The similarity between parasites i and j is thus given by

$$S_{i,j} = \frac{|A_i \cap A_j|}{R},$$
(5.1)

where A_i and A_j are sets containing antigen types in the repertoires of parasites i and j, respectively. A network can then be constructed by drawing edges between parasite genomes i and j if their similarity is greater than the threshold, i.e. $\eta_{i,j} > \tau$. Note that the parasite population is filtered to remove any antigen repertoires with a frequency of less than five prior to clustering. This limits noise which arises from recombinant antigen repertoires that have low abundance and are therefore unlikely to be important in terms of the competitive interactions between parasite lineages. Each strain is represented by a separate sub-graph that is disconnected from other strains. The number of sub-graphs (number of strains) is strongly dependent on the similarity threshold used to determine whether or not an edge should connect two antigen repertoires, τ. Since one of the aims of this study is to test the hypothesis that immune selection promotes non-overlapping strain structure, $\tau = 50$ was chosen as a compromise between tolerating a relatively large degree of overlap between strains, while preventing antigen repertoires with only a small number of differences from being considers separate strains.

Two metrics are used to quantify the parasite population structure; the number of strains (the number of clusters) and the degree of inter-strain overlap. Both of these describe the degree to which the parasite population is organised into immunogenically distinct strains. The average inter-strain overlap is calculated by taking the average number of antigens shared between each pair of strains in the host population. A strain is considered to contain an antigen if any antigen repertoire in the corresponding cluster contains the antigen, i.e.

$$interstrain\ overlap = \frac{|C_i \cap C_j|}{n}, \tag{5.2}$$

where C_i and C_j are the set of antigens contained within clusters i and j, n is the number of pairs given by $n = \frac{k(k+1)}{2} - k$, and k is the number of clusters.

5.2.2 Simulated challenge experiments

Virtual challenge experiments were performed to look for global temporal patterns in the parasite population over time. This involved sampling the host and parasite populations at regular time intervals. The ten most abundant antigen repertoires along with a random sample of 100 hosts were taken to generate a snapshot of the host and parasite populations at each point in time. The average parasite within-host fitness was calculated for each host-parasite pair, whereby each host sample was tested against all historic, contemporary and future parasite samples. A parasite sample was considered 'historic' if it was taken from a point in time before the host sample, and it was considered a 'future' sample if it was taken from a point in time after the host sample. This produced a matrix of parasite fitness from which large scale changes in the parasite population structure, for example cycling between regions of antigen space, can be observed.

The within-host fitness is relative with respect to a particular host because it is dependent on both the specific variants contained in the parasite's antigen repertoire and the particular immune state of the host. It is calculated as

$$f = \sum_{i=1}^{R} a_i * (1 - h_{a_i})/R, \tag{5.3}$$

where a_i is the ith antigen in the parasite's repertoire, h_{a_i} is the value of the host's immune vector corresponding to antigen a_i, and R is the parasite repertoire size. Note that while within-host fitness is proportional to the induced duration of infection, it is normalised to between 0 and 1.

Seasonality

Seasonal fluctuations in transmission intensity was modelled by adding a seasonal term to the mosquito biting rate given by

$$\Delta_b = m \sin(\frac{2t\pi}{365}) \tag{5.4}$$

where t is the time in days and m is the magnitude of seasonality. The biting rate is therefore calculated as

$$b = b_0 + \Delta_b, \tag{5.5}$$

where b_0 is the baseline value for mosquito biting rate.

5.3 Results

5.3.1 Visualising strain structure

To illustrate the effect of diversity on the emergence of strain structure within the parasite population, we first looked at the simplest possible case involving multiple antigen repertoires. By setting $R = 2$ and $A_{max} = A_{init} = 4$ the model can be used to reproduce Gupta's (Gupta et al., 1996) findings, which demonstrated the that immune selection can organise parasite populations into non-overlapping strains. In contrast to previous work, however, our model uses a repertoire of antigens in which alleles do not have fixed loci. The four antigen variants can therefore be combined to create 6 possible antigen repertoires which can be grouped into three discordant (non-overlapping) pairs. There are four additional antigen repertoires which can be generated through recombination, but which are disadvantaged because they contain the same antigen twice and therefore forgo any advantage associated with expressing a second antigen.

Regardless of these differences, theory predicts the emergence of strain structure, and this is observed in our model (shown in figure 5.1a). A single discordant pair dominates the transmission dynamics, and the dynamics of the pair are periodic and out of phase with one another. Invasion by recombinant antigen repertories is inhibited, leading to a stable strain structure. While not an unexpected result, this is an important validation step. It demonstrates that non-allelic repertoires of antigens are subject to the same sort of evolutionary forces to form minimally overlapping strains as parasites with loci-specific antigens. In addition, neither (mitotic or meiotic) recombination nor concurrent infection are fundamental barriers to the emergence of strain structure.

Both the size of the dominant sets and the number of possible antigen repertoires

increase exponentially as repertoire size and antigenic space increase. At the same, time the mean abundance of any particular antigen repertoire decreases. This makes it difficult to visualise dominant sets of strains. Figure 5.1b illustrates this for just a slight increase in repertoire size ($R = 4$) and low diversity ($A_{max} = 32$). Clear structure nevertheless emerges in this scenario, as it did in the simpler scenario, with a dominant set of antigen repertoires accounting for over 90% of infections (figure 5.1c and d). With $R = 4$ and $A_{max} = 32$ discordant sets can contain up to 8 parasite genotypes, however in the particular instance shown in figures 5.1b-d there are 6. Given sufficient time it is likely that a full set of eight discordant strains would evolve, although this may be limited by the ability of the population to maintain sufficient diversity. The remaining parasite population consists of antigen repertoires which exist at much lower frequencies and appear to be largely recombinant antigen repertoires of dominant repertoires (the most abundant two 'minor' antigen repertoires and their overlaps are highlighted in red and orange in figure 5.1d).

In order to visualise repertoire size and global antigen diversity for studying *var* genes, different visualisation methods are required. Additionally, it is questionable whether recombinant repertoires, which only differ from their parent repertoire by a small degree (as seen in figure 5.1d), should be considered as separate strains or simply a variant of their parent strain. To overcome this problem, antigen repertoires were clustered using networks to group similar repertoires into strains. Henceforth 'strain' is used specifically to refer to sets of antigen repertoires which have been clustered together (see Methods). We note that the absolute number of strains is highly dependent on the repertoire similarity threshold, τ. Without a clear biologically motivated value for τ, the absolute number of strains should be treated as qualitative. However, this does not prevent meaningful investigation of the processes that affect the dynamics and maintenance of structured parasite populations.
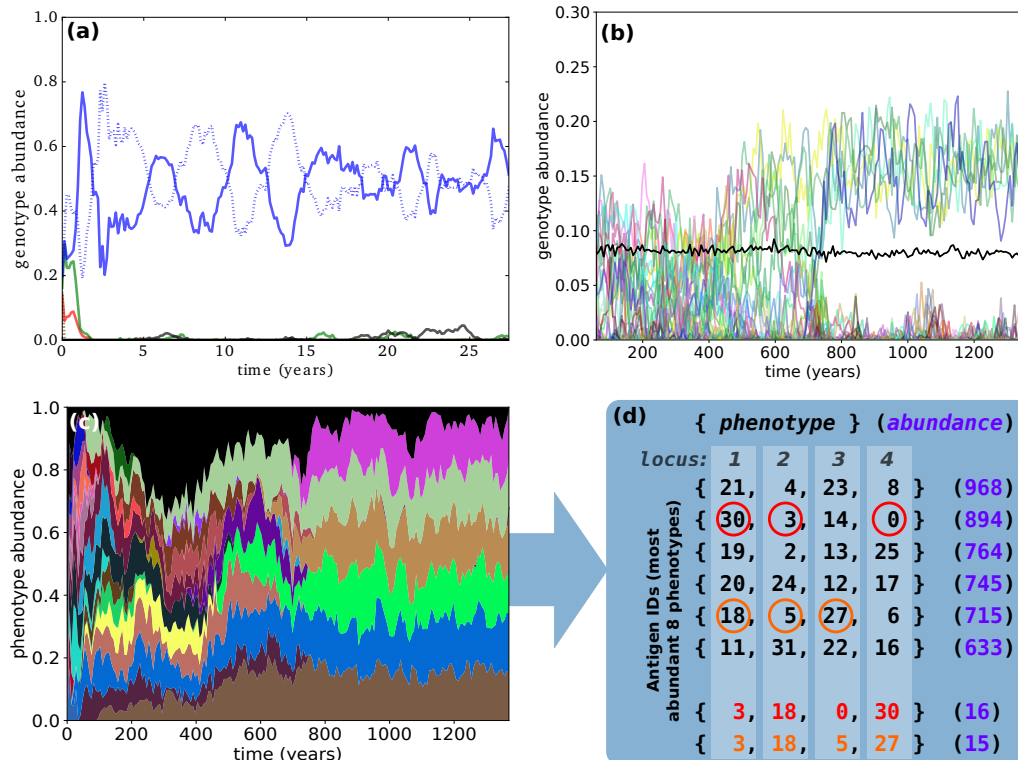
FIGURE 5.1: **(a)** Timeseries showing strain abundance for a single simulation when antigen space is restricted to four antigen variants and the parasite repertoire size is set to two ($R = 2$ and $A_{max} = A_{init} = 4$). There are six heterogeneous antigen repertoires that can be grouped into three discordant pairs (red, green and blue). Grey indicates sub-optimal homogeneous antigen repertoires (i.e those containing two of the same antigen type). One discordant pair dominates the infection dynamics and inhibits invasion by the other strains. Stochastic fluctuations determine which discordant pair dominates. For illustrative purposes the contribution each antigen makes to infection length was set to $\alpha = 90$ (in days). Other parameters: $H = 20000, M = 20000, b = 0.16, P_s = 0, P_c = 0.002$. **(b)** When $R = 4$ and $A_{max} = A_{init} = 32$ a larger dominant set of antigen repertoires persist, with theory predicting discordant sets of up to eight repertoires. Host prevalence is shown in black. **(c)** The composition of the parasite population over time. Specific antigen repertoires are visualised individually over time if their abundance exceeds 0.1 (all others are pooled into an 'other' category shown in black). A set of six repertoires eventually emerges to dominate the parasite population. The antigen repertoires and frequency of the eight most abundant strains are shown in **(d)**. Amongst the six dominant repertoire there are no overlaps. The next two most abundant repertoires exist at very low frequency and are likely to be recombinant progeny of two dominant repertoires (highlighted in red and orange). Other parameters used (for **(b - d)**: $H = 20000, M = 20000, b = 0.14, P_s = 5E - 4, P_c = 2.5E - 4, \alpha = 90$

.

5.3.2   The emergence of strain structure

To understand how strain structure emerges in our model, we estimated the number of strains over time and used two different metrics to quantify the degree of strain structure (figures 5.2 a-c). We found that the number of strains rapidly increases shortly after the burn-in period but consistently decreased before reaching an equilibrium. This decrease occurred despite the number of antigen variants in circulation remaining approximately constant (figure 5.2a). To understand this we need to consider the degree of strain structure, which we estimated using two measures of antigen overlap. In the first we calculated the mean number of antigens which overlap between strains (figure 5.2b), and in the second we took a simple average of the overlap between each antigen repertoire (figure 5.2c). In both cases, smaller antigen overlap is taken to be an indication of greater strain structure. An increase in overlap (both metrics) is observed at the same times as there is an increase in the number of strains. This indicates a low degree of strain structure during this period. The degree of overlap peaks shortly after diversification starts and drops to a new equilibrium as immune selection begins to shape the parasite population and strain structure emerges. During this process some strains are driven to extinction, which causes the drop in the number of strains seen in figure 5.2a.

For comparison, and to control for non-immune factors, we repeated this analysis without host immunity, in which case the parasite population is able to maintain a much higher level of diversity (figure 5.2d-f). This is because the parasite population becomes much larger and can therefore maintain a greater diversity. In the absence of immunity all antigen variants and antigen repertoires have equal fitness. The emergence of separate clusters (strains) occurs due to diffusion-like diversification
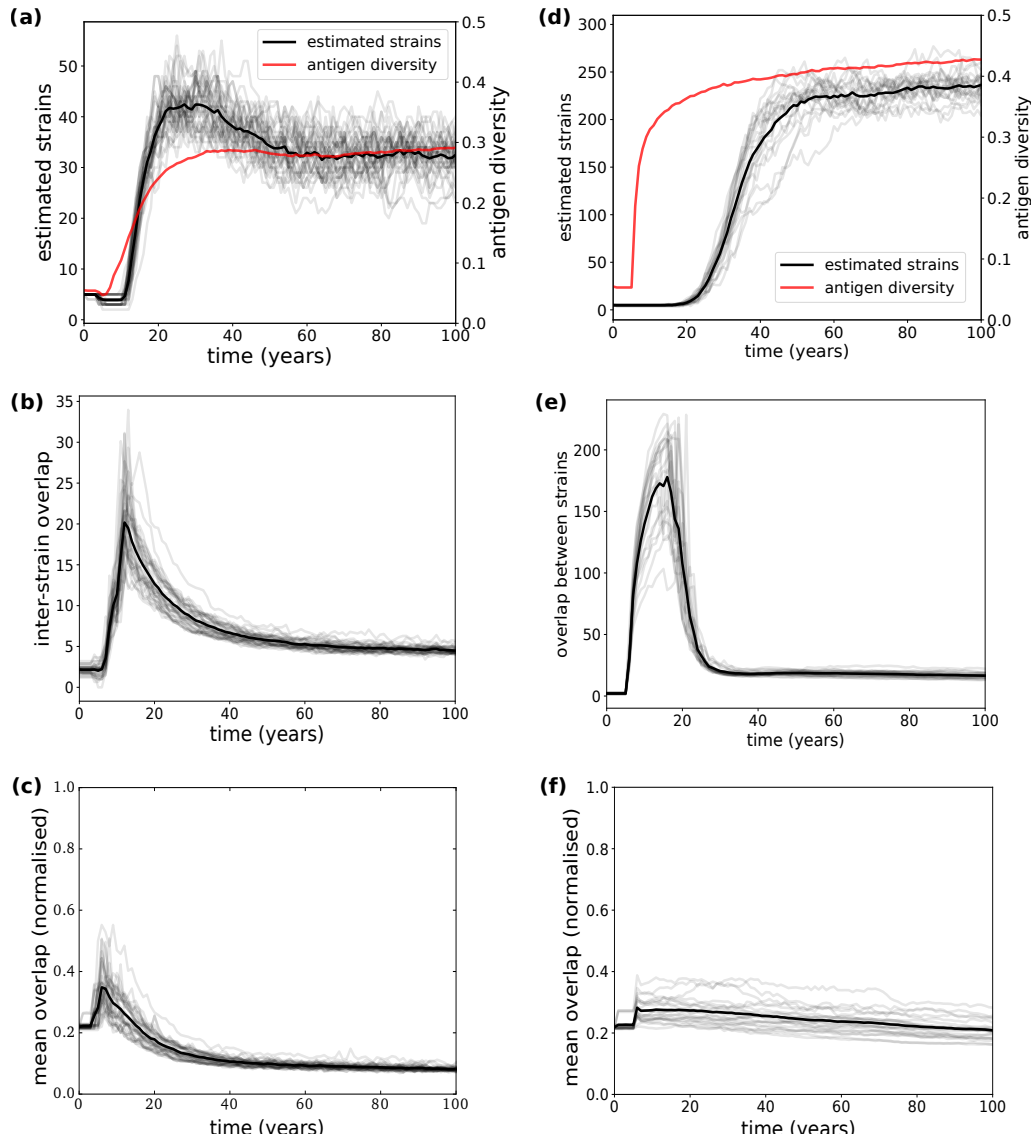
FIGURE 5.2: Changes in the parasite population structure over time with and without selection by host immunity. **(a, b and c)** With immune selection: The number of strains and mean inter-strain overlap initially increases over time as recombination generates new antigen repertoires and antigen variants. As the rate of diversification slows and immune selection acts to reduce the overlap between strains, the number of strains reduces until a stable population structure emerges. **(d, e and f)** Without immune selection: In the absence of immune selection all parasite strains are equally fit. Recombination means the parasite population diffuses out from the initial strains, and slowly utilise new regions of antigen space. Inter-strain overlap peaks before an there is appreciable increase in the number of strains. Parameters used: $M = H = 10000$, $b = 0.12$, $S_{init} = 5$, $A_{max} = 5000$, $P_c = 0.0008$, $P_s = 0.002$, $\tau = 50$. Black lines show the mean of 50 repeats. Light gray lines show timeseries of individual simulations.

spreading across the antigen space, in contrast to the directional selection experienced by parasites under immune selection. This results in the larger lag between the

number of circulating antigens and the number of strains (figure 5.2a). Additionally, the mean number of strains increases monotonically before reaching an equilibrium. The diffusion-like nature of neutral diversification results in a qualitatively different dynamic behaviour with regards to repertoire overlap. There is a more gradual peak in inter-strain overlap as the wave of diversification spreads across antigen space and before separate strains emerge (figure 5.2e). Furthermore, when separate strains do emerge, there is much higher mean overlap between strains since there is no immune selection pressure. The difference between the scenarios with and without immunity is more clearly seen from the mean antigen repertoire overlap (figure 5.2c and f).

The difference in the way the parasite population diversifies with and without immune selection is probably best illustrated using network diagrams generated by clustering antigen repertoires. Figures 5.3a and b show examples of such diagrams. Each antigen repertoire is represented by a node connected to similar repertoires. The network therefore represents a snapshot of the parasite population at a single point in time (25 years) for two simulations, one with immunity (5.3a) and one without (5.3b). Immune selection acts to facilitate the 'budding' of new strains by selecting against overlaps between antigen repertoires. Existing strains grow and eventually split to form two separate strains. In the absence of immune selection new antigen repertoires 'diffuse' through antigen space and remain connected to their parent repertoires for much longer, delaying the emergence of new strains.

One reason for computing both the inter-strain overlap and overlap between antigen repertoires was to determine which metric provides a clearer indication of the degree of strain structure. Although both metrics appear to be equally capable of indicating strain structure, calculating the overlap between antigen repertoires has several advantages. Firstly, it is much quicker to compute, and secondly, it is more

similar to the way antigenic diversity is measured and compared in the field. Additionally, the maximum possible overlap is simply given as R, such that repertoire overlap is easily normalised. As such we will continue the analysis using antigen repertoire overlap rather than inter-strain overlap.
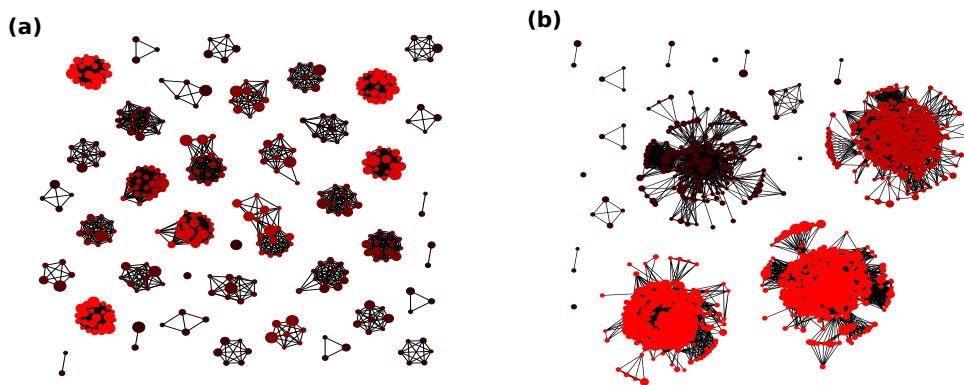
FIGURE 5.3: Network diagrams showing the strain structure of two parasite populations which have evolved for 25 years with **(a)** and without **(b)** immune selection. Each simulation was started with four initial strains ($S_{init} = 4$). Antigen repertoires are represented by nodes, with edges connecting antigenically similar nodes. Node colour indicates the abundance of the strain (cluster) as a whole (brighter red clusters are more abundant), while node size indicates the abundance of individual antigen repertoires. Parameters used: $M = H = 10000$, $b = 0.12$, $S_{init} = 5$, $A_{max} = 5000$, $P_c = 0.0008$, $P_s = 0.002$, $\tau = 50$.

5.3.3 Effect of recombination and transmission intensity

Next we explored the sensitivity of strain structure to changes in transmission intensity and recombination rate. As we expected, we find that high transmission intensity leads to better defined strain structure, shown in figure 5.4. This is because greater transmission intensity both increases the parasite population size (therefore increasing the capacity to retain diversity, see Chapter 3) and increases between-host immune selection pressure. Very high recombination rates were found to reduce the degree of strain structure when transmission intensity was low because recombination can easily disrupt advantageous associations between genes. However, for this to happen required recombination rates much higher than are biologically reasonable.
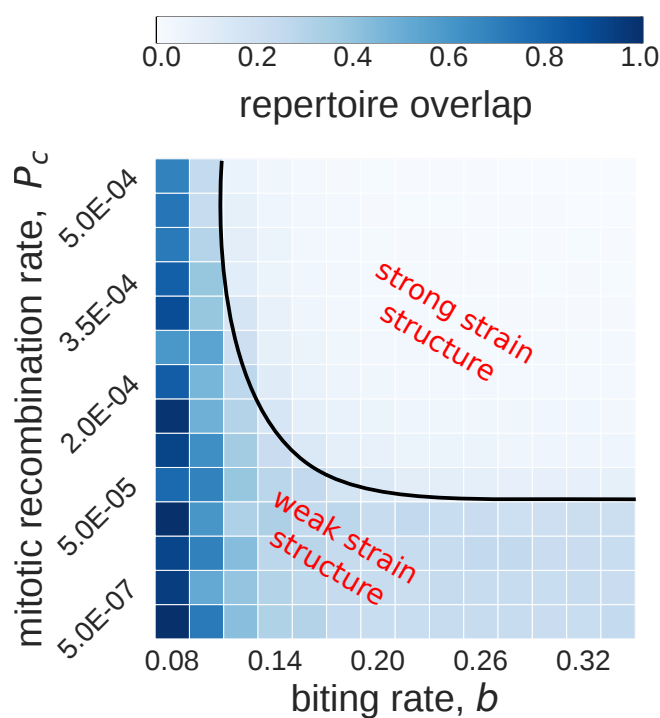
FIGURE 5.4: Repertoire overlap is used as a metric of the degree to which a parasite population is structured into strains. Greater transmission intensity increases the strength of between-host immune selection and strongly selects for non-overlapping antigen repertoires between strains. Greater mitotic recombination facilitates exploration of antigen space, allowing strains to more easily evolve to utilise different antigens to one another. Very low mitotic recombination slows this exploration, and prevents strong strain structure from emerging.

5.3.4   Stability of strain structure to temporal perturbations

In Chapter 3 we discussed the role of transmission-diversity feedback in causing hysteresis in the response of diversity and prevalence to changes in transmission intensity. We highlighted the potential importance this can have for assessing the efficacy of control measures, because the system's inertia makes it difficult to discern whether changes in prevalence are caused by recent changes in transmission, or whether the system is still equilibrating from a past fluctuation/s. Here we investigated whether changes in the strain structure could persist after a perturbation in transmission, and whether this could be used to indicate the recent history of transmission intensity in a region.

A similar protocol was used as in Chapter 3, in which an equilibrium is established before perturbing the system by increasing or decreasing transmission over a short number of years (in this case approximately 7 years). As expected, there is a temporal lag in the time taken for the number of strains to respond and equilibrate to the new transmission setting. This behaviour closely follows the previously described lag in total antigen diversity (figure 5.5a-d). Interestingly, when transmission intensifies there is a peak in the mean overlap and evidence of a relatively long period of disruption to the strain structure in the parasite population. When transmission intensity decreases, there was no sudden disruption to the strain structure. Instead the average overlap gradually increased to its new equilibrium point over a long period of time (in the same order of magnitude as host lifespan). This suggests that in some cases the parasite population structure could be used to indicate recent changes in transmission intensity. Specifically, high inter-strain overlap in the antigen repertoires of parasites may indicate a recent increase in historic transmission intensity to which the system has not yet equilibrated.
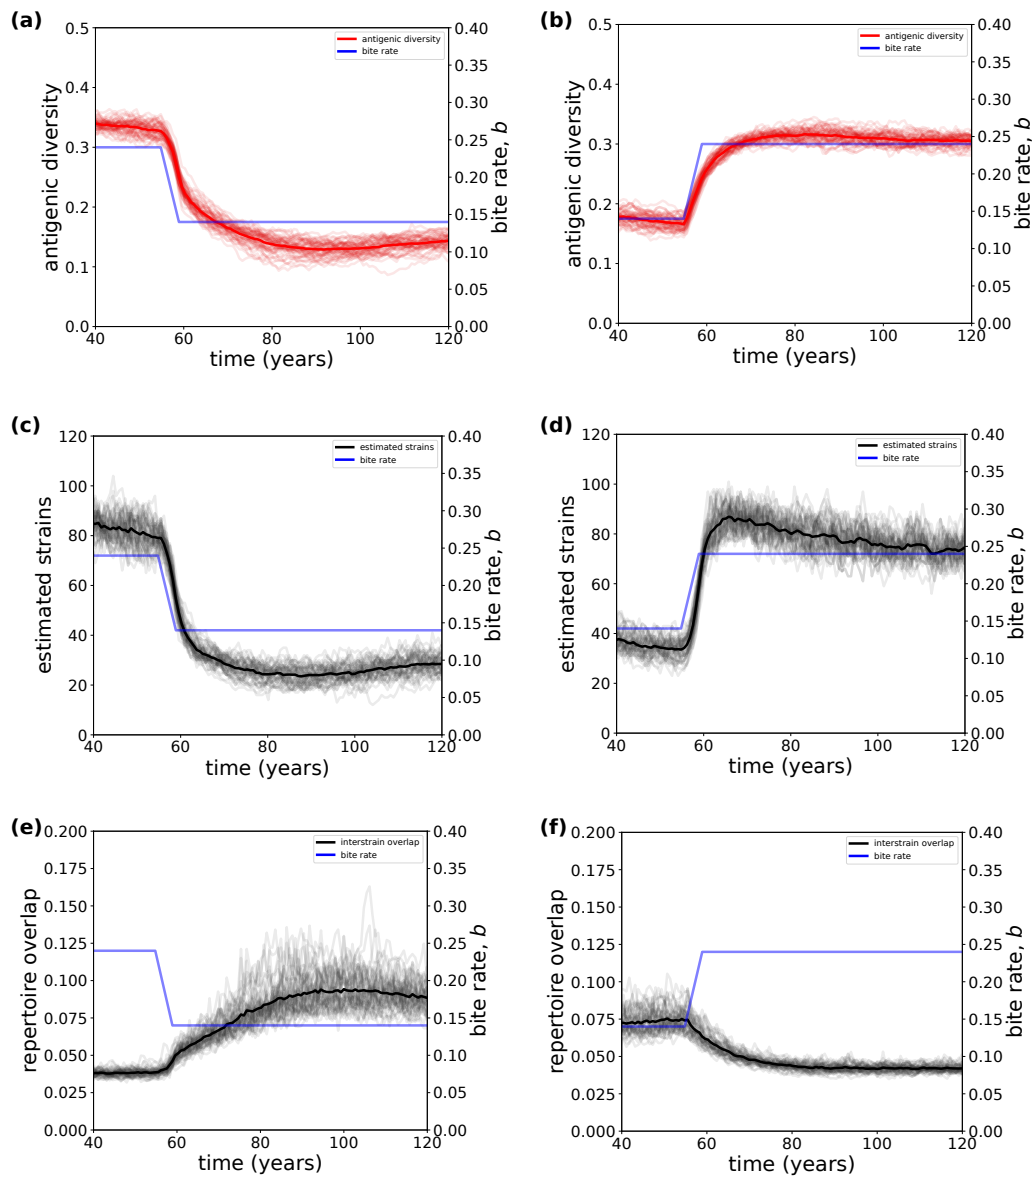
FIGURE 5.5: When transmission intensity is perturbed over a short period by decreasing **(a, c, e)** or increasing **(b, d, f)** the transmission intensity by modulating mosquito biting rate. A temporal lag is seen in the number of strains and mean repertoire overlap. The number of strains appears to closely match the change in diversity, but continues to equilibrate for some time after diversity has stopped changing (**(c)** and **(d)**). This is especially pronounced following an increase in transmission intensity. Sub-figures **(e)** and **(f)** show the mean repertoire overlap in the parasite population, and demonstrates that the strain structure is still changing long after the perturbation. All means are calculated from 50 repeats. Individual time series are shown in light grey (or red). Parameters as follows; $M = H = 10000$, $S_{init} = 40$, $A_{max} = 10000$, $P_c = 0.0008$, $P_s = 0.002$

5.3.5 The effect of seasonality on strain structure

Seasonal variation in the abundance and lifespan of vectors make annual variations in transmission intensity commonplace. In the extreme case, transmission can be interrupted for many months, and parasites rely on their ability to induce long infections in hosts in order to re-establish transmission when conditions become more favourable. This causes a bottleneck in the parasite population. Here we investigated the effect of this seasonal bottleneck on the evolutionary dynamics of parasite population diversity and structure.

Simulations were allowed to reach an equilibrium at the baseline bite rate ($b = 0.15$) without seasonal variation in order to provide a point for comparison. Mosquito biting rate (b) varied annually between $\pm 20\%$ of the baseline rate. As expected, the overall effect of seasonality is to reduce prevalence and reduce diversity (measured by either the total number circulating antigens or number of strains), shown in figure 5.6a and b. However, the strain structure of the parasite population remained stable despite the population bottleneck. While seasonality was associated with a decrease in the number of strains that the parasite population could support, the degree of overlap was not strongly affected by annual fluctuations (figure 5.6c and d).

5.3.6 Long-term temporal patterns

Few studies have considered the effect of temporal changes on the antigenic composition of *P. falciparum* populations. A notable exception is work by Nielsen et al (Nielsen et al., 2004), which compared the ability of contemporary sera and sera
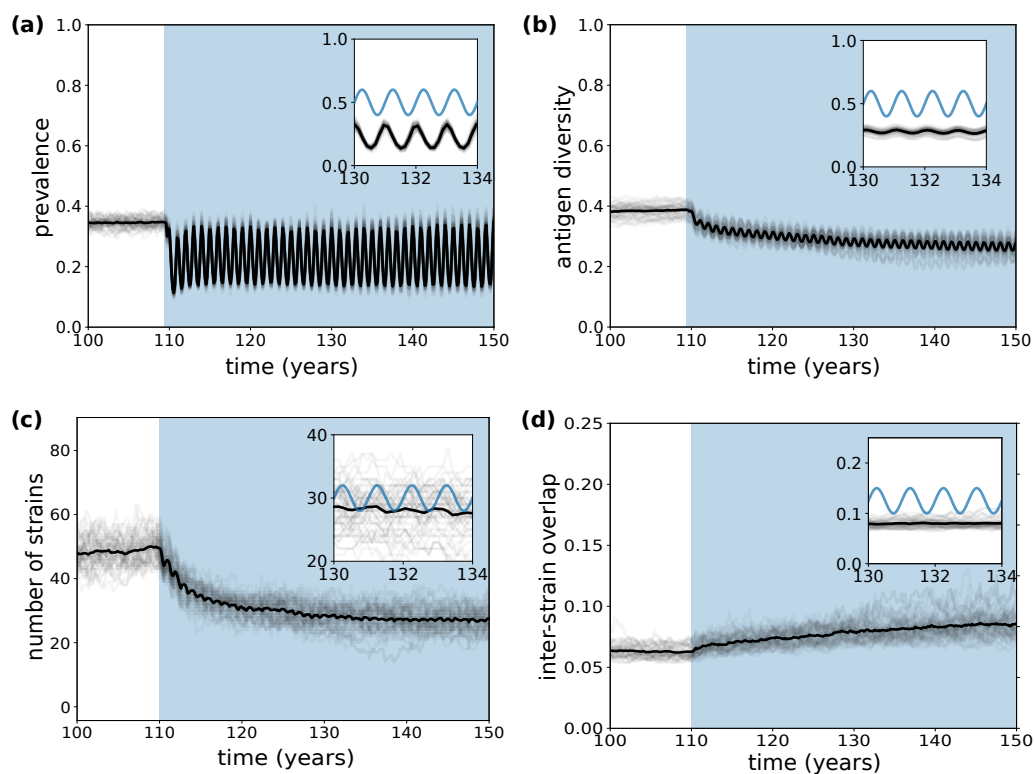
FIGURE 5.6: The effect of seasonality on **(a)** prevalence, **(b)** antigenic diversity, **(c)** number of strains **(d)** inter-strain overlap. Insets show a zoomed in section of the main figures, with light blue line showing seasonal change in mosquito biting rate. Seasonality was turned on (blue shaded area) after an equilibrium was reached without seasonality to provide a point for comparison. Black lines indicate the mean of 50 repeat simulations and individual simulations are shown by the faint grey lines. Seasonal variation in *b* fluctuated between 0.12 and 0.18.

taken ten years previously to recognise antigens expressed by current parasite populations. The authors found little difference between the ability to recognise contemporary antigens, suggesting that the parasite population was relatively stable during this period. More recent work used a similar approach, but analysed parasite samples from severe and non-severe infections separately (Warimwe et al., 2016). It was found that, while agglutination rates from contemporary sera was higher for samples from severe infections than non-severe infections, sera taken ten to twelve years previously showed no difference in the their agglutination scores for samples taken from severe and non-severe infections. This contrasts the Neilsen's findings,

and could suggest that antigens associated with severe disease do not exhibit temporal stability in parasite populations, but gradually change through recombination. To look for long-term population-level patterns in the antigens used by the parasite population we conducted virtual challenge experiments. The motivation behind this was to look for cyclic patterns or directional changes in the abundance of different antigens within the parasite population over time. This would allow us to determine whether the parasite population is stable over long periods of time, or whether new antigen variants are constantly emerging to replace older variants.

Simulations were ran for a duration which spanned several host lifespans using different sized antigenic spaces and transmission intensities. We did not observe cyclic patterns in the circulating antigens. However, the antigenic composition of the parasite population was not entirely stable over time. Predictably, the lowest parasite fitnesses were obtained when hosts were paired against parasite samples from the recent past, presumably because these hosts have a high probability of having been exposed to these strains. This is contrary to the results reported by Nielsen (Nielsen et al., 2004), which found parasite populations to be temporally stable. However, this work was conducted over a ten year period, and the change in the parasite population that we observe occurs rather more slowly (on the order of 50 to 100 years). The degree of change in the parasite population over a ten year period is thus very small and may be difficult to detect in the field.

Alternatively, temporal patterns in antigen prevalence may only be apparent for a sub-group of antigens (e.g. those associated with severe disease). While we did not test this explicitly, the overall result could equally apply to a sub-set of antigens as to the whole repertoire. It is possible that temporal instability is only apparent in a sub-group of less diverse antigens because population level immunity can reach higher levels with relatively few exposure events. However, our results suggest

that the prominence of long term dynamics depend strongly on the size of the antigenic space and capacity of the system to maintain diversity. Larger antigenic spaces favoured a more pronounced long term trend in parasite fitness. This meant that the antigen composition of the parasite population could continually evolve to exploit new regions of antigen space to evade acquired immunity. This can be seen in figure 5.7a, which shows the fitness of parasites paired against hosts from different points in time. When the antigen space was constrained to approximate a transmission-saturated state, temporal structuring was less apparent because the parasite population was more constrained in its ability to find novel antigens (figure 5.7b). This suggests another possible explanation for the discrepancy between Nielsen's study and our results in that the parasite population could already be effectively utilising all viable antigenic variants. If temporal patterns are more prominent in a sub-group of more conserved antigens then a different mechanism would need to be involved. This might involve, for example, the gradual breaking up of antigens through the exchange of domains due to mitotic recombination, as suggested by (Warimwe et al., 2016).
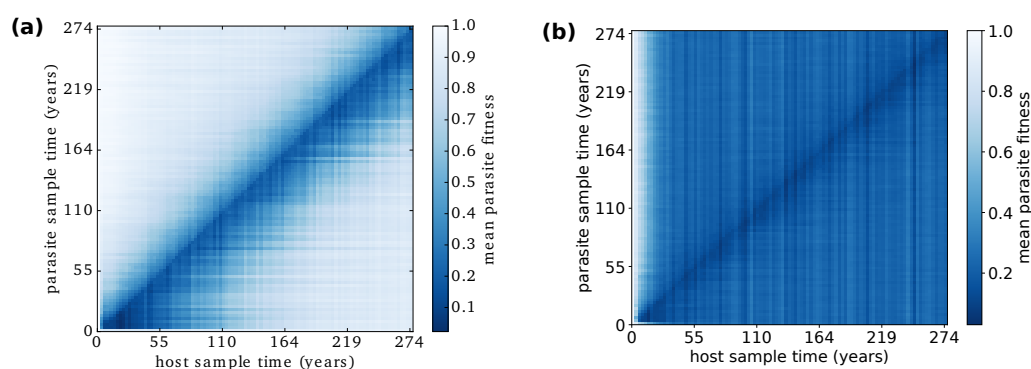
FIGURE 5.7: Virtual challenge experiments. Parasite fitness is lowest for contemporary and future hosts, suggesting the absence of a fixed and stable antigen composition in the parasite population. However the strength of this relationship was strongly dependent on the ability of the parasite to diversity and of the population to retain diversity. The effect of constraints on diversification can be seen by comparing **(a)** and **(b)** which used a largely unconstrained antigen space ($A_{max} = 50000$) and a strongly constrained antigen space ($A_{max} = 5000$), respectively.

5.3.7 Over-utilisation of antigens associated with severe-disease

In Chapter 4 we showed how a trade-off between within- and between-host immune selection can select for partitioned antigen repertoires. The ratio of $A : B$ antigens was sensitive to constraints on the global diversity, such that antigen repertoires which utilise larger numbers of the more diverse group were selected for. This resulted in the efficient use of the globally diversity and therefore minimised between host immune selection. We hypothesised that this same mechanism would select for antigen repertoires which over-utilise variants that might confer a within-host advantage. If this fitness advantage is sufficiently large, the disadvantage due to between-host immune selection could be overcome. This in turn would cause uneven abundances of antigens belonging to different phenotype groups, which may have consequences for population-level strain structure.

To test this we first simulated the dominance frequency of different antigen repertories using the methodology described previously in Chapter 4. Figure 5.8a shows the baseline case in which there are no differences between the innate fitness of A and B groups. Restricted antigen space in group A leads to the selection for antigen repertoires which utilise greater numbers of group B antigens. In figure 5.8b the innate fitness of phenotype group A is increased by doubling the contribution to infection length made by this group. This is sufficient to reverse the skew and favour repertoires which utilise larger numbers of group A antigens. We refer to this situation as group A being over-utilised, because antigen repertoires utilise a larger proportion of group A antigens than would be expected based on the global number of A and B antigens variants alone. This demonstrates the ability of innate fitness advantages to overcome the disadvantage associated with reduced diversity and to select for antigen repertoires which over-utilise one phenotype group.
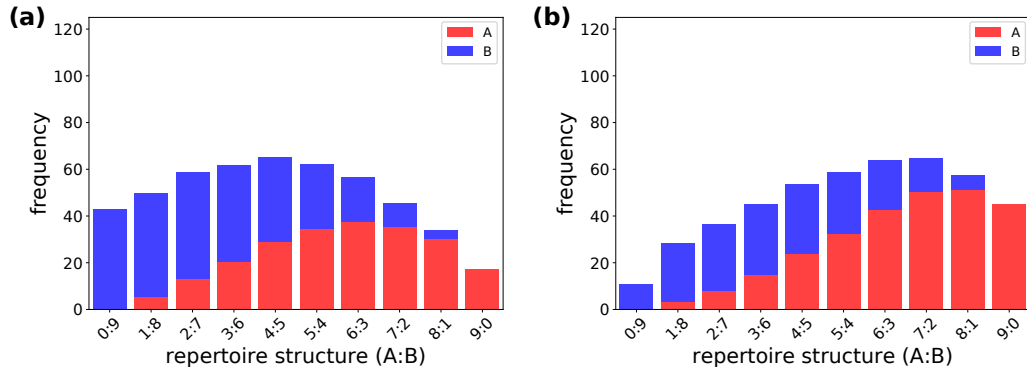
(a)

(b)

FIGURE 5.8: The distribution of dominant repertoires structure: **(a)** When phenotype group B is more diverse antigen the distribution is skewed toward greater utilisation of group B antigens. **(b)** When phenotype group A is associated with greater within-host fitness antigen repertoires which over-utilise group A are selected for, despite the greater degree of diversity in group B. In this case phenotype group *A* antigens contributed twice as many infection-days to an infection as phenotype group *B* antigens. Other parameters used: $M = H = 10000$, $S_{init} = 40$, $A_{max,A} = A_{init,A} = 9$, $A_{max,B} = A_{init,B} = 27$, $S_{init} = 10$, $R = 9$, $P_c = 0.0$, $P_s = 0.0$, $b = 0.16$.

To understand the effect of over-utilisation on strain structure we fixed the repertoire structure such that each antigen repertoire was comprised of 30% group *A* antigens and 70% group *B* antigens (18:42). We then tested the effect of over-utilisation by changing the partitioning of antigenic space (different $N_A : N_B$ ratios). An antigen group would thus be over-utilised if its proportion in antigen repertoires was larger than the proportion of antigenic space (i.e. for group *A*: $\frac{n_A}{n_A+n_B} > \frac{N_A}{N_A+N_B}$). The total size of antigen space was kept constant. As expected, the maximum number of strains and the lowest repertoire overlap occurs when the ratio of *A* and *B* antigens in antigen space corresponds to the proportions of the fixed repertoire structure (figure 5.9). This is because, as shown analytically in (Recker, Arinaminpathy, and Buckee, 2008), equal proportions lead to the most efficient utilisation of global diversity. As the partitioning of antigenic space deviates from that of the antigen repertoires, antigen repertoires over-utilise one or the other phenotype group. This results in a decrease in the number of strains (figure 5.9a) and an increase in repertoire overlap (figure 5.9b). Unsurprisingly, when we calculated the degree overlap

in *A* and *B* groups separately there was greater overlap in the more constrained phenotype groups (red and blue lines in figure 5.9b), indicating that the degree of strain structuring was lower in some parts of the repertoire than others.
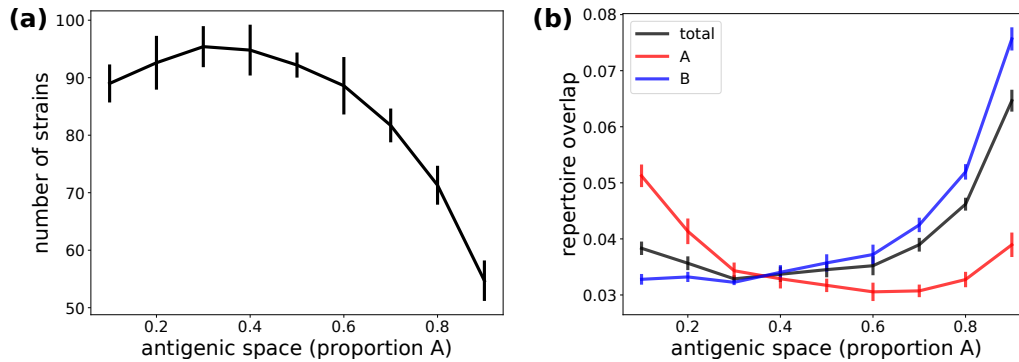

FIGURE 5.9: The effect of different antigen space partitioning ratios on the strain structure of parasite populations with fixed 18 : 42 partitioned antigen repertoires structure. **(a)** Number of strains peaks when the partitioning of antigen space matches the partitioning of antigen repertoires. **(b)** Mean repertoire overlap is lowest when the partitioning of antigen space matches the partitioning of antigen repertoires. Values were calculated from 20 repeats. Runs in which the parasite population went extinct were removed before calculating means. Error bars show the 95% confidence interval. Parameters used: $M = H = 10000$, $S_{init} = 40$, $A_{max} = 12000$, $P_c = 0.0008$, $P_s = 0.002$.

5.3.8   Age-distribution of antigen presentation

There is a strong correlation between age and severity of infection in endemic regions, with hosts acquiring immunity against severe disease relatively quickly. It has been hypothesised that severe disease may be associated with the expression of a restricted subset of *var* genes (Kaestli et al., 2006; Kyriacou et al., 2006; Jensen et al., 2004; Bull et al., 2000; Gupta et al., 1999; Kraemer and Smith, 2006). We investigated whether over-utilisation of antigens associated with severe disease, which was found to lead to different degrees of strain structure with respect to different phenotype groups, could explain this age-distribution in the acquisition of immunity. To do this we calculated the age-distributions for two metrics: host susceptibility and the infection phenotype. Host susceptibility is defined previously (Chapter 3).

Infection phenotype is defined as the average contribution to infection length that each phenotype group makes in hosts of a given age. It is determined by summing the contribution to infection (in days) attributed to each phenotype group and then calculating the mean across all infections, stratified by host age. These means are normalised by the maximum duration of infection across all age groups. This metric can be used to determine the important antigen group/s in host infections for different ages groups. Under the assumption that group A antigens are more likely to cause severe disease, the degree to which an infection is dominated by phenotype group A can then be used as a proxy for risk of severe disease.

We ran simulations using the baseline assumption that the global antigen space was partitioned in the same proportions as the antigen repertoires of individual parasites (i.e. a 3 : 7 ratio of A and B groups). The age distributions of susceptibility and infection phenotype were calculated at equilibrium and shown in figure 5.10. Hosts gained immunity to both phenotype groups at equal rates, as indicated by the decrease in susceptibility (figure 5.10a). Similarly, while the average infection length decreases as hosts gain immunity through exposure, indicated by a downward trend in figure 5.10b, the proportion that A and B phenotypes contribute to infection remains approximately constant, and hence there is no change in infection phenotype with age.

Next, we explored the effect that over-utilisation of severe disease causing antigens would have on the acquisition of immunity by setting the $A : B$ ratio of antigen space to be smaller than that of the fixed antigen repertoire partitioning structure (i.e. 1 : 9 versus 3 : 7). In contrast to the baseline scenario, over-utilisation means that hosts gain immunity to group A more quickly than they do to group B (figure 5.11a). This is because hosts are exposed to a greater proportion of the possible group A diversity with each infection event. Different rates of acquisition of immunity for
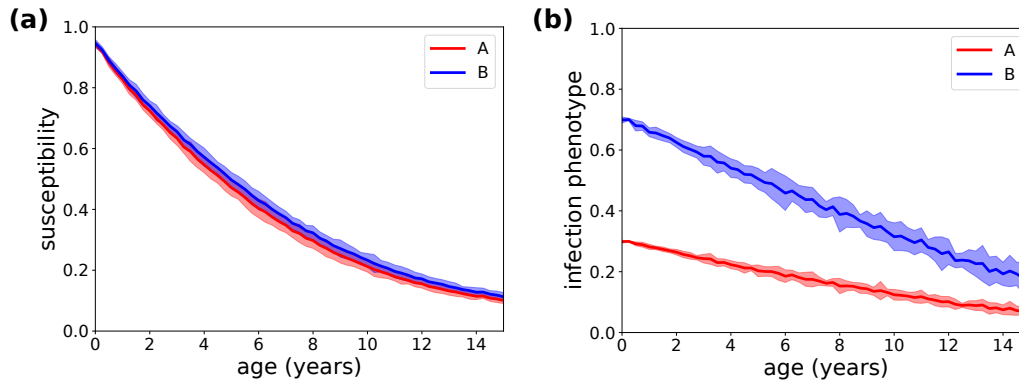
FIGURE 5.10: Age distribution of host **(a)** susceptibility and **(b)** infection phenotype, at equilibrium using baseline assumption that antigenic space is partitioned in the same $A : B$ ratio (18 : 42) as the parasite repertoire structure. Age distributions were produced using ten repeat simulations which were ran to equilibrium and the output of which was aggregated to produce a single mean age distribution for host susceptibility and infection contribution. Bold lines show the mean of 10 repeats and the shaded region indicates 95% confidence intervals. Other parameters used: $H = M = 10000$, $A_{max} = 10000$, $S_{init} = 25$, $P_c = 0.0008$, $P_s = 0.002$, $N_A = 18$, $N_B = 42$, $b = 0.16$

different phenotype groups leads to changes in infection phenotypes as hosts get older under continuous exposure (figure 5.11b), with group A antigens contributing less to infections in older hosts.

Over-utilisation of phenotype group A means that the system is in a more transmission saturated (or diversity limited) state with respect to phenotype group A than phenotype group B. This is an important factor, which prevents the parasite population from generating greater diversity in phenotype group A in response to immune selection. When the size of the antigenic space as a whole ($A_{max} = 1500$) is reduced (maintaining the same partitioning proportion between A and B phenotypes), the system moves toward a more strongly transmission saturated state and the age-dependent change in infection phenotype is more pronounced (figures 5.11c and d). Similarly, increasing the size of the antigen space ($A_{max} = 100000$) allows unrestricted diversification in response to immune selection and the parasite population generates more diversity to mitigate the over-utilisation of group A (figures 5.11e and f).

While we showed that selection can lead to over-utilisation of one phenotype group over the other when there are differences in innate fitness between the phenotype groups (figure 5.8), we do not consider differences in innate fitness for the rest of the analysis (including the simulations used to produce figure 5.11). Therefore, these age-dependent patterns in infection phenotype are generated solely from differences in the utilisation antigenic diversity within different parts of the antigen repertoires.

5.3.8.1   Transcription hierarchies

Next we investigated the effect of expression order on the age distribution of susceptibility and infection phenotype by making two assumptions. Firstly, we assume that phenotype group A antigens are preferentially expressed in hosts, for example because they have a within-host selective advantage and hence parasites that switch to group A antigens predominate. Although switching networks appear to be highly structured and variable between *var* genes (Horrocks et al., 2004; Recker et al., 2011), the relatively high rate with which merozoites switch *var* gene expression between generations means that it is expected that following release from heptocytes merozoites will initially express all of their *var* repertoires to some extent. *In vivo* studies support this, but have further shown that after just a few generations of merozoites selection acts to strongly shape expression, and that in naive individuals this predominantly involved the expression of group A and B genes (Wang et al., 2009). For the purposed of this analysis we describe this selective expression as a 'transcription hierarchy'. Transcription hierarchies were implemented such that A group antigens were expressed before B group antigens provided that the host was not already immune to them. This differs from our previous assumption that antigen groups that are associated with a within-host fitness advantage are over-utilised by parasites.

Secondly, in order for this transcription hierarchy to have any effect it is necessary to relax our previous assumption that parasites express their entire antigen repertoires during an infection. If parasites express their entire repertoire in every infection then it is impossible for expression order to affect our model. This is because in our simplified within-host model of infection the duration of infection is calculated as an aggregate of the contributions made by individual antigens. Allowing for the possibility that infections are sometimes cleared before the whole antigen repertoire has been expressed means that the order of expression becomes an important consideration. This was implemented by introducing a probabilistic element to the acquisition of immunity and clearance of an infection whereby there is a small fixed probability (P_{clear}) of the infection ending after the expression of each antigen variant.

Figure 5.12 compares the effect of hierarchical transcription with random transcription when infection clearance is probabilistic. When transcription order is random there is little qualitative change from the baseline scenario (with full repertoire expression). If infections can be cleared before parasites have the chance to express their entire antigen repertoire, the average number of antigens that hosts are exposed to with each infection is lower, and immunity is acquired more slowly than in the baseline case. It is important to note that host susceptibility decreases at the same rate for each phenotype group (figure 5.12a), and the infection phenotype remains the same for hosts of all ages. In other words, phenotype *A* and *B* contribute to infections in approximately constant proportions, despite an overall reduction in infection length (figure 5.12b). When a transcription hierarchy is used such that parasites preferentially express phenotype *A* antigens, the rate at which hosts gain immunity diverges. This can be seen in the rapid decrease in susceptibility to phenotype *A* antigens as compared to phenotype *B* antigens (figure 5.12c). This directly leads to an exposure-dependent (and hence age-dependent) change in infection phenotype,

with older hosts experiencing infections in which fewer antigens of phenotype A are expressed (figure 5.12d). Under the assumption that phenotype group A is associated with severe disease, this could explain the observed age-dependent pattern in morbidity.

We also tested an additional mechanism in which parasites express a fixed proportion of their antigen repertoires during an infection. This has an almost identical effect to probabilistic immunity, shown in figure 5.13. As parasites express fewer antigens the distinction between infection phenotype in young and older children increases (compare figure 5.13a and b, in which parasites express a half verses a quarter of their repertoires, respectively). Similarly, as the probability of early parasite clearance (P_{clear}) increases, the mean number of antigens that are expressed during an infection decreases, and the distinction between infection phenotypes in young and older children becomes more pronounced (figure 5.13c and d).
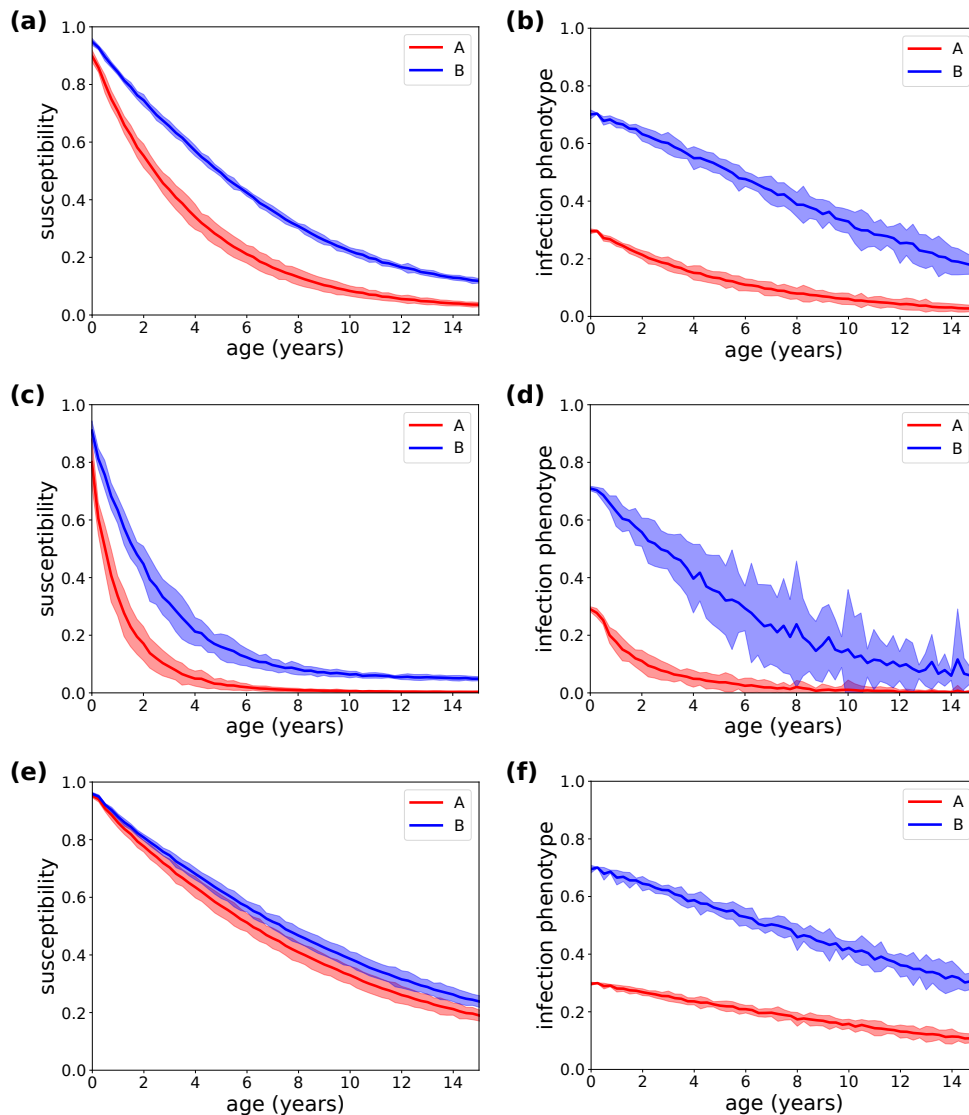
FIGURE 5.11: The age distribution of host susceptibility and infection phenotype with respect to each antigen group. Parasites over-utilise phenotype group A. i.e. the $A : B$ partitioning ratio of antigenic space is $1 : 9$ whereas antigen repertoire partitioning structure ratio is fixed at $3 : 7$ (A:B: repertoire structure of $18 : 42$). Sub-figures demonstrate the effect of transmission saturation. **(a)** and **(b)** $A_{max} = 10000$ (partial transmission saturation). **(c)** and **(d)** $A_{max} = 1500$ (full transmission saturation). **(e)** and **(f)** $A_{max} = 100000$ (little transmission saturation). Bold lines show the mean of 10 repeats and 95% confidence intervals are indicated by the shaded redions. Other parameters used: $M = H = 10000$, $S_{init} = 25$, $b = 0.16$, $P_c = 0.0008$, $P_s = 0.002$, $N_A = 18$, $N_B = 42$.
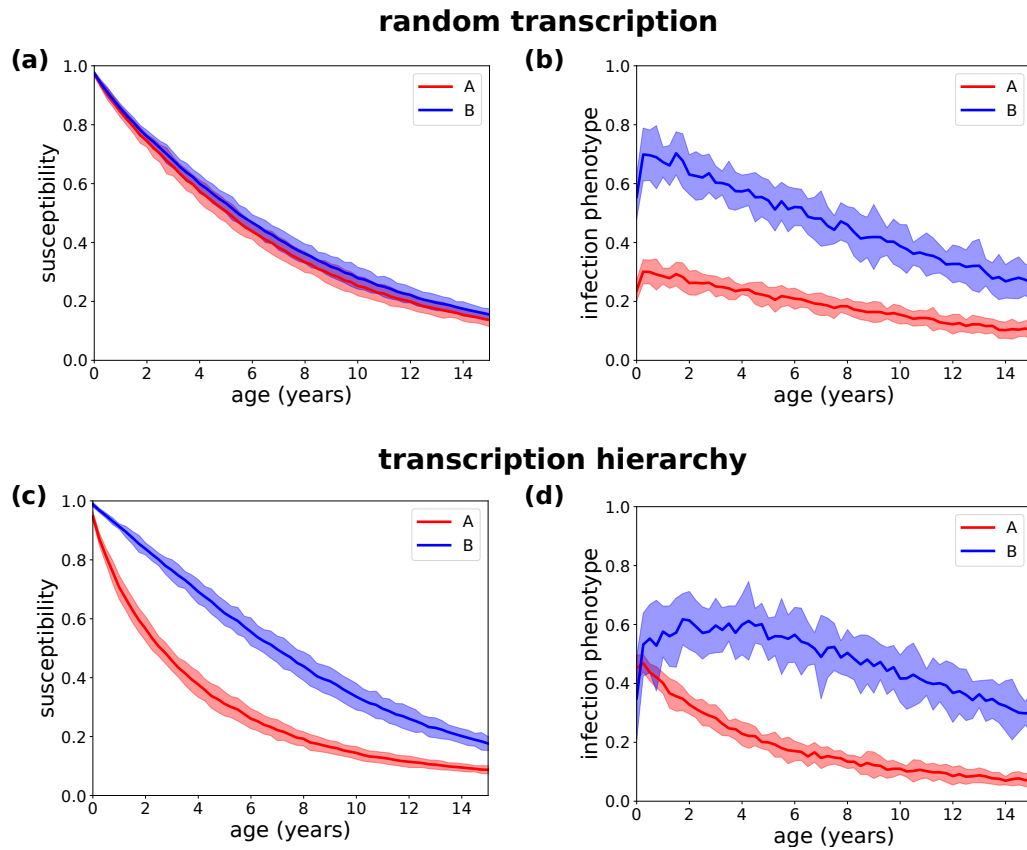
FIGURE 5.12: Susceptibility and infection phenotype distributions for hosts of different ages when there is a probability that infections are cleared before all antigen variants have been expressed. **(a)** and **(b)** Probabilistic parasite clearance does not in itself lead to a qualitative change in either host susceptibility or infection phenotype with age. However, hosts do acquire immunity more slowly because they are exposed to fewer antigens per infection. When transcription hierarchies are used: **(c)** hosts gain immunity to phenotype groups at different rates and, **(d)** the infection phenotype changes with host age. Bold lines show the mean of 10 repeats and the shaded region indicates the 95% confidence interval. Parameters used: $M = H = 10000$, $S_{init} = 40$, $A_{max,A} = 3000$, $A_{max,B} = 7000$, $P_c = 0.0008$, $P_s = 0.002$, $N_A = 18$, $N_B = 42$, $P_{clear} = 0.05$.
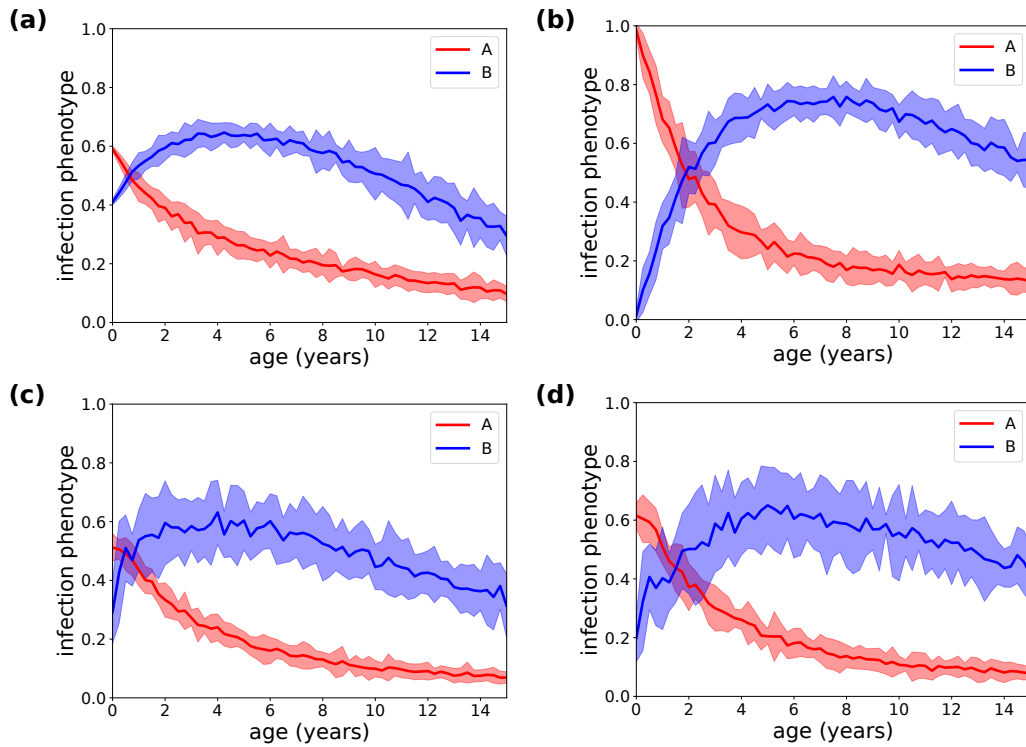
FIGURE 5.13: Partial expression of parasite repertoires in conjunction with a transcription hierarchy leads to changing infection phenotypes with age. **(a)** Parasites randomly choose half of their antigen repertoires to express. **(b)** Parasites randomly choose a quarter of their antigen repertoires to express. Increasing the probability of early parasite clearance (p_{clear}) has a similar effect: **(c)** On average approximately 15 antigens are expressed ($p_{clear} = 0.067$). **(d)** On average approximately 10 antigens expressed ($p_{clear} = 0.1$). Note that the difference in the magnitude is because each infection contribution is normalised by the maximum possible infection contribution which is different for probabilistic clearance (equal to R) and expression of a fixed number of antigens. Lines show the mean of 10 repeats and shaded region shows indicates the 95% confidence interval. Other parameters: $M = H = 10000$, $S_{init} = 40$, $A_{max,A} = 3000$, $A_{max,B} = 7000$, $P_c = 0.0008$, $P_s = 0.002$, $N_A = 18$, $N_B = 42$.

5.3.8.2   Transmission intensity

It has been observed that the age distribution in disease severity is strongly influenced by transmission intensity (Snow et al., 1997). To investigate the effect of transmission intensity in our model, we simulated age distributions of infection phenotypes under different transmission settings.

Under static diversity we find a strong transmission intensity dependence on infection phenotype with age (figure 5.14a and b). Increasing transmission intensity increases the rate of exposure and hosts develop immunity faster. The expression of large proportions of phenotype group A is therefore restricted to younger age groups. The opposite occurs for lower transmission intensity because it takes longer for hosts to become exposed to the majority of phenotype group A antigens.

We tested two scenarios in which diversity was dynamically generated through mitotic recombination. In the first, antigenic space was large ($A_{max} = 10000$) and led to a diversity saturated equilibrium state. Under these conditions no change in the distribution of infection phenotypes was observed for different transmission intensities (figure 5.14c and d). This is because in a diversity saturated setting, the amount of diversity in the system scaled with transmission rates and parasite population size, as discussed previously in Chapter 3. Hence, the parasite population was able to maintain greater diversity under higher transmission intensity, and hosts did not develop immunity more quickly despite being exposed to antigens at a greater average rate. In the second dynamic diversity scenario, the antigenic space was much more restricted ($A_{max} = 1500$), which limited the amount of novel diversity which could be generated and resulted in a transmission saturated equilibrium state at higher transmission intensity (figure 5.14e and f). Under transmission saturated conditions, the age-distribution of infection phenotypes was qualitatively similar to

those of the static diversity scenario and highly dependent on the exposure rate.
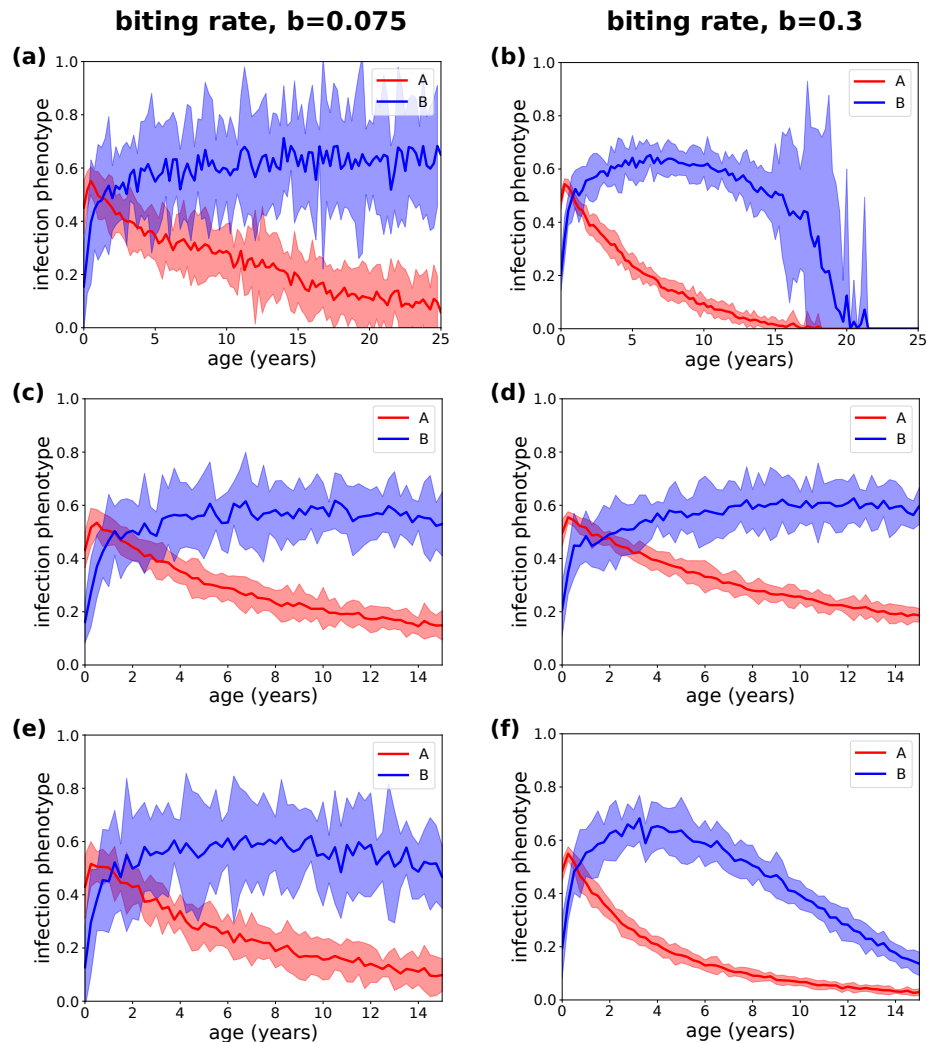
FIGURE 5.14: Mean age-distribution of infection phenotype under a transmission saturated and diversity saturated settings. The left and right-hand columns show the mean age-distributions for a low ($b = 0.075$) and high ($b = 0.3$) transmission settings respectively. **(a)** and **(b)** Static diversity (no recombination) is used in figures and represents a transmission saturated setting. **(c)** and **(d)** When novel diversity can be generated through mitotic recombination in a largely unrestricted fashion ($A_{max} = 10000$) the equilibrium state is one of diversity saturation even at high transmission rates. **(e)** and **(f)** When diversification is restricted to a smaller antigenic space ($A_{max} = 1500$) then transmission saturation is re-established at higher transmission intensity. Each subplot shows the mean of ten repeat simulations. Shaded region shows the 95% confidence interval. The default infection contribution was raised from the default value of 3 to 7 for all simulations. This was necessary to maintain viable transmission in the static diversity scenario while allowing comparison over the same range of transmission intensity as the dynamic diversity scenarios. Other parameters: $M = H = 20000$, $S_{init} = 25$, $N_A = 18$, $N_B = 42$, and where used recombination rates $P_c = 0.0008$, $P_s = 0.002$.

5.4   Discussion

P. falciparum has evolved sophisticated immune evasion strategies, which utilise diversity at multiple levels. At the population level, parasites are organised into strains which minimise immune interference to one another, while at the level of individual infections parasites escape immune recognition by switching expression between different *var* gene variants. As a result, hosts develop immunity slowly through the gradual exposure to a broad range of circulating antigens. One of the most striking epidemiological characteristics of *P. falciparum* malaria is the strong age-dependent pattern in disease severity, with the greatest burden of disease often falling on young children. This change in pathology has been linked to the frequency with which different *var* subgroups are expressed in hosts of different ages and has led to the hypothesis that different cytoadhesion phenotypes which may have evolved to exploit different host niches (Kraemer and Smith, 2003).

var genes can be categorised into different groups based on upstream promoter regions, chromosomal location, transcription direction and number of binding domains. Of these groupings, the UpsA group has been linked to increased risk of severe disease (Kaestli et al., 2006; Kyriacou et al., 2006), being more frequently expressed in individuals with severe disease and in young children. In order to elucidate possible mechanisms which could lead to the changes in infection phenotype (or infection severity), we explored the effect of different diversity constraints on the strain structure and transmission of parasite populations with partitioned repertoires. We further considered the possible role of transmission hierarchies in explaining these patterns.

We first established metrics to quantify the degree of strain structure in our simulated parasite populations and characterised the effect of global antigen constraints

on strain structure. In line with previous work by others (Recker, Arinaminpathy, and Buckee, 2008), we found that over-utilisation of one phenotype group reduced the ability of parasites to fully utilise the global diversity. Both the number of co-existing strains and the mean overlap between strains increased as the $A : B$ partitioning of antigen space deviated from the $A : B$ partitioning of parasite repertoires. We showed that antigen repertoires, which over-utilise one antigen group, would be selected for if there is an innate within-host fitness advantage associated with that group, for example due to differences in binding affinity and/or growth rate.

Next, we explored the conditions under which over-utilisation of one phenotype group could explain observed patterns in the age-distributed acquisition of immunity. We demonstrated that over-representation of one phenotype group can occur when the group confers a within-host fitness advantage. We found that over-utilised groups were more transmission saturated. Transmission saturation prevents diversification in response to increased immune selection. Furthermore, this led to unequal between-host immune selection and different degrees of strain structure with respect to the different phenotype groups and this caused immunity to develop more rapidly to the over-utilised group. This in turn resulted in changes to infection phenotype with host age, in which phenotype group A antigens were expressed to a lesser extent in older individuals. Interestingly, this effect was only apparent when the system was transmission saturated with respect to group A, i.e. when group A antigens experience greater between-host immune selection than group B. This agrees well with other work which suggests that different *var* groups or sub-groups might be under different degrees of immune selection (Buckee, Bull, and Gupta, 2009). While we only used two antigen groups in this study, we see no reason these results wouldn't generalise to three or more antigen groups. For example, cerebral malaria is known to peak later than other forms of severe disease. Under this

framework we might expect that antigen variants which are associated with cerebral malaria exhibit lower fitness than those associated with other forms of severe malaria, but greater fitness than those expressed in non-severe infections. More generally, the age distribution of different clinical presentations may reflect the relative fitness of the antigen variants which cause them.

We also considered transcription hierarchies which favoured the expression of group *A* antigens during an infection and investigated the effect this would have on the acquisition of immunity. This led to hosts developing immunity to group *A* antigens more rapidly, with immunity to phenotype *B* being acquired considerably slower. Furthermore, there was a strong correlation between exposure / age and infection phenotype, with infections in young hosts more likely to involve the expression of larger numbers of phenotype *A* antigens.

Finally we investigated the effect of different transmission intensities on the emergence of age-dependent infection phenotypes. Under transmission saturated settings, our model agreed qualitatively with field studies which show that higher transmission intensity is associated with more rapid acquisition of immunity to severe disease, as well as a more exaggerated change in infection phenotype with age. However, under diversity saturated conditions increases in transmission intensity and prevalence had little effect because the parasite population could mitigate increases in selection pressure by diversification. Given that transmission intensity is a strong factor determining the rate at which hosts develop immunity to severe disease (Reyburn et al., 2005; Snow et al., 1997), our results suggest that *P. falciparum* epidemiology is often transmission saturated. This is most likely to occur in regions with higher transmission, but considerable variation is likely due to ecological factors, such as host population size. Identifying whether a region is transmission or diversity saturated is of potential importance for understanding the way in which

its epidemiology is expected to respond to natural or control-induced perturbations in transmission rate (see Chapter 4). This could have important considerations for understanding how control will impact on the burden of severe disease.

Furthermore, the fact that age-dependent infection phenotype could not be reproduced in diversity saturated settings reinforces that, while vast, PfEMP1 diversity might not be unconstrained. Further work should focus on identifying structural associations between *var* genes and *var* groups, rather than simply comparing sequence similarity. In particular, more work is needed to understand the associations at the domain level and their implications for binding phenotype and serotype.

Chapter 6

Summary

Throughout this thesis we argued that understanding antigenic diversity, and critically the structure contained within this diversity, is centrally important to understanding *P. falciparum* epidemiology. We focused on the diversity exhibited by the PfEMP1 encoding *var* multi-gene family because of its importance in naturally acquired immunity, immune evasion, and as a determinant of virulence. This work undertook two primary avenues of investigation. Firstly, we sought to understand the implications of considering parasite diversity as a dynamic property. We identified ecological and epidemiological parameters which are important determinants of diversity and strain structure and investigated the effect of the feedback between diversity and transmission on disease prevalence. Secondly, we aimed to understand the evolutionary forces which maintain structured diversity at different scales and how different constraints on diversity affect parasite fitness and host immunity.

Each parasite genome encodes multiple variants of PfEMP1 that are expressed in a mutually exclusive fashion in a process referred to as clonal antigenic variation. This helps the parasite to evade immunity and prolong infection. Furthermore, there are a large number of distinct *var* gene variants circulating in natural *P. falciparum* populations (Barry et al., 2007; Chen et al., 2011), which implies that hosts probably encounter new antigen variants throughout their lives. While hosts never develop

full sterile immunity, anti-disease immunity is acquired relatively quickly through repeated exposure. It has been hypothesised that this naturally acquired immunity which changes malaria infection pathology from severe and life-threatening disease in young hosts to asymptomatic, chronic infections in other individuals, arises through the piecemeal acquisition of a broad collection of protective antibodies to the circulating antigenic diversity in parasite populations (Newbold et al., 1992; Bull et al., 1998).

Different PfEMP1 variants exhibit different cytoadhesion phenotypes, causing infected erythrocytes to sequester in different host tissues. This is an important determinant of virulence and is key to understanding the range of pathologies associated with malaria (Kyes, Kraemer, and Smith, 2007; Kraemer and Smith, 2006). One of the most striking and tragic aspects of *P. falciparum* epidemiology is that severe and life-threatening disease is predominantly experienced by young children. The risk of severe disease decreases with exposure, and the age-distributed risk of disease is therefore thought to be strongly dependent on transmission intensity (Reyburn et al., 2005).

The fact that protection against severe disease is acquired relatively quickly, and that some antigen variants appear to be recognised more commonly than others (Bull et al., 1999; Bull et al., 2000), has been used as evidence that severe disease is associated with the expression of a restricted sub-set of *var* genes. It has been hypothesised that different *var* gene groups are evolving in response to different selection pressures and may be specialised to particular niches depending on the immune status of the host (Kraemer et al., 2007). For example, the expression of UpsA and UpsB *var* genes is associated with an increased risk of severe disease (Kaestli et al., 2006; Kyriacou et al., 2006). This has led to the suggestion that this *var* group may include antigens which are adapted to relatively naive hosts, while other *var* groups

may be specialised to semi-immune hosts.

The host immune system is thought to play a role in determining the specific antigen variants which are expressed during an infection, with long-lasting but predominantly variant-specific immunity preventing viable expression of antigens to which the host has previously been exposed to. Furthermore, immune selection is thought to shape parasite populations such that they exhibit strain structure with respect to their *var* gene repertoires (Kraemer et al., 2007; Rask et al., 2010; Day et al., 2017; Chen et al., 2011; Barry et al., 2007; Aguiar et al., 1992). This minimises immune interference and competition between strains by reducing potential overlaps in their antigen repertoires (Gupta et al., 1996), even to the point where hosts may experience concurrent or overlapping infections (Ntoumi et al., 1995; Smith et al., 1999b). Interestingly, while there is large diversity in the specific *var* genes between individual parasite repertoires, the proportion of different Ups groups within repertoires is remarkably conserved (Buckee and Recker, 2012). Assuming that different Ups groups are loosely associated with different host niches, this may suggest that parasites have evolved the ability to exhibit flexible phenotypes depending on the immune status of their host.

While highly polymorphic, PfEMP1 exhibits mosaic domain level structure (Bull et al., 2008; Ward et al., 1999), and the requirement to maintain functional binding may lead to relatively conserved regions. This potentially leads to shared or cross-reactive epitopes between variants and might underpin observations of some degree of short-lived cross-immunity (Ofori et al., 2002; Giha et al., 1999b). Furthermore, it has been hypothesised that weakly immunogenic epitopes may help facilitate the synchronised expression of PfEMP1 during infections and could explain the establishment of chronic infections in individuals with well developed immunity (Recker et al., 2004).

There has been considerable published work characterising the extent of *var* gene diversity (Barry et al., 2007; Rask et al., 2010; Day et al., 2017; Chen et al., 2011). Progress has also been made linking structural and phenotypic diversity in PfEMP1 to disease outcomes (Hviid, 2004; Beeson et al., 2001; Salanti et al., 2004; Rowe et al., 1995; Chen, Schlichtherle, and Wahlgren, 2000; Baruch et al., 1997; Mo et al., 2008; Clausen et al., 2012; Kyriacou et al., 2006). While anti-disease immunity is becoming more clearly understood, little is known about the effect of (variant-specific) immunity on transmission. We argue that models of *P. falciparum* transmission often neglect or oversimplify this complexity, and that immunity, diversity and transmission are intrinsically interlinked, with assumptions about one having important, but often implicit, implications for the others. To address our research questions we developed three modelling frameworks which we referred to as: the cumulative exposure (Chapter 2), the dynamic diversity model (Chapters 3-5), and the repertoire structure model (Chapter 4).

In Chapter 2 we explored some simple deterministic modelling, frameworks which test common assumptions about the action of immunity on transmissibility, and discussed their ability to reproduce field observations regardless the relationship between transmission and prevalence. In particular, we demonstrated that modelling frameworks which completely ignore diversity and assumed no change in transmissibility with exposure (e.g. the SIS framework) are unable to accurately replicate the transmission-prevalence curve. In contrast, frameworks in which exposure can lead to some degree of immunity were able to reproduce this relationship well. One way to interpret this sort of immunity is that the rate at which hosts transition from a non-immune to semi-immune state depends on the diversity in the parasite population, with larger diversity presumably causing hosts to transition more slowly. To investigate this further we developed a more flexible deterministic model

by extending this framework to include an arbitrary number of exposure levels (the cumulative exposure model). This model allowed us to explore the interaction between exposure, immunity and diversity in more detail by directly specifying how key epidemiological parameters change as a function of exposure. We used the cumulative exposure model to demonstrate two mechanisms by which prevalence can be limited, which we referred to as transmission saturation and diversity saturation. A transmission saturated system is characterised by the insensitivity to changes in transmission rate because prevalence is predominantly limited by the available antigenic diversity. In diversity saturated systems the opposite is true, and prevalence is predominantly limited by the host's contact with pathogens.

We went on to describe three possible effects that immunity could have on transmission: (1) the piecemeal acquisition of variant-specific immunity preventing re-infection by strains to which the host has previously been exposed to, (2) exposure to one strain inducing some degree of cross-reactive immunity arising from shared epitopes between strains, and (3) exposure leading to changes in the duration of infection, e.g. through more rapid parasite clearance. All three of these scenarios could equally reproduce the empirically observed transmission-prevalence curves but resulted in different constraints on the predicted values for diversity, transmission rate, recovery rate, and, where appropriate, the strength of cross-immunity. We characterised an interaction between diversity and the duration of infection in which similar transmission-prevalence curves can be generated by different combinations of these parameters. Field observations could thus be explained by high diversity and moderate infection lengths, or lower diversity and longer infections. Better estimates of the duration of infection (and in particular, the effect of exposure on the duration of infection) might therefore help us to better estimate the degree of antigenic diversity which exists in natural *P. falciparum* populations. More generally, a

better understanding of the effect of exposure on the susceptibility to infection and duration of infection is needed. This will help to identify the particular immune processes that are important for future theoretical work and move towards a situation where we can begin to make more qualitative predictions about the influence of antigenic diversity on prevalence.

Few previous models of *P. falciparum* consider diversity explicitly, and many treat diversity as a static quantity whereby hosts are associated with a small probability of attaining a semi-immune state after (repeat) infection. We argued that diversity is better interpreted as a dynamic property, which responds to changes in disease transmission and the local ecology. To explore the epidemiological implications of this we developed the dynamic diversity model, a stochastic individual-based model in which antigens and antigen repertoires could evolve in response to immune selection. This meant that diversity emerged naturally from the underlying transmission setting. In Chapter 3 we used this model to test the epidemiological implications of this. We argued that when diversity is allowed to respond dynamically to transmission, a positive feedback loop is formed in which transmission creates opportunities for diversification, which in turn aids immune evasion and thus facilitates further transmission.

When the parasite population could generate diversity through mitotic and meiotic recombination we found that diversity did not increase indefinitely, even when diversification was essentially unconstrained. Instead, diversity reached an equilibrium level which depended on the specific ecological and epidemiological conditions. In particular, the equilibrium level of diversity was strongly dependent on transmission intensity, recombination rate and population size. When we simulated transmission-prevalence curves, the ability of the parasite population to diversify in

response to immune selection meant that diversity continued to increase as transmission intensity increased. As a result transmission saturation did not occur and prevented prevalence from plateauing at high transmission intensity.

While both population size and recombination rate had a strong effect on the amount of diversity which can be maintained in the parasite population (with larger populations and higher recombination rates being able to maintain greater diversity), neither were able to induce transmission saturation at high transmission intensity. Rather than bringing about a plateau in prevalence at high transmission intensity they simply delayed the onset of 100% prevalence. We found that transmission saturation could only be established when diversification is constrained, e.g. by limiting the potential number of immunologically distinct antigen variants. In our model this could be induced either by increasing the strength of cross-immunity or by reducing the size of the antigenic space. By controlling the point at which transmission saturation occurs, evolutionary restrictions caused prevalence to plateau at levels similar to those observed empirically.

Next we investigated the response of diversity and prevalence to perturbations in transmission intensity. The transmission-diversity feedback introduced a temporal lag in the way that diversity (and thus prevalence) responded to perturbations in transmission intensity. One consequence of this is that historic changes in transmission could influence the epidemiology of a region for much longer than previously thought. The diversity-transmission feedback therefore has potentially important influence on *P. falciparum* epidemiology both in terms of equilibrium prevalence and in the response of the system to changes in transmission rates.

In Chapter 4 we focused on diversity within the *var* gene repertoires of individual parasite genomes. We aimed to elucidate the evolutionary processes which select for and maintain phenotypically diverse antigen repertoires. Past theoretical

work predicts that immune selection should lead to the emergence of strain structure (Gupta et al., 1996). The existence of a strongly conserved repertoire partitioning structure seems to contradict this, since it potentially leads to the inefficient use of globally available diversity (Recker, Arinaminpathy, and Buckee, 2008). To investigate the evolutionary processes which maintain partitioned antigen repertoires we developed the repertoire structure model. This was a stochastic individual-based model in which strains, defined by their antigen repertoires, compete for susceptible hosts. Our main findings were that a combination of variant-specific immunity and cross-immunity could select for phenotypically mixed antigen repertoires. Constraints on diversity lead to uneven between- and within-host immune selection in different phenotype groups and could select for skewed antigen repertoires dominated by the least constrained phenotype group. This mirrored the observations of skewed antigen repertoire proportions in natural *var* gene repertoires (Buckee and Recker, 2012). This mechanism relied on cross-immunity acting in a phenotype specific manner, i.e where exposure to an antigen from one phenotype group had no effect on the immune status of the host with respect to other phenotype groups. While we argued that this is a reasonable assumption to make, simply because we would expect structural similarities between antigens which target the same host receptor, further work is required to confirm this.

In Chapter 5 we explored the effect of partitioned antigen repertoires on strain structure and the acquisition of immunity. We extended the dynamic diversity model to support antigen repertoires comprised of two phenotype groups. This was used to test the effect of differences in the way global antigen diversity and individual parasite repertoires were partitioned. First we described a method to quantify the degree of strain structuring and the number of circulating strains in the dynamic diversity model. We found that the degree of strain structuring and the number of

circulating strains was maximised when the antigenic space was partitioned in the same proportions as that of individual parasite repertoires which is in line with empirical observations. This led to efficient use of the globally available diversity and is in agreement with previous theoretical work (Recker, Arinaminpathy, and Buckee, 2008). We also showed that, as predicted by others (Gupta et al., 1996; Gupta, Ferguson, and Anderson, 1998), strain structure was most pronounced in high transmission settings where between-host immune selection is strongest.

We showed that when one phenotype group is associated with an innate fitness advantage over the other, for example because of greater binding affinity, there is selection for antigen repertoires which 'over-utilise' this group. This causes there to be stronger between-host immune selection on the over-utilised antigen group and meant that a system could be transmission saturated with respect to this group of genes whilst simultaneously being diversity saturated with respect to other group.

We hypothesised that over-utilisation of antigens associated with severe disease could explain age-dependent changes in disease severity. To test this we assumed that one phenotype group was associated with an increased risk of severe disease and were therefore over-utilised in the parasite population. We ensured over-utilisation of this group by fixing the repertoire structure such that repertoires utilise a greater proportion of antigens from this group than would be expected from the globally available diversity alone. Other than this there were no innate differences between the two groups. Under these assumptions we found that hosts gained immunity to the over-utilised group more rapidly than to the other group. Crucially, this only occurred when the system was transmission saturated with respect to the over-utilised group. Interestingly, the other group did not need to be transmission saturated for this effect to manifest, which could offer a possible explanation for how hosts appear to develop immunity against severe disease relatively early in life

but remain susceptible to infection by parasites expressing novel antigen variants throughout their lives. This concurs with analyses of serology networks, which suggests that not all *var* gene groups are under the same degree of immune selection (Buckee, Bull, and Gupta, 2009). Additionally, antigens from the over-utilised group did not need to impart a relative fitness advantage for parasites as long as they were over-utilised in antigen repertoires. We simply used this as a plausible hypothesis to explain why a parasite population might over-utilise one phenotype group in the first place.

Finally, we considered the ability of the antigen expression order to explain age-dependent differences in infection phenotype. Here, no phenotype group was over-utilised but we assumed that one antigen group was preferentially expressed in hosts over the other, e.g. because they confer a selective advantage. We referred to as a transcription hierarchy. We also assumed that parasites typically express only a proportion of their antigen repertoire. When transcription hierarchies were used, the dominant infection phenotype varied strongly with age if the system was transmission saturated, at least with respect to the preferentially expressed group. On the other hand, and in agreement with previous results, we found little change in infection phenotypes with age when the system was diversity saturated. This suggests that transmission saturation is an important prerequisite for age-dependent changes in infection phenotype (taken here to be analogous with infection severity), regardless of whether transcription hierarchies or over-utilisation are responsible for the differing rates of immunity acquisition between antigen groups. In reality, hosts may acquire protection from severe disease as a result of a combination of both mechanisms.

We developed the theoretical concepts of transmission and diversity saturation to help explain the effect of diversity on transmission and prevalence. Much of the

work in this thesis describes mechanisms which behave differently in transmission saturated systems than they do in diversity saturated systems. A greater understanding of how these concepts relate to the epidemiological reality of *P. falciparum* is therefore critical. In particular, fieldwork is needed to characterise the degree to which *P. falciparum* epidemiology is dominated by transmission versus diversity saturation over a range of ecological and transmission settings. This will more fully elucidate the implications of our results and help to develop a more complete understanding of how diversity is being used by the parasite. The distinction between transmission and diversity saturation is, of course, part of a continuous spectrum, and the local ecology and recent transmission history will be a source of considerable variation between regions. However, the extent to which immune selection drives the generation of diversity and the extent to which the system responds to control measures crucially depends on the position of the system along this continuum.

We predict that high transmission regions tend to approximate a transmission saturated state, at least with respect to a subset of *var* genes which are associated with severe disease. This prediction is supported by several aspects of our work. Firstly, using the dynamic diversity model we demonstrated in (Chapter 5) that for the age-distribution of immunity to be sensitive to transmission intensity, the system had to be transmission saturated with respect to at least one phenotype group. Secondly, when we simulated transmission-prevalence curves (in Chapters 2 and 3), prevalence only plateaued at high transmission intensity when the system was transmission saturated. This was true in both the cumulative exposure model and the dynamic diversity model. Thirdly, in Chapter 4 we showed that differences in the strength of between-host immune selection could explain the maintenance of antigen repertoires with skewed proportions. However, in diversity saturated settings between-host immune selection would likely be considerably reduced, because the

probability of a particular parasite encountering hosts with previous exposure to its antigen repertoire is lower.

We argue that a great deal of information could be gained from efforts to characterise the extent of constraints on the ability of *var* genes to diversify and to catalogue the reservoir of diversity under different ecological settings. This could provide significant leverage in the fight to control malaria and help to develop control strategies which are specific to the local ecology. For example, control measures which aim at reducing host contact with parasites (e.g. bed nets) are likely to be limited in their ability to affect prevalence in transmission saturated settings but have a high potential in diversity saturated conditions. Similarly, a better understanding of how the immense diversity of PfEMP1 is structured could also have important implications for vaccine development and deployment. A partially effective vaccine, which for example imparts immunity to a sub-set of PfEMP1 variants (say to a set of highly virulent antigens), may not only be able reduce the burden of severe disease but has the potential to reduce prevalence in transmission saturated settings where other methods would have limited effect. This is because such a vaccine could effectively reduce the usable diversity available to the parasite population in a region where diversity is the primary factor limiting prevalence. In contrast, we would expect this hypothetical vaccine to have virtually no effect on prevalence in a diversity saturated setting.

A logical next step to the work presented in this thesis is to assess the impact of control on host prevalence, strain structure and the acquisition of immunity when diversity is treated as a dynamic property. In particular, future work should focus on characterising the effect of different types of control under a range of ecological and epidemiological settings, which should include both diversity saturated and transmission saturated states.

Much work to date has focused on measuring the number of *var* gene variants in circulation by comparing sequence similarity between DBLα domains. Unfortunately, it is still not clear to what extent this maps to distinct antigen variants which elicit independent immune responses. It is therefore difficult to generate quantitative predictions from models which explicitly incorporate diversity at the antigen level. Recently there has been progress in identifying more detailed structural similarities and associations between conserved blocks within *var* genes and *var* gene groups, as well as in parasite populations (Rask et al., 2010). However, more work is needed to understand the consequences of these associations for the acquisition of immunity. In particular, we require a better understanding of the role of phenotypic diversity in PfEMP1 binding and how domain level structure is associated with different binding phenotypes. This is currently an active area of research and there has been significant progress in recent years. Further work will help us to classify diversity in PfEMP1 in a way which is relevant to understanding host immunity and pathology. This will aid in identifying associations between PfEMP1 structure, pathology and immunity. In turn, this will facilitate understanding of the role that immunity plays in shaping antigen repertoires and how parasites use phenotypic diversity as part of their transmission strategy.

In summary, this work demonstrates that the structured diversity exhibited by *P. falciparum* is an essential component to explain several important aspects of malaria epidemiology. We showed that the fact that diversity is a dynamic emergent property of the underlying transmission dynamics has important implications. Immune selection structures parasite diversity both in terms of strain structure and at the level of individual antigen repertoires. It is therefore of crucial importance to consider evolutionary processes acting at different ecological scales to elucidate the complex immuno-epidemiology of *P. falciparum* malaria.

Appendix A

Sensitivity analysis of the dynamic diversity model

Monofactorial sensitivity analysis was performed for the dynamic diversity model to characterise the relative importance of individual parameters with respect to prevalence, transmission, diversity and immunity. The aim was to investigate the ways in which the main parameters influence the system of feedback loops described in figure 3.7 in Chapter 3 (reproduced in figure A.1 below for convenience) and to quantify the parameter space over which the behaviour described in previous chapters can be considered robust.

The parameters used, intervals considered and baseline values are shown in table A.1. The analysis was conducted by keeping all other parameters at their baseline values and simulating at fixed points over the interval. To characterise the variance caused by stochasticity in the agent interactions, ten repeats were performed for each parameter set.
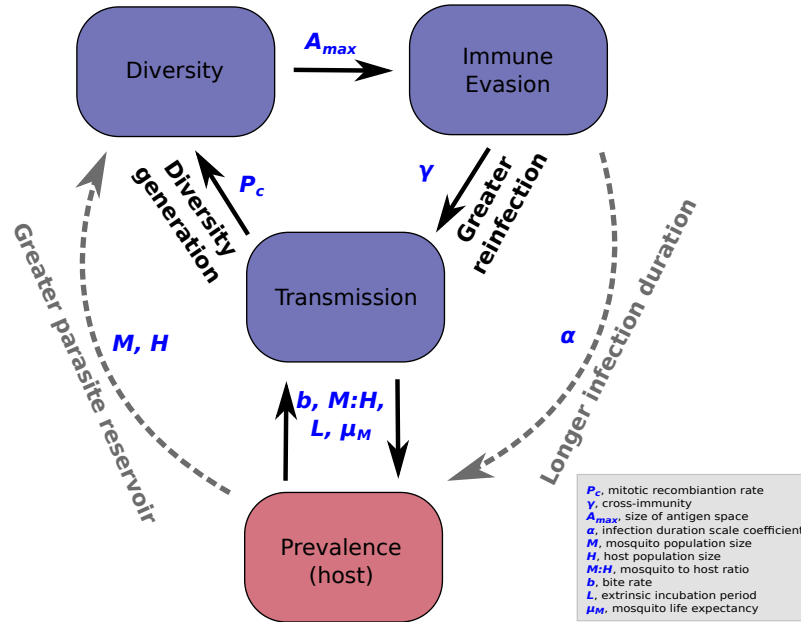
FIGURE A.1: Reproduction of figure 3.7: Interaction between the main processes in the dynamic diversity model. The model parameters which directly affect the strength of each interaction are shown in blue.

TABLE A.1: Parameters used in mono-factorial sensitivity analysis, the ranges considered and their baseline values.

parameter	interval	baseline value
Population size (constant $M:H=1$)	$(4000, 12000)$	8000
$M:H$, mosquito to host ratio	$(0.5, 1.5)$	1
b, mosquito bite rate	$(0.05, 0.25)$	0.12
p_{trans}, infectivity of infected hosts to mosquitoes	$(0.05, 1)$	0.5
α, baseline antigen infection contribution (days)	$(0.4, 3.2)$	2.0
p_c, mitotic recombination rate	$(0, 0.02)$	0.002
θ, scale of change due to mitotic recombination	$0.1, 100$	50
p_s, meiotic recombination rate	$(0, 0.1)$	0.01
γ, strength of cross-immunity	$(0, ???)$	0.0

A.1 Total Population size

In Chapter 3 we demonstrated that larger population sizes increased the capacity of the parasite population to retain antigenic diversity in the absence of the generation of new diversity (figure 3.9). Here, we consider the effect of total population size with diversity generation in effect, and hence with feedback allowed between transmission, diversity generation and immune evasion. Population size was varied over the interval $(8000, 12000)$, while keeping the mosquito to host ratio constant $(M : H = 1)$. All other parameters were set to their baseline values (see table A.1). Ten repeat simulations were conducted for each parameter set, shown in figure A.2.

As expected, larger population sizes were associated with greater equilibrium antigenic diversity, with a linear relationship clearly visible in figure A.2b. Both the transmission rate and prevalence were also positively associated with population size (figure A.2a). This is because in larger host/mosquito populations, which can support greater parasite diversity, more exposure events are needed to acquire a given level of immunity. Figure A.2c reflects this: higher transmission rates in larger populations means that hosts experience greater total exposure (absolute immunity increases), however, greater antigenic diversity means that hosts have immunity to, on average, a smaller proportion of circulating diversity. Despite the increase in exposure the additional diversity associated with larger population sizes appears to be sufficient to prevent transmission saturation, at least in the absence of other factors limiting diversity generation (e.g. as discussed in section 3.3.3).
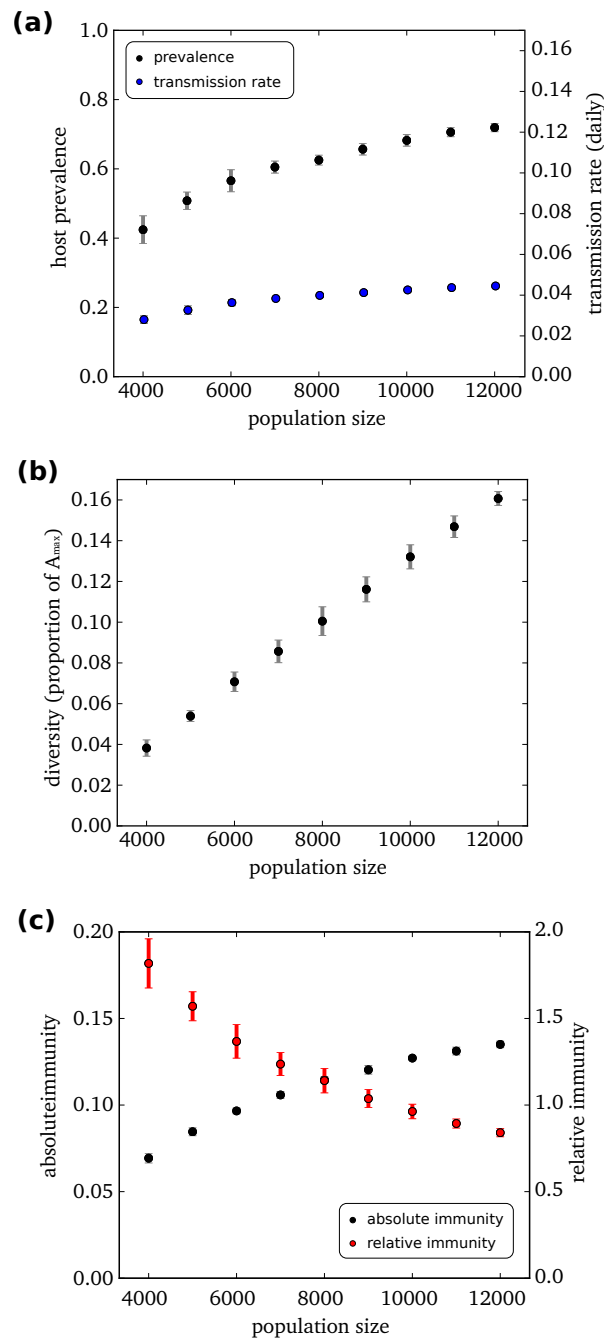
FIGURE A.2: The effect of changes in total population size on **(a)** prevalence, **(b)** antigenic diversity and **(c)** immunity. The sensitivity of prevalence to population size is driven by the ability of larger populations to support larger and more diverse parasite populations, with population size showing a linear positive relationship with the degree of antigenic diversity circulating in the parasite population. This in turn leads to greater total exposure in hosts. Lower relative immunity suggests that larger population sizes reduce the transmission saturation of the system, prevalence from being limited by diversity. See section 3.2.9 for a description of how relative and absolute immunity are calculated. Error bars indicate the 95% confidence interval, and was calculated from ten repeat simulations.

A.2 Mosquito biting rate

Mosquito biting rate (b) was varied over the interval $(0.05, 0.25)$ to assess the effect on prevalence, diversity and immunity. Changing the biting rate directly scales the transmission rate, as can be seen by the linear relationship with transmission rate in figure A.3a. At biting rates below 0.07 transmission could not be sustained using the baseline parameters. As expected, increases in transmission rate lead to an increase in prevalence, which begins to plateau as the system becomes limited by the number of uninfected hosts (figure A.3a). Feedback between transmission and diversity, but also population size and diversity capacity, mean that parasite populations exhibited greater diversity at higher mosquito biting rates (figure A.3b). This diversification appears to mitigate transmission saturation, as can be seen by the fact that relative immunity remains approximately constant despite hosts experiencing greater exposure (higher absolute immunity) when biting rate increases (figure A.3c).
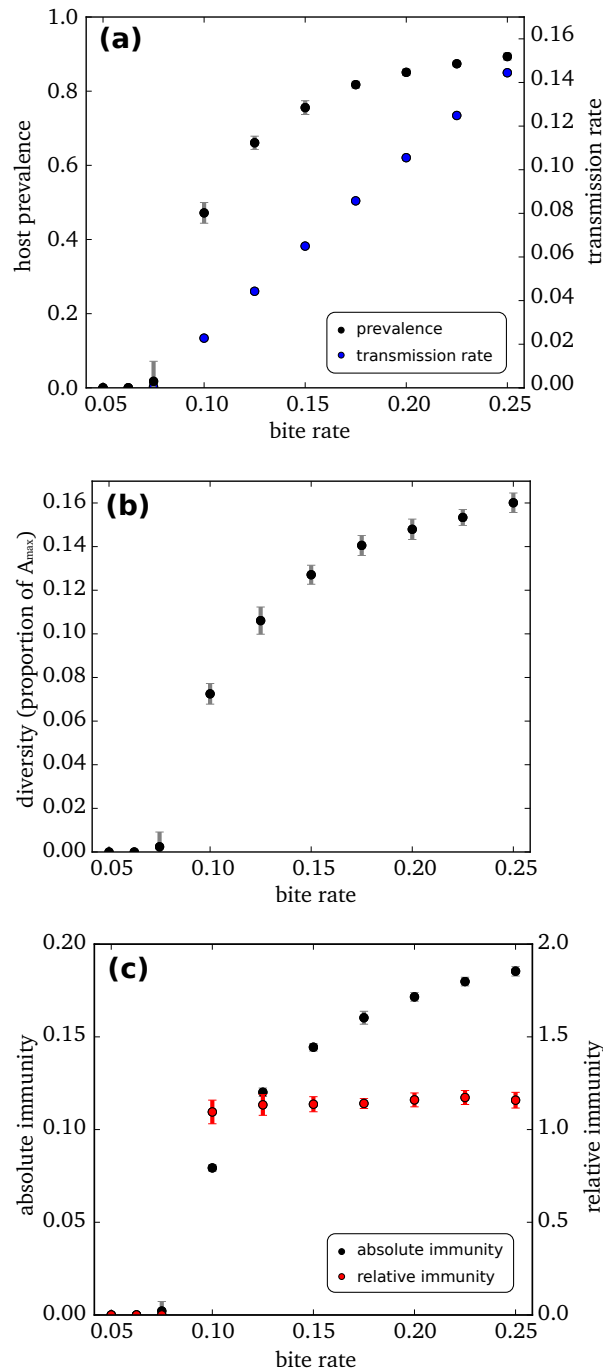
FIGURE A.3: The effect of mosquito biting rate (*b*) on **(a)** preva-
lence and transmission rate, **(b)** circulating antigenic diversity, and
(c) immunity. Transmission rate is linearly and positively related to
mosquito biting rate, which in turn drives prevalence and the gen-
eration of new antigenic diversity. Feedback between transmission
and diversity means that relative immunity remains constant, despite
increasing total exposure. This prevents the system from becoming
transmission saturated. Error bars indicate the 95% confidence inter-
val, calculated from ten repeats.

A.3 Mosquito to host ratio

The ratio of the number of mosquitoes to hosts is an important determinant of transmission in vector borne diseases. It determines the mean number of bites received by hosts as well as being linked to the host and mosquito population size, and therefore has the potential to influence both the ability of the system to maintain diversity and the transmission rate. The mosquito to host ratio, $M : H$, was manipulated in two ways. Firstly, the host population was kept constant while the mosquito population is varied (mosquito driven). Secondly, the host population is varied while the mosquito population remains constant (host driven). In both cases this was performed to generate $M : H$ ratios over the interval $(0.5, 1.5)$, with ten repeat simulations for each value tested, shown in figure A.4.

Both mosquito and host driven changes in $M : H$ ratio cause linear changes in transmission and are therefore also positively associated with prevalence A.4a and d). This arises from the simple fact that higher numbers of mosquitoes relative to hosts increases the rate at which hosts experience bites. It would usually be expected that diversity will increase along with transmission and prevalence, as indicated in the interacting processes diagram (A.1). This is only seen in the mosquito driven scenario, however (compare A.4b and e). This is can be explained by considering the effect of changes in population size, which are necessary to achieve changes in $M : H$ ratio. In the mosquito driven scenario increases in $M : H$ ratio increase the total population size meaning that both population size and transmission rate act in the same direction, to increase the intensity of feedback between transmission, diversity and immune evasion. In contrast, the host driven scenario achieves increases in $M : H$ ratio by decreasing the number of hosts. This has the opposite effect - increases in transmission are only achieved at the expense of population size, and the two

effects work against one another. The net result is that diversity stay approximately constant despite changing transmission rate.

The difference between mosquito driven and host driven scenarios is illustrated by the different way in which host immunity responds to each scenario (A.4e and f). In both cases absolute immunity is positively associated with increases in $M : H$ ratios, reflecting the changes in transmission rate. Relative immunity increases with increases in the $M : H$ ratio for the host driven scenario but stays approximately constant (or decreases slightly) in mosquito driven scenario. The host driven scenario is therefore more prone to becoming transmission saturated, while the mosquito driven scenario is able to generate enough diversity to compensate for this.

Note that we did not attempt to change the $M : H$ ratio in such a way as to keep the total diversity capacity constant. This is because differences fundamental differences between in the life expectancy, number of infections, duration of infection and handling of immunity in hosts and mosquitoes prevent meaningful comparison of their contributions to maintaining the parasite population. It can be said, however, that each host will contribute more to maintaining a parasite population (and the diversity contained within) owing to fact that durations of infection in hosts tend to be many times the life expectancy of mosquitoes, as well as the fact that hosts can be concurrently infected whereas mosquitoes can only have a single infection.
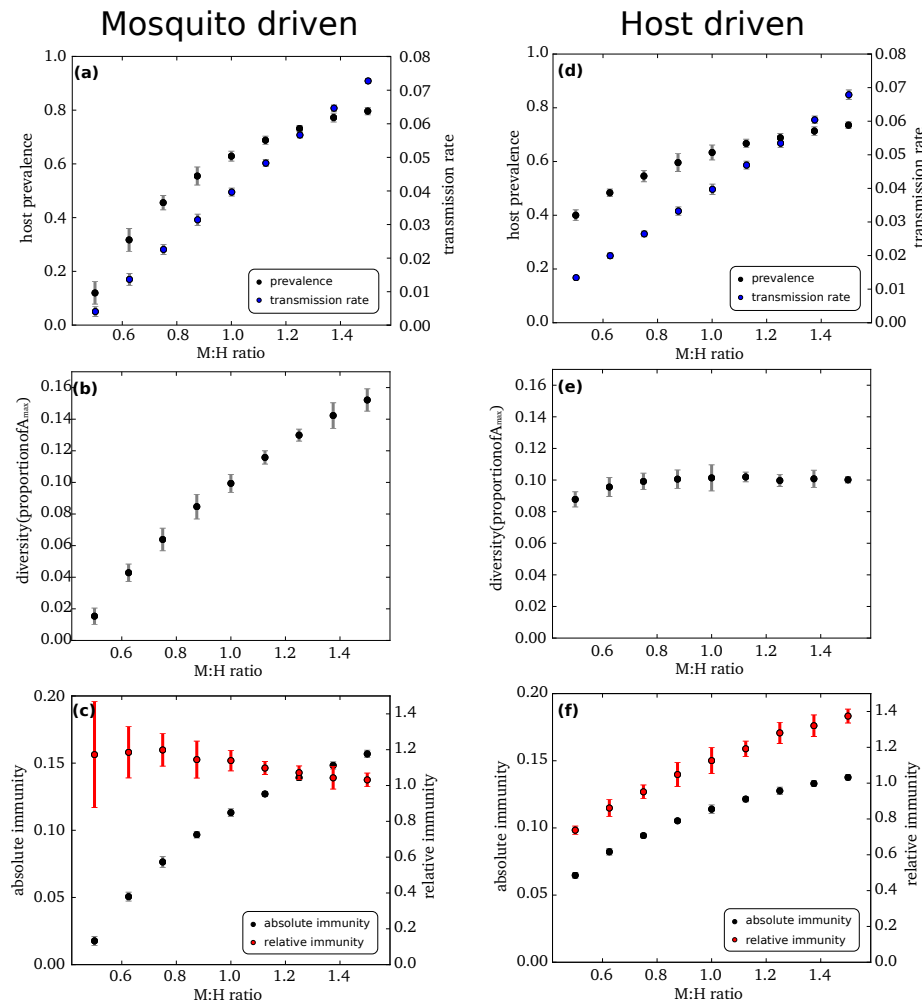
FIGURE A.4: The sensitivity of prevalence, diversity and immunity to mosquito to host ratio ($M : H$) driven by changes in the mosquito population **(a-c)** and the host population **(d-f)**. Changes to the $M : H$ ratio drives changes in transmission, resulting in a positive linear relationship between $M : H$ and transmission rate (**(a)** and **(d)**). Greater transmission results in higher prevalence, which increases the parasite population size. A secondary effect of this is that the diversity capacity of the system also increases (see section 3.3.1). **(b)** In the case of the mosquito driven scenario, this either acts synergistically with changes in number of available hosts (in this case, mosquitoes). This results in rapid increases in parasite diversity with increasing $M : H$. **(e)** In the host driven scenario however, increases in transmission rate are accompanied by decreases in the number of available (human) hosts. These counteract one another and the antigenic diversity of the parasite population stays approximately constant. **(c)** and **(f)** Absolute immunity increases in both scenarios, reflecting that greater transmission rate means hosts are exposed to more antigens over their lifetimes. Relative immunity differs between the two scenarios, however. The extra diversity generated by the mosquito driven scenario prevents the onset of transmission saturation, while the relative immunity increases in the host driven scenario, showing that the system is becoming more transmission saturated at higher $M : H$ ratios. Error bars indicate the 95% confidence interval and was calculated from ten repeat simulations.

A.4 Duration of infection

Duration of infection was changed by altering the baseline contribution to infection made by each antigen, α (figure A.5). Longer infection lengths mean that, once infected, hosts stay infected for longer, and hence we would expect a direct increase in the prevalence. Since hosts which are infected for longer are also infectious for longer in our model, and we would therefore also expect a proportional increase in the transmission rate. Both of these effects are clearly visible in figure A.5a. Interestingly, there is not a linear relationship between duration of infection and EIR (transmission rate), which closely matches the plateauing curve of prevalence. Because of the transmission-diversity feedback, and the ability of larger parasite populations to retain greater diversity, equilibrium diversity is higher for longer infection lengths (figure A.5b), which aids transmission by preventing transmission saturation (figure A.5c). Absolute immunity peaks at moderately high α because, for very long infection lengths, hosts are exposed to fewer strains over the course of their lives.
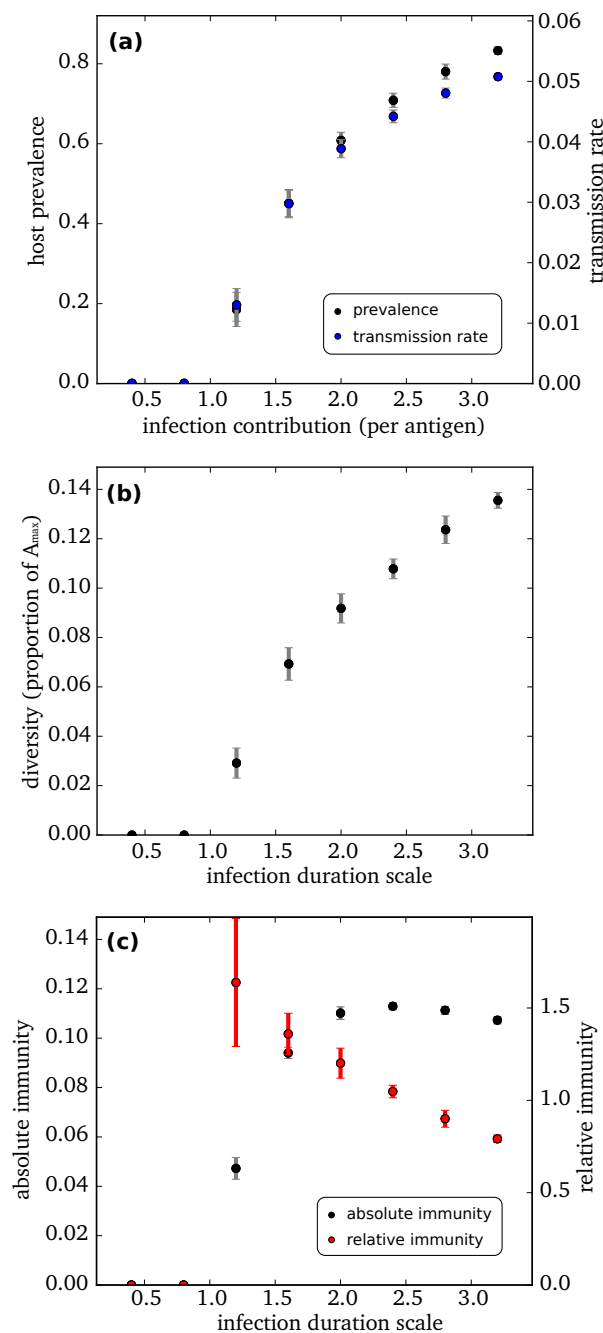
FIGURE A.5: The effect of infection duration on **(a)** prevalence and transmission rate, **(b)** antigenic diversity, and **(c)** immunity. Infection length is a strong determinant of both prevalence and transmission rate, because hosts which are infectious for longer periods of time have more opportunities to infect others. Consequently, feedback between transmission and diversity generation, as well as there being a greater capacity to retain diversity (owing to there being a larger parasite population) cause the equilibrium level of antigenic diversity to increase for longer infections. In the absence of strong constraints on diversification, transmission saturation is prevented, as can be seen by the decrease in relative immunity for larger α. Error bars indicate the 95% confidence interval and was calculated from ten repeats.

A.5 Infectivity of hosts to mosquitoes

The infectivity of hosts to mosquitoes is another parameter which scales transmission intensity. Higher infectivity leads to more mosquito infections and a larger number of infectious bites to hosts (figure A.6a). The transmission-diversity feedback increases with an increase in transmission and prevalence, as in previous examples (figure A.6a). This increases the capacity of the parasite population to generate and maintain diversity (figure A.6b), delaying the onset of transmission saturation (figure A.6c) despite increased exposure.
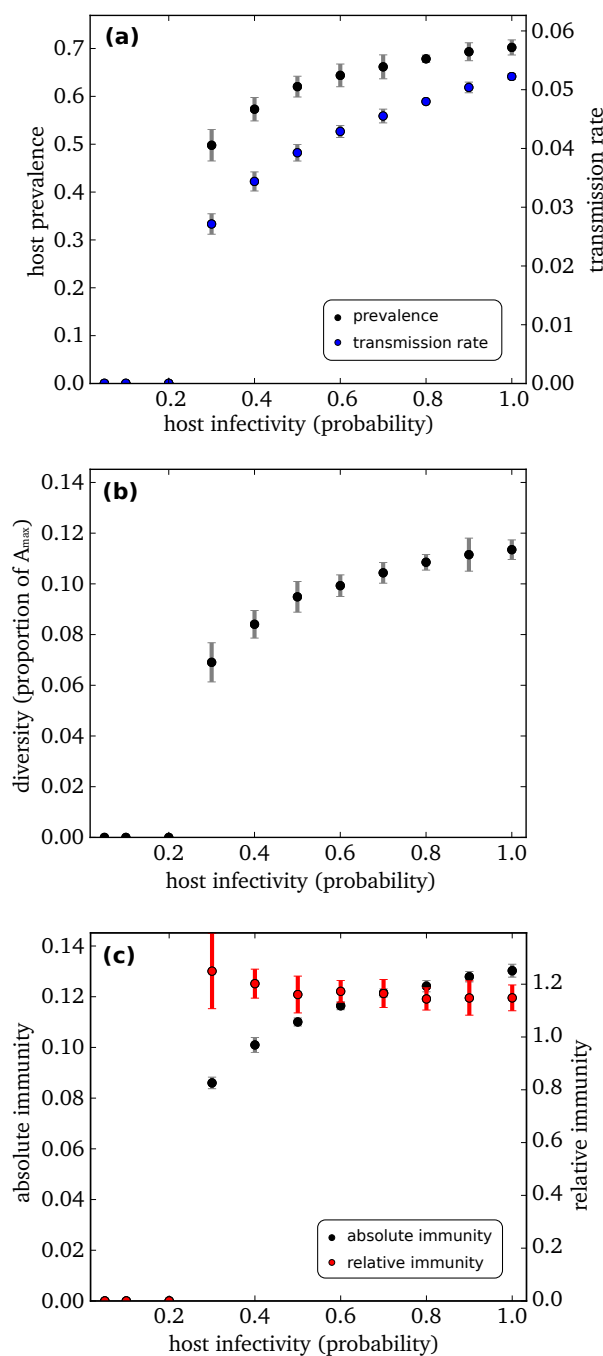
FIGURE A.6: The effect of infectivity of hosts to mosquitoes on **(a)** prevalence and transmission, **(b)** antigenic diversity and **(c)** host immunity. Higher infectivity increases transmission with predictable effects on the interlinked processes of diversity generation and immune evasion: Greater transmission increases diversity generation which causes the equilibrium point of diversity to increase. Absolute immunity increases, because hosts are exposed to a greater number of antigens over their lifespan. Relative immunity decreased, because the diversity generated offsets the extra exposure, and prevents transmission saturation. Error bars indicate the 95% confidence interval, which was calculated from ten repeats.

A.6 Mitotic recombination rate

The effect of mitotic recombination rate on determining degree of antigenic diversity at equilibrium was discussed in Chapter 3. This can be seen again in figure A.7b, in which relatively small increases in the rate of mitotic recombination increases the equilibrium point of antigenic diversity. As a result, there is a sharp drop in the relative immunity of hosts (figure A.7c) and the parasite population is much better at evading host immunity. Transmission rate thus also increases for higher recombination rates, which leads to greater prevalence (A.7a). The primary effect of changes in the mitotic recombination rate is to influence diversity, while the interlinked processes of transmission, diversity generation, immunity and prevalence mean that these changes propagate through each of the systems of feedback shown in figure A.1.
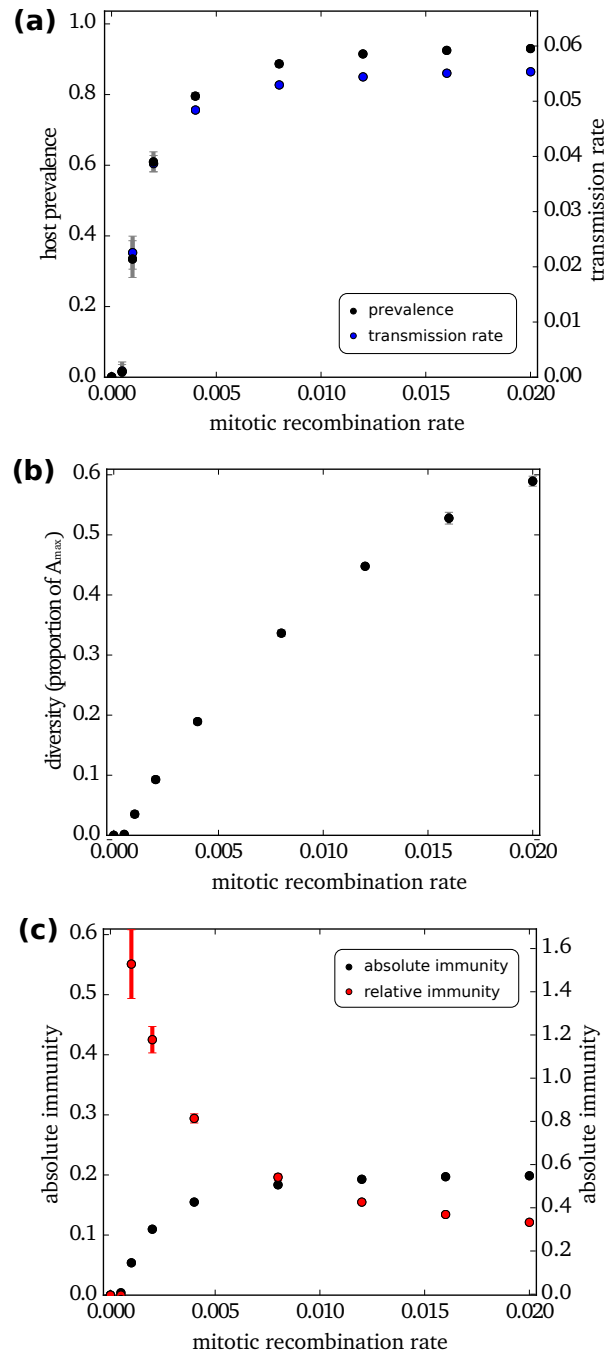
FIGURE A.7: The sensitivity of **(a)** prevalence and transmission, **(b)** antigenic diversity and **(c)** immunity to changes in the rate of mitotic recombination. Mitotic recombination is an important factor determining equilibrium diversity, with higher recombination rates leading a greater degree of diversity in the parasite population. This in turn facilitates immune evasion, greater transmission and higher prevalence. Error bars indicate the 95% confidence interval and was calculate from ten repeats.

A.7   Mitotic recombination scale (θ)

θ scales how different recombinant antigens are from their parents (produced by mitotic recombination). Higher θ effectively increases the rate of exploration of antigenic space, and hence also the adaptability of the parasite population. In practice, since simulations are ran to equilibrium we found no effect of θ on prevalence, transmission, diversity or immunity (figure A.8). Because lower θ decreases the rate of exploration of antigenic space, it increased the time taken to reach equilibrium, but did not affect the final value at equilibrium (see figure A.8)a-c. The exception to this is where θ approaches 0, which effectively turns mitotic recombination off and prevents the parasite population from changing in response to immune selection.
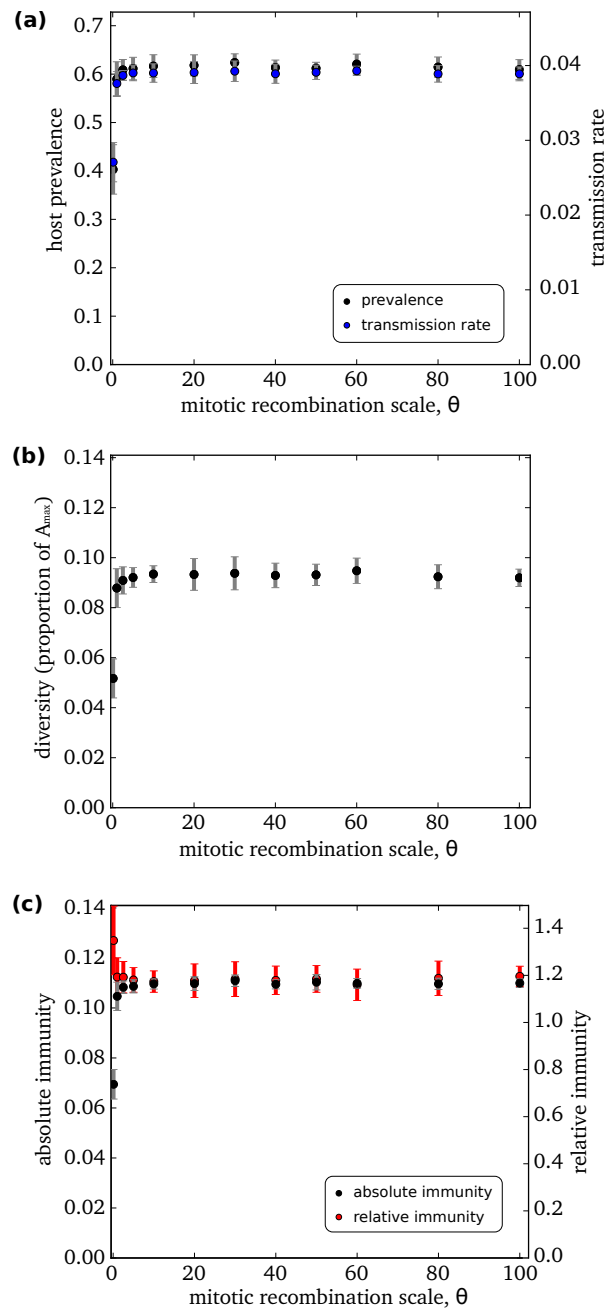
FIGURE A.8: The effect of scale of change resulting from mitotic recombination, controlled by θ, on **(a)** prevalence and transmission rate, **(b)** diversity and **(c)** immunity. In general there was little to no effect of θ on the equilibrium state. Note that for very low θ mitotic recombination is essentially inhibited, resulting in a parasite population which could not adapt to host immunity. In all other cases there was no observed effect. Error bars indicate the 95% confidence interval and was calculated from ten repeats.

A.8 Meiotic recombination rate

When we tested the sensitivity of prevalence, diversity and immunity to changes in meiotic recombination the effect size was consistently much smaller compared to mitotic recombination (figure A.9). Meiotic recombination acts to reorganise existing diversity but does not contribute to the generation of new antigens. The advantage of meiotic recombination to the parasite population is likely to be in enabling rare/adaptive antigens to be exchanged between lineages (thus spreading more rapidly), and in facilitating the emergence of strain structure. However, meiotic recombination is likely to occur considerably less frequently than mitotic recombination because (at least in our model) it relies on mosquitoes biting hosts with concurrent infections. Despite this, higher rates of meiotic recombination were associated with a small increases in prevalence and transmission rate (figure A.9 a). It is likely that the increase in transmission and prevalence is, at least in part, related to the degree of strain structure in the parasite population, although it is not immediately clear why antigenic diversity would be lowest at intermediate levels of meiotic recombination (see figure A.9b). The emergence of strain structure and it's effect on host immunity is discussed in more depth in Chapter 5.
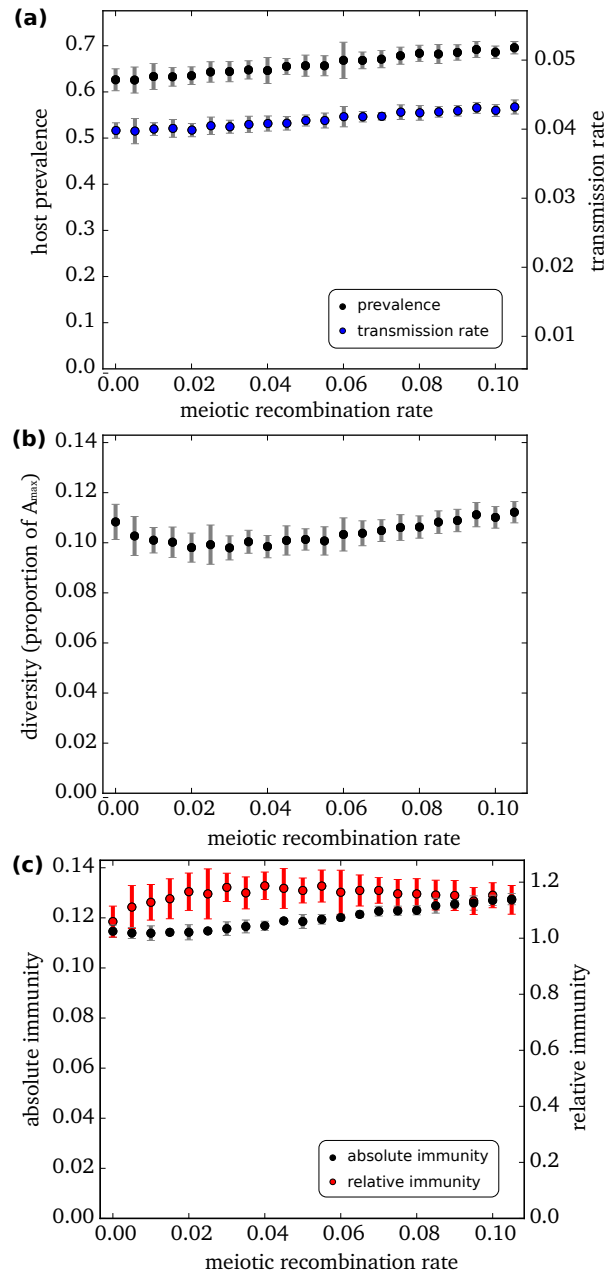
FIGURE A.9: Meiotic recombination, which facilitates the exchange of whole antigens between repertoires, has a much smaller effect on diversity, prevalence and transmission in comparison to mitotic recombination. **(a)** prevalence and transmission rate, **(b)** antigenic diversity and **(c)** immunity. Error bars indicate the 95% confidence interval and was calculated from ten repeats.

A.9   Strength of cross-immunity

When a host is exposed to an antigen, cross-immunity acts in a way which means that they acquire immunity to similar antigens, despite never having been exposed to them. The similarity of one antigen to another is determined by their distance in antigen space. See section 3.2.6 for a full description of how cross-immunity is implemented. The strength of cross immunity is controlled by γ, which scales the distance in antigen space over which cross-immunity acts.

Figure A.10 shows the effect of changes in γ on prevalence and transmission (A.10a), diversity (A.10b) and immunity (A.10c). The primary effect of stronger cross-immunity is to increase the rate immune acquisition (figure A.10c). When cross-immunity is stronger both absolute and relative immunity rapidly increases to very high levels (note the difference in scale compared to previous figures). Much of this immunity is against antigens which are not in circulations in the parasite population, but can still act to restrict the ability of the parasite to diversity. Unsurprisingly this leads to a negative relationship between γ and equilibrium diversity A.10b. In other words, this constrains the parasite's ability to diversify and leads a transmission saturated setting where transmission is primarily limited by diversity. Stronger cross-immunity leads to more severe limitations on diversity, and reduces transmission (and therefore prevalence) further A.10a. For a more in-depth discussion and comparison with other potentially diversity-limiting scenarios, see section 3.3.3.
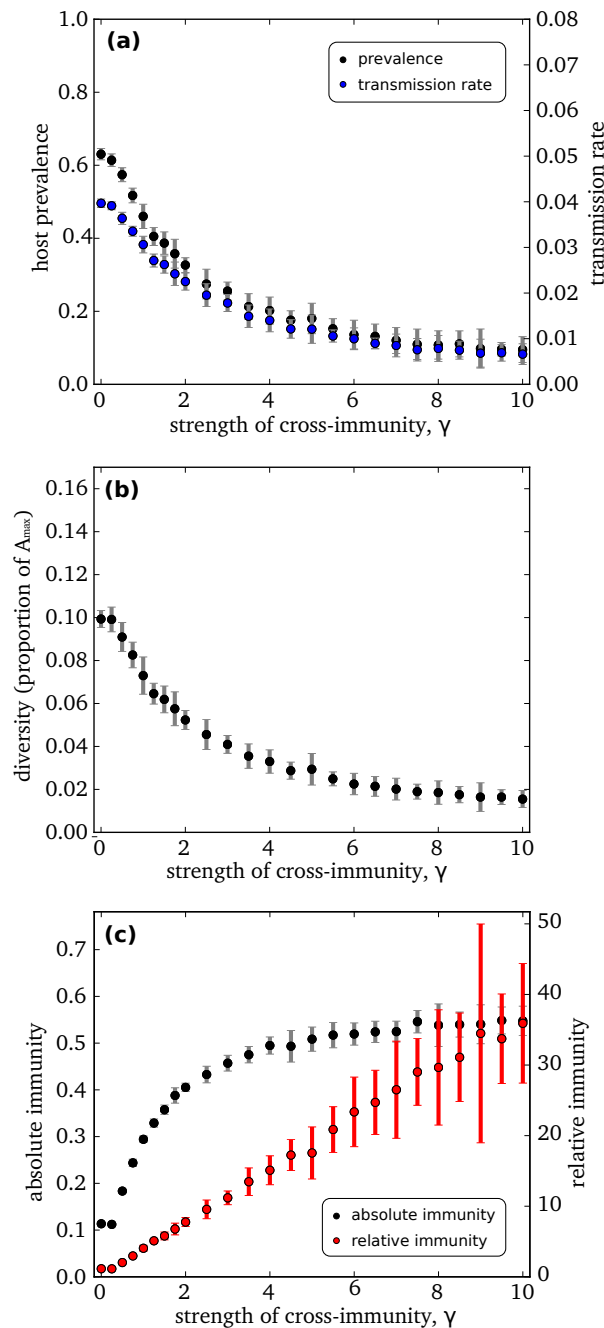
FIGURE A.10: The effect of cross-immunity on **(a)** prevalence and transmission rate, **(b)** antigenic diversity, and **(c)** host immunity. The strength of cross-immunity (γ) primarily acts to change the rate of immune acquisition, and breadth of immunity. For strong cross-immunity (large γ) this leads to a transmission saturated state in which prevalence and transmission are limited by the available diversity. The parasite population cannot escape acquired immunity in the host population. This situation is confounded by smaller parasite populations (lower prevalence) which limits the amount of diversity which can be maintained in the parasite population. Error bars indicate the 95% confidence interval and was calculated from ten repeats.

A.10 Summary

Transmission, prevalence, diversity and immunity are closely interlinked, with changes in on one having important consequences for the others. This is made more complex by the feedback between these processes illustrated in figure A.1. However, different parameters influence this system in different ways depending on which processes are directly altered, and which are affected by secondary effects. Despite the relatively large number of parameters to the model, their effect on the epidemiology can be broadly categorised based on which processes they directly effect: transmission, diversity, immunity or prevalence.

Mosquito biting rate (b), mosquito to host ratio (particularly mosquito driven), host infectivity to mosquitoes and duration of infection directly modulate the transmission rate. They effect the system in a similar way and are generally characterised by strong positive (often linear) relationship with transmission rate, increasing transmission intensity, which in turn increases prevalence and diversity through the mechanisms depicted in figure A.1.

Rate of mitotic recombination and total population size primarily affect diversity: mitotic recombination is responsible for the generation of new antigenic diversity, while larger populations can maintain greater antigenic diversity. Their effect on the system is characterised by a positive linear or near-linear relationship with equilibrium diversity, with increases in these parameters quickly leading to a high degree of diversity saturation and driving increases in prevalence through greater immune evasion. Interestingly the system was relatively insensitive to changes in the rate of meiotic recombination, presumably because it doesn't facilitate the generation of new diversity in and of itself, but may help to spread or maintain structured diversity once generated.

Host driven $M : H$ ratio exhibited a mixture of transmission driven and diversity driven characteristics. More mosquitoes per host meant greater transmission potential, but to achieve this the host population size had to be smaller which reduces the capacity of the population to maintain diversity. When it came to the level of diversity in the parasite population these two effects appeared to approximately balance one another, with the reduced capacity to maintain diversity negating the increased number of opportunities to diversity owing to greater transmission intensity.

Strength of cross-immunity (γ) is the only parameter which we tested that affects the degree of immunity directly. High γ constrains the ability of the parasite population to adapt to host immunity, effectively reducing the size of antigenic space. A similar result was found in section 3.3.3 (both for size of antigenic space and strength of cross-immunity), and was notable for it's ability to induce transmission saturation and therefore explain the plateauing of prevalence at high transmission intensity.

Appendix B

Publications

Author for correspondence:
Mario Recker
e-mail: m.recker@exeter.ac.uk

THE ROYAL SOCIETY
PUBLISHING

Maintenance of phenotypic diversity within a set of virulence encoding genes of the malaria parasite *Plasmodium falciparum*

Thomas Holding and Mario Recker

Centre for Mathematics and the Environment, University of Exeter, Penryn Campus, Penryn TR10 9EZ, UK

Infection by the human malaria parasite *Plasmodium falciparum* results in a broad spectrum of clinical outcomes, ranging from severe and potentially life-threatening malaria to asymptomatic carriage. In a process of naturally acquired immunity, individuals living in malaria-endemic regions build up a level of clinical protection, which attenuates infection severity in an exposure-dependent manner. Underlying this shift in the immunoepidemiology as well as the observed range in malaria pathogenesis is the *var* multigene family and the phenotypic diversity embedded within. The *var* gene-encoded surface proteins *Plasmodium falciparum* erythrocyte membrane protein 1 mediate variant-specific binding of infected red blood cells to a diverse set of host receptors that has been linked to specific disease manifestations, including cerebral and pregnancy-associated malaria. Here, we show that cross-reactive immune responses, which minimize the within-host benefit of each additionally expressed gene during infection, can cause selection for maximum phenotypic diversity at the genome level. We further show that differential functional constraints on protein diversification stably maintain uneven ratios between phenotypic groups, in line with empirical observation. Our results thus suggest that the maintenance of phenotypic diversity within *P. falciparum* is driven by an evolutionary trade-off that optimizes between within-host parasite fitness and between-host selection pressure.

1. Introduction

Plasmodium falciparum is the most virulent human malaria parasite. A key virulence determinant is the ability of infected red blood cells to bind to a variety of host receptors, causing sequestration in the deep vasculature and obstruction of blood flow in vital organs. Cytoadhesion is mediated by the highly polymorphic *Plasmodium falciparum* erythrocyte membrane protein 1 (PfEMP1), expressed on the surface of infected red cells and encoded by the *var* multigene family [1–4]. These proteins are also prominent targets for host immune responses, and each parasite genome contains a different repertoire of ≈60 *var* genes, of which only one gene is actively transcribed at a time. In a process known as clonal antigenic variation, parasites switch expression between individual *var* genes in a mutually exclusive fashion, causing new variants to appear over the course of an infection [5,6].

Different PfEMP1 variants bind to different host receptors [3] and can lead to parasite sequestration in different host tissues, such as the brain or placenta [7]. Probably, the best known example of this is the involvement of *var2csa* in pregnancy-associated malaria, which can be found in every *var* gene repertoire sequenced to date and whose protein product (VAR2CSA) binds exclusively to placental CSA, leading to severe complications during pregnancy [8]. Therefore, *in vivo* antigenic variation not only aids immune evasion, but can also evoke temporal changes in the phenotypic profile of the parasite population during

infection, with important consequences for infection outcome. Note, for the purpose of this work, we refer to phenotype or phenotypic group, as a subset of *var* genes with similar binding characteristics.

Although hosts living in malaria-endemic areas may never attain a state of sterile immunity, they acquire a broad spectrum of PfEMP1-specific antibodies with repeated exposure [9–12]. This has been shown to attenuate infection severity, such that severe and fatal malaria is usually confined to the youngest age groups, whereas older individuals only rarely experience clinical malaria episodes [13]. Underlying this transition in the immunoepidemiology of *P. falciparum* malaria appears to be a qualitative change in *var* gene expression during infection mediated by the host's immune repertoire [14,15], indicating an exposure-mediated link between acquired immunity, phenotype selection and malaria pathology.

Despite their enormous sequence diversity, *var* genes can be grouped according upstream promoter sequence (Ups) or chromosomal location, with a remarkably conserved distribution of these groups within parasite repertoires [16,17]. This conserved partitioning appears to be facilitated by recombination being largely confined to within groups, meaning that groups may be evolving independently of one another [7,18]. Furthermore, these genetically defined groups seem to associate with binding phenotype [19], which suggests that *P. falciparum* parasites have not only evolved to maximize antigenic diversity, but also evolved a mechanism by which they maintain phenotypic diversity within their repertoires of antigen encoding and virulence-associated genes.

While the evolutionary forces that favour phenotypically diverse antigen repertoires remain largely unknown, the central role that PfEMP1 plays in acquired immunity and immune evasion makes selection by the immune system a likely candidate. Previous work exploring the effects of immune selection predominantly focused on between-host interactions and predicted that parasite populations will self-organize into antigenically discordant strains, thereby minimizing the detrimental effect of previous exposures on subsequent infections owing to cross-immunity [20]. The highly conserved partitioning of *var* gene repertoires, which has also been shown to restrict the number of unique antigen combinations and thus only poorly uses the available global pool of antigenic variants [21], is difficult to reconcile with this hypothesis, however.

More recently, we have postulated how evolutionary trade-offs at multiple ecological levels (within-host and within-populations) can explain observed distributions of *var* genes and gene domains with respect to their degree of sequence conservation [22]. Crucially, while this framework explored the evolution of gene and domain distributions within already partitioned repertoires, it did not address why parasites should encode multiple phenotypic groups, rather than evolving into separate and independently transmissible phenotypes.

Here, using an evolutionary framework, we show how a trade-off between within-host fitness and between-host selection pressure can robustly select for phenotypically diverse antigen repertoires. Furthermore, we show that functional constraints that limit the degree by which different phenotypes can diversify without losing binding functionality are sufficient to explain observed *var* gene repertoire structures in *P. falciparum*.

2. Methods

We developed a stochastic, individual-based model to investigate the evolutionary factors that select for phenotypically diverse antigen repertoires. We defined parasites strains by their antigenic repertoire, which itself comprised up to two phenotypic groups. That is, each repertoire was made up of antigenic variants of either phenotype *A* or phenotype *B*. We assumed that all strains have the same repertoire size, referred to as *r*, but could differ in the number of variants belonging to either group *A* or *B*. We refer to a particular repertoire structure as $(A:B)$, indicating the ratio of antigens within the repertoire belonging to phenotype *A* or phenotype *B*. Note, repertoire structure can be considered as a collective term, referring to all strains with the same distribution of genes of type *A* and type *B*.

Owing to the high dimensionality of the potential antigenic space that can be generated simply through recombination of the global pool of antigenic variants, we investigated gene repertoire evolution by repeatedly simulating population infection histories with a fixed set of parasite strains. At the beginning of each simulation, we initialized the parasite population to contain exactly one member of every possible repertoire structure (i.e. one of each $(r:0)$, $(r-1:1), \ldots, (0:r)$) and filled the repertoires by randomly sampling from the two global variant pools (of sizes N_A and N_B). Each simulation was run until it reached a dynamic equilibrium (around 10 000 simulated days), at which point all but a small subset of strains are driven to extinction. We refer to those strains that are maintained within the population as 'dominant', irrespective of their population-level prevalence. That is, dominance is here defined as a purely qualitative measure referring to those strains that have competitively excluded all other strains.

For the majority of simulations, we restricted the genomic repertoire sizes to $r = 9$ and antigenic pool sizes to $N_A = 13$ and $N_B = 13$, which are all significantly smaller than the assumed diversity of *var* genes both within individual repertoires and at the population level. These numbers were chosen for computational reasons only (larger numbers take significantly longer to converge), however, and we note that our results were independent of repertoire and global variant pool sizes. Examples of larger repertoire and variant pool sizes can be found in electronic supplementary material, figure S1.

Hosts are modelled individually with the host population being kept constant at 10 000 and deaths being replaced by newborns. Host mortality was assumed to be age-dependent, given as

$$p_{\text{death}}(T) = \frac{0.002 e^{0.07T}}{365},\tag{2.1}$$

where *T* is the host age in years. We ignored maternal protection and further assumed that infections do not contribute to death. At the start of each simulation, hosts are initialized such that the population is demographically at equilibrium.

For the sake of simplicity, we kept the total rate of infection constant over time and instead assumed transmission to be solely frequency dependent, where the (daily) probability of a host getting infected with strain *i* depends only on the relative frequency of strain *i* in the population. The daily infection rate of strain *i* can then be given as

$$p_{\text{inf}}^i = \lambda_0 \frac{\pi_i}{\sum_j \pi_j},\tag{2.2}$$

where $\lambda_0 = 0.033$ is the daily force of infection and π_i is the proportion of hosts currently infected with strain *i*.

Upon infection, it was assumed that the infecting strain will express its antigenic repertoire in a random order, where each variant contributes positively to infection length. Hosts acquire immunity in a variant-specific manner, which prevents the specific variant from contributing to future infections. In this

fashion, hosts might acquire full immunity to a particular strain without necessarily having been exposed to it. Where cross-immunity is considered, not all variant genes will contribute equally to infection length, in which case variant-specific immunity is acquired probabilistically and in proportion to their contribution to infection length, L.

Within this framework, a strain's fitness solely refers to infection length and is therefore not a fixed quantity, but dependent on the immune repertoire of the host it currently infects. Infection length was determined by the number of novel antigens the infecting strain presents to the host, with no explicit assumptions about expression order and no intrinsic differences between the two phenotypic groups. This might be seen in slight deviation from hypothesized fitness differences between UpsA and UpsB/C *var* genes, for example, where the former is often observed to dominate infections in young and immunologically naive hosts. However, by assuming an equal contribution by all variants, independent of phenotype, we create a more robust null-model, which does not rely on specific gene expression orders or growth rate differences in differently pre-exposed hosts.

With regards to immune interactions, we considered three different scenarios (i) variant-specific immunity only, (ii) phenotype-specific cross-immunity, and (iii) phenotype-transcending cross-immunity.

2.1. No cross-immunity

In the absence of any immune interactions, i.e. where immunity is solely variant-specific, we can define the length of infection, L_n, simply by a linear function of the number of novel antigens of type A and B (n_A and n_B, respectively), which in its simplest form could just be given as

$$L_n = L_A + L_B = l_0(n_A + n_B), \qquad (2.3)$$

with l_0 being the contribution of each antigen to infection length. As mentioned above, we did not assume any intrinsic difference between genes or phenotypes with regards to growth rate or immunogenicity, so each variant gene contributes equally to total infection length.

2.2. Phenotype-specific cross-immunity

We argued that antigens belonging to the same phenotype will have higher sequence similarities, and hence antigenic similarities, than antigens binding to different host receptors. We therefore considered the case whereby phenotype-specific cross-reactive immune responses would build up during the infection, which, in turn, would reduce the contribution of each additionally expressed variant—*of the same phenotype*—to infection length. The total infection length is then governed by a law of diminishing returns, given as

$$L_s = L_A + L_B = \sum_{j=0}^{n_A-1} l_0 e^{-j\sigma} + \sum_{k=0}^{n_B-1} l_0 e^{-j\sigma} \text{ and} \qquad (2.4)$$

$$= l_0\left(\frac{1-e^{-n_A\sigma}}{1-e^{-\sigma}} + \frac{1-e^{-n_B\sigma}}{1-e^{-\sigma}}\right), \qquad (2.5)$$

with σ being the degree of cross-reactivity ($\sigma \in [0, 1]$). For most of this work, we assumed cross-immunity to be transient and reset to zero between consecutive infections; this was later relaxed, however, without changing the qualitative nature of our results.

2.3. Phenotype-transcending cross-immunity

Finally, we also considered the case where cross-immunity would build up during the infection in relation to any antigen presented, i.e. independent of their phenotypic group. In this case, the total infection length can be calculated as

$$L_t = \sum_{i=0}^{n_A+n_B-1} l_0 e^{-i\sigma} = l_0 \frac{1-e^{-(n_A+n_B)\sigma}}{1-e^{-\sigma}}. \qquad (2.6)$$

In order to investigate the robustness of evolved repertoire structures under different contributions of phenotype-specific and phenotype-transcending cross-immunity, we defined

$$\Gamma := l_0 \left(\sum_{j=0}^{n_B-1} e^{-j\sigma} - \sum_{i=n_A+1}^{n_A+n_B-1} e^{-i\sigma}\right), \qquad (2.7)$$

such that the infection length under the consideration of both phenotype-specific and phenotype-transcending cross-immunity becomes

$$L_c = \sum_{j=0}^{n_A-1} l_0 e^{-j\sigma} + \sum_{k=0}^{n_B-1} l_0 e^{-j\sigma} - \delta\Gamma \text{ and} \qquad (2.8)$$

$$= L_A + L_B - \delta\Gamma. \qquad (2.9)$$

Therefore, $\delta = 0$ corresponds to phenotype-specific cross-immunity only and $\delta = 1$ corresponds to phenotype-transcending cross-immunity. Equally, setting $\sigma = 0$ recovers the infection length considering no cross-immunity at all.

3. Results

3.1. Variant-specific immunity selects against phenotypic diversity

We first considered the case where host immunity is solely variant-specific, such that infection length is linearly correlated with the number of novel variants a parasite strain presents to its host (see Methods). Previous theory predicts that immune selection can cause parasite populations to organize into sets of dominant strains with minimal antigenic similarity. In this case, the immune responses induced by these strains cover a sufficiently large part of the available antigenic diversity to suppress the emergence or invasion of other (recombinant) strains [20,23]. We find a similar behaviour in our model system, where pairs of strains consisting of two homogeneous repertoires containing only one or the other phenotype, i.e. with a (phenotype A : phenotype B) repertoire structure of $(r:0)$ and $(0:r)$, which by definition share no antigen variants, more often end up being dominant with all other strains competitively excluded. Note that we consider a strain to be dominant if it is not extinct by $t = 10\,000$. An example time series of this scenario is shown in figure 1*a*, where strains with phenotypically heterogeneous repertoires are driven to extinction and a set of strains with repertoire structures $(r:0)$ and $(0:r)$, and hence no antigenic overlap, become dominant.

Although the particular outcome of each simulation with respect to the repertoire structure that ends up dominating the population is highly stochastic and strongly dependent on the initial composition of strains, we find that on average strains whose repertoires are more biased towards one phenotype or another have a significant competitive advantage, simply because the chance of antigenic overlap between their respective repertoires is minimal compared with strains with a more even distribution of phenotypes within their repertoire. In other words, as strains become more balanced in terms of their antigens belonging to one or the other phenotypic group, the chance of overlap between them increases. This,
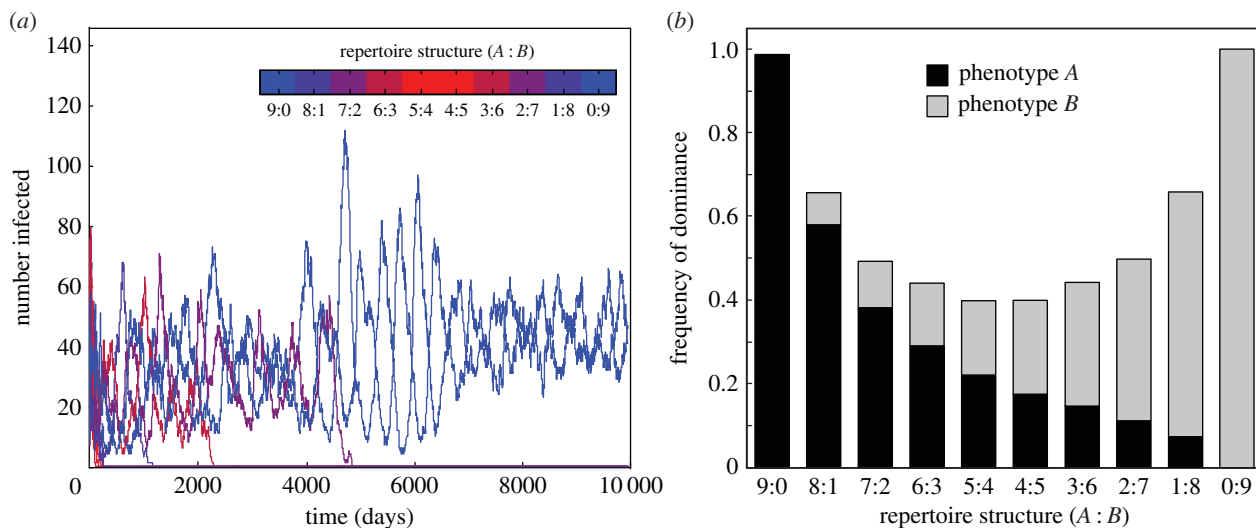
Figure 1. Default model behaviour under variant-specific immunity only. (a) Time series of a single model run show the changing prevalence of strains with differently structured repertoires, with strains coloured according to their repertoire structure (A : B). Phenotypically diverse strains are driven to extinction as a pair of strains consisting of only one or the other phenotype (i.e. with repertoire structures (9 : 0) and (0 : 9)) start to dominate. (b) Bar graph of the frequency by which different repertoire structures dominate the population, based on 10 000 model runs, shows how phenotypically less diverse parasites are favoured under variant-specific immunity. Stacked bars are used to visually represent within-repertoire phenotype distribution. Other parameter values: $r = 9$, $N_A = 13$, $N_B = 13$.

on the other hand, reduces their population-level fitness as infection by one strain is more likely to induce partial immunity to another. This is shown in figure 1b, where we used average dominance frequency as an estimate of the (population-level) fitness of a particular repertoire structure. As mentioned in Methods, dominance is a qualitative measure referring to a strain or set of strains that is maintained in the population and has competitively excluded all other strains, as exemplified in figure 1a.

3.2. Phenotype-specific cross-immunity selects for partitioned gene repertoires

An important part of immune protection in malaria is the ability of anti-PfEMP1 antibodies to inhibit and potentially reverse binding of the infected red cells to host receptors [24,25]. It is therefore conceivable that antibodies raised against a variant of one phenotype should exhibit some cross-inhibition of binding of other variants of the same phenotype, but not necessarily inhibit the binding of variants that target a different host receptor. We therefore examined the effects of phenotype-specific cross-immunity on *var* gene repertoire evolution by further assuming that cross-immunity accumulates over the course of infection with each expressed variant of the same type. This, in turn, reduces the contribution of every additionally expressed variant to infection length, resulting in the total infection length, L_n, in relation to the number of novel antigens exposed to the host, being governed by a law of diminishing returns (see Methods for more details).

Depending on the degree of cross-immunity, σ, we note that a strain's infection length, even in naive hosts, quickly plateaus with the number of expressed variants, as shown in figure 2a, implying there is an upper limit to the length of an infection which cannot notably be extended by expressing additional variants of the same phenotype. Crucially, however, this limitation can be circumvented by the parasite by expressing members of an alternative phenotypic group.

Under the addition of phenotype-specific cross-immunity, antigen repertoires now experience selection pressure by two conflicting forces. At the population level, the pressure to reduce the probability of encountering hosts with prior immunity selects for homogeneous repertoire structures and low phenotypic diversity, which minimizes antigenic overlap with other strains (as shown above, figure 1). At the level of individual infections, on the other hand, repertoires are under selection pressure to maximize infection length, which is best satisfied by having phenotypically diverse repertoires containing equal numbers of each phenotypic group. As a result, under moderate-to-high degrees of cross-immunity, σ, we now find phenotypically diverse strains dominating the population, as shown by an example time series and repertoire structure dominance frequency in figure 2b and c, respectively.

To further understand and illustrate the evolution of different repertoire structures under changes in the degree of cross-immunity, we can calculate the relative fitness of a particular strain by means of its infection length in a host previously infected by any of the other strains. For illustration purposes only, we restrict the analysis to a limited antigenic pool where the total number of variants of phenotype A or B equals the parasite's repertoire size, such that a pair of phenotypically homogeneous strains would completely exhaust the available antigenic space. The matrix plots in figure 3 show the infection length of strain i in a host with previous exposure to strain j for all (i,j)-pairs, which clearly demonstrate how phenotypically diverse repertoires start to gain a competitive advantage as the degree of cross-immunity increases (from $\sigma = 0$, (a), to $\sigma = 0.7$, (d)), leading to their preferred selection and dominance within the population.

Importantly, the selection of phenotypically diverse strains is robust to parameter changes in host population sizes, repertoire size, force of infection and antigen expression order and further persists, regardless of whether we assume cross-reactive immunity to be transient and lost between infections or whether it is permanent and accumulates over repeated infections (shown in electronic supplementary
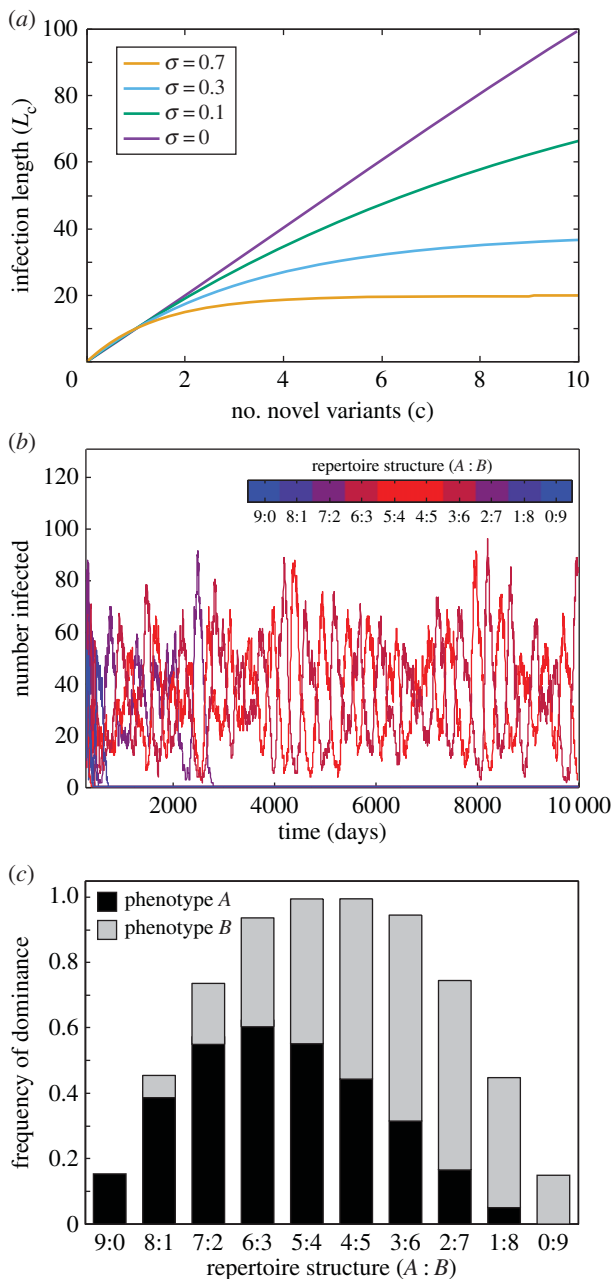
Figure 2. Cross-immunity selects for phenotypically diverse repertoires. (*a*) Total infection length as a function of the number of antigenic variants presented over the course of infection, *c*, under variation in the degree of cross-immunity, σ. For high levels of cross-immunity, infection length starts to plateau well before the parasite can exhaust its antigenic repertoire and additionally expressed genes will not contribute to further infection length. (*b*) Time series of a single model run with temporary cross-immunity ($\sigma = 0.3$) showing the competitive exclusion of phenotypically uniform parasites (blue lines) in favour of strains containing partitioned and hence phenotypically diverse repertoires (red and purple lines). (*c*) Bar graph of the average strain dominance frequency based on 10 000 model runs shows how phenotypically diverse strains gain a selective advantage under phenotype-specific cross-immunity ($\sigma = 0.3$). Stacked bars are used to visually represent within-repertoire phenotype distribution. Other parameter values: $r = 9$, $N_A = 13$, $N_B = 13$.

material, figures S1–S5). We also tested the robustness of these results by relaxing our assumption of cross-immunity to be restricted to the respective phenotypes (see Methods). We now find that as cross-immunity shifts away from being confined to variants of the same phenotype to become more predominantly phenotype transcending, parasite

repertoires become phenotypically less diverse. As shown in electronic supplementary material, figure S4, even moderate-to-high degrees of phenotype-transcending immunity can select for within-repertoire phenotypic diversity as long as there is a non-negligible degree of phenotype-specific cross-immunity.

3.3. Functional constraints influence repertoire partitioning

Data from fully sequenced *var* gene repertoires reveal a conserved structure with very uneven group sizes. We therefore investigated possible conditions that could offer an evolutionary advantage of skewed phenotype distributions, concentrating again on the functionality of PfEMP1 in terms of receptor binding.

It can be argued that the evolution and diversification of PfEMP1 variants is constrained by the necessity to retain functional binding to their target receptors. It is further possible that different binding phenotypes, which target different host receptors, are constrained to different degrees, for example, owing to the specific structural characteristics of the binding targets. On the one hand, this may limit the global antigenic diversity of each phenotype to a different degree, i.e. resulting in different global pool sizes N_A and N_B. On the other hand, it may result in different degrees of cross-reactive inhibition experienced by variants of each phenotype, i.e. where the degree of cross-immunity, σ, becomes phenotype-specific (σ_A and σ_B).

We first examined the influence of global antigen diversity on the repertoire partitioning structure by considering differently sized antigen pools of the two phenotypic groups *a* and *b*. We note that by increasing the size of one phenotype's antigen pool, the likelihood of overlaps within a particular phenotype group decreases. In the absence of cross-immunity, this results in a competitive advantage to repertoire structures that use larger numbers of this (more diverse) phenotype (figure 4*a*). With the addition of cross-reactive immunity, the extent to which large numbers of antigens from this phenotype group can be usefully exploited is curtailed, however. This results in a skewed distribution in repertoire structure dominance frequency, where homogeneous repertoire structures (with low phenotypic diversity) are strongly disadvantaged while moderately heterogeneous repertoires with a bias towards the more diverse phenotype group are favoured (figure 4*b*).

Next, we considered differential degrees by which antibodies can inhibit PfEMP1 binding in a phenotype-specific manner. We argued that because PfEMP1 is a modular protein, distinct variants may still share whole domains or exhibit regions of high similarity between domains. Furthermore, domains can be organized into motifs comprised of specific domain subclasses in tandem that are commonly referred to as domain cassettes. The above-mentioned functional constrains on domain diversity together with observed correlations between groups of domain subclasses, which may act as functional units, could limit the number of antigenically distinct binding epitopes and thus increase the degree of cross-reactivity within variants containing similar domains. Broad inhibitory cross-reactivity at the domain/domain cassette level has been shown previously [25] and while we do not model the domain-level structure of PfEMP1 explicitly, we incorporated differences in
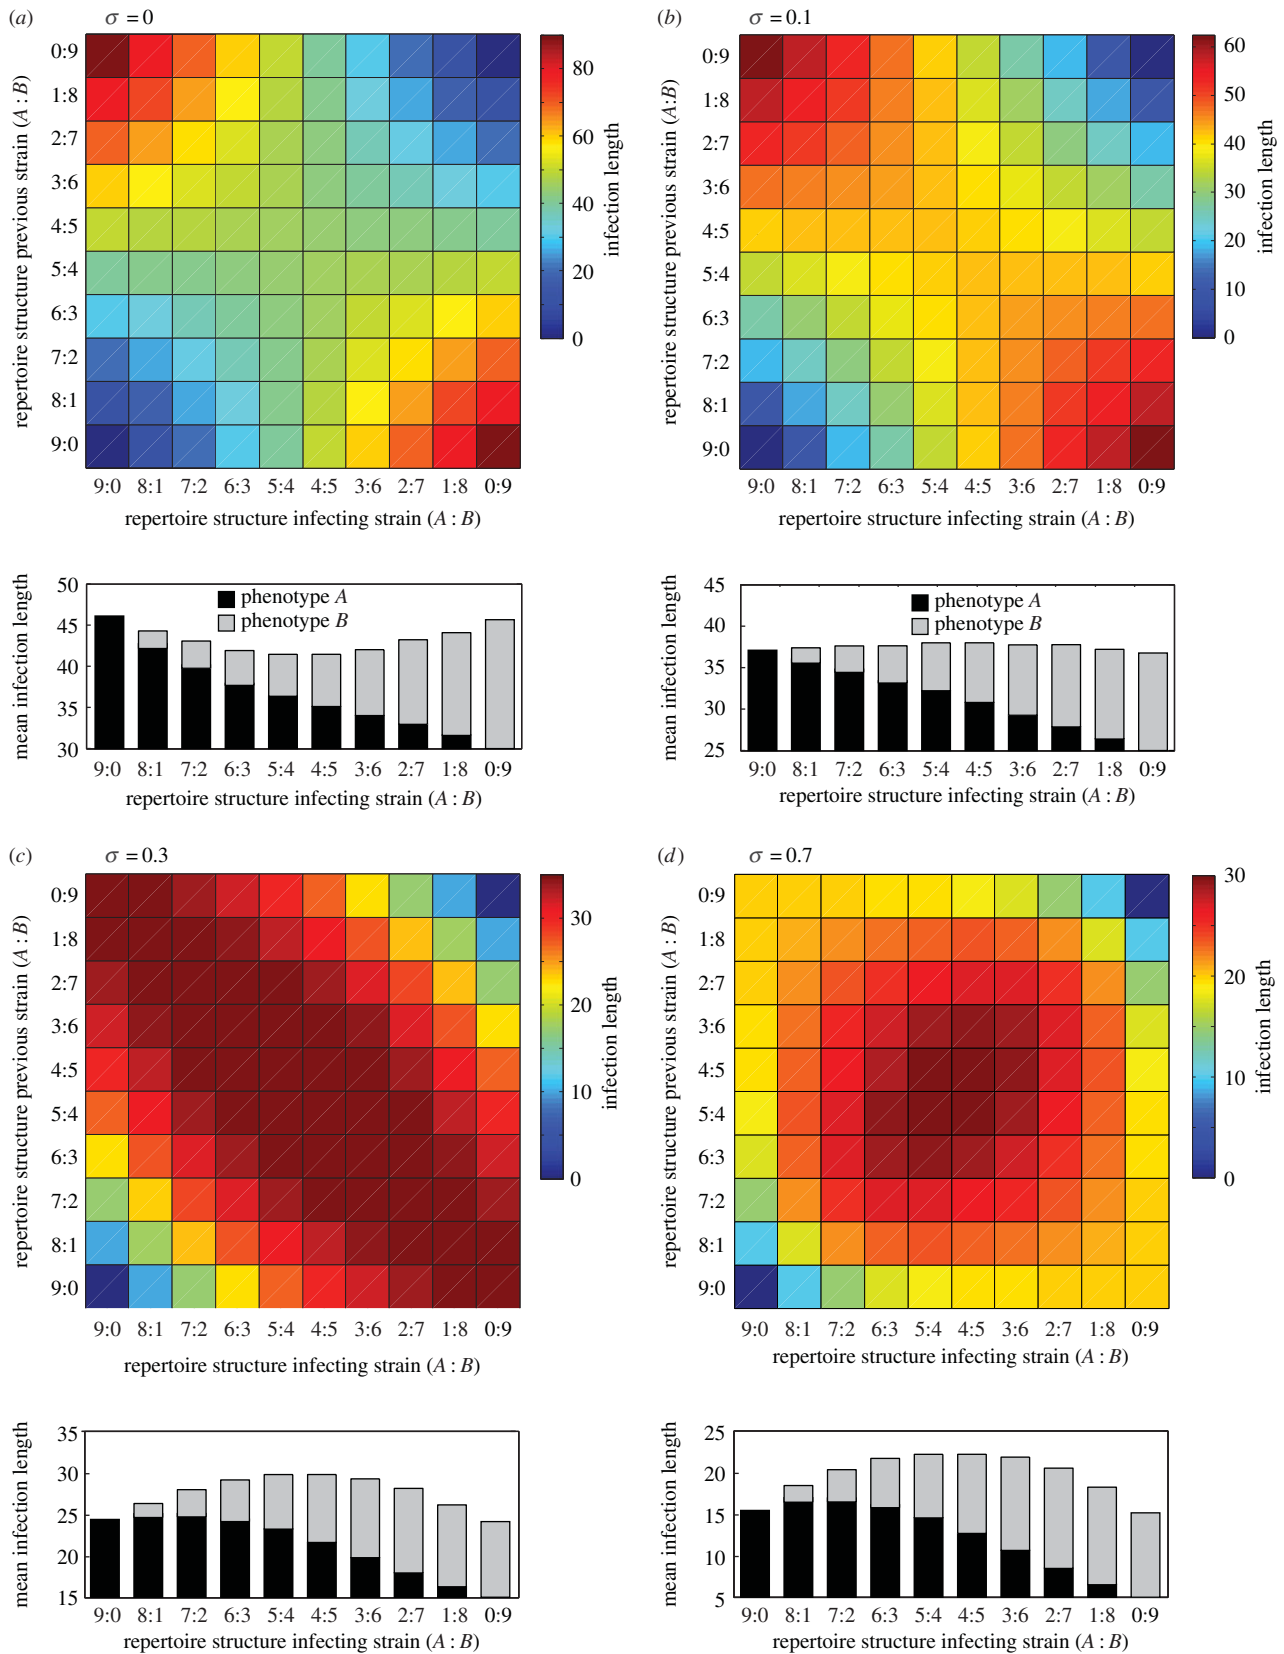
Figure 3. Pairwise and mean infection length under changes in cross-immunity. (*a* – *d*) Matrix plots show the (*i,j*) pairwise infection lengths, resulting from an infection of strain *i* in a host previously exposed to strain *j*. The bar graphs depict the mean infection length of each repertoire structure (represented as stacked bars to represent within-repertoire phenotype distribution), derived by averaging the strains' length of infection against all non-self-strains (i.e. the column mean of the pairwise infection matrix). As cross-immunity is increased ((*a* – *d*)), phenotypically diverse repertoires start to gain a competitive advantage. Note that different scales are used to illustrate the qualitative change in repertoire fitness. (*a*) $\sigma = 0$, (*b*) $\sigma = 0.1$, (*c*) $\sigma = 0.3$, (*d*) $\sigma = 0.7$. Other parameter values: $r = 9$, $N_A = 9$, $N_B = 9$.

cross-reactivity by using phenotype-specific values of σ. The results are similar to those obtained by considering different antigenic pool sizes, leading to skewed phenotype distributions whereby strains with moderately heterogeneous

repertoires can be seen to enjoy an evolutionary advantage, resulting in frequent selection and dominance in the population (figure 4*c*), in line with empirical observations (inset in figure 4*b*).
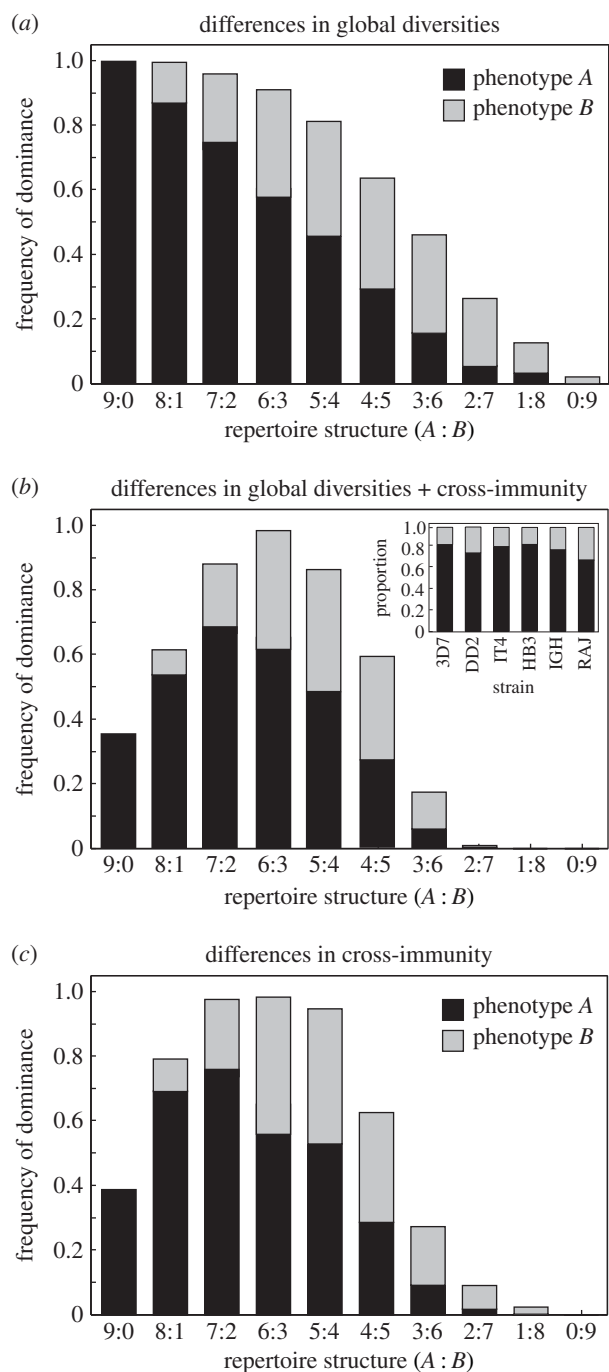
Figure 4. Functional constraints on PfEMP1 select for skewed phenotype distributions. Bar graphs show the average strain dominance frequencies based on 10 000 model runs, with stacked bars used to represent within-repertoire phenotype distribution. (*a*) Considering different antigenic pool sizes of phenotypes *A* and *B* ($N_A = 36$ and $N_B = 9$) with variant-specific immunity only selects for parasites with mono-phenotypic repertoires. (*b*) Including cross-immunity ($\sigma = 0.1$) restricts the benefit of expressing more variants of the same phenotype, leading to the evolution of phenotypically diverse repertoires skewed towards the more diverse phenotype (*A*). The inset shows the repertoire distribution of UpsA versus non-UpsA *var* genes in seven sequenced *P. falciparum* isolates, shows a skewed and conserved repertoire partitioning. (*c*) Differences in the degree of cross-reactivity, with $\sigma_A = 0.1$ and $\sigma_B = 0.8$, has the same effect as different antigenic pool sizes and select for repertoires with uneven numbers of phenotypic variants. Other parameter values: $r = 9$.

4. Discussion

The highly immunogenic nature of the *var* gene-encoded malaria virulence factors PfEMP1 makes them important targets for host protective immune responses and therefore subject to intense selection pressure. Previous theoretical work on the evolution of *var* gene repertoires has predicted that strong immune selection will act to structure the parasite population into strains with minimal antigenic similarity [20], such that an infection by one strain has little or no detrimental impact on the fitness of subsequently infecting strains. Since then, further empirical work on *var* genes and *var* gene repertoires has highlighted a conserved genomic structuring, whereby gene repertoires are not just random collections of antigenic variants, but contain distinct groups of genes differentiated for example by their genomic location, upstream promoter sequence, number of binding domains or even sequence diversity [6,17,22]. These characteristics are interrelated, however, and there appears to be some mechanism in place that maintains the observed repertoire structure over evolutionary time [26] despite high rates of recombination. This, on the other hand, strongly suggests that gene repertoires containing distinct groups of genes have a competitive advantage despite the poor utilization of the globally antigenic diversity this entails [21], especially when considering that the ratios of genes belonging to the different groups also appear highly conserved [6,16,17].

In line with previous predictions, we found that variant-specific immune responses can select for parasite strains with minimally overlapping antigenic repertoires. Considering that antigenic variants can be of two different phenotypes, this implied an evolutionary trend towards the dominance of pairs of strains comprising variants of only one or the other phenotype. We note that this particular outcome was partly driven by our model set-up to contain one representative strain of every possible repertoire structure. In this case, two strains whose repertoires contained variants of only one or the other phenotypic group will be (antigenically) non-overlapping by default, whereas the probability that two strains containing variants of both phenotypic groups have overlapping repertoires increases as the phenotype distribution becomes more even. On the other extreme, when we initialized the model such that each strain has a non-overlapping counterpart, there are no longer population-level fitness differences and all strains will be equally likely to become dominant. One way or the other, this clearly illustrates that variant-specific immunity alone cannot account for the empirically observed phenotypic diversity within *var* gene repertoires.

We therefore concentrated on the functional importance of PfEMP1, i.e. adhesion to host cell receptors, as a potential factor in *P. falciparum* evolution. Because of its crucial contribution to within-host parasite fitness, by removing the infected red cells from circulation and thus avoiding clearance by the spleen, we argued that (antigenic) diversification of PfEMP1 may be limited for the protein to retain functional viability, and we further argued that this limitation would increase the chance of antibody-mediated cross-inhibition of parasite binding. The presence of transient cross-reactive immune responses has been well established in longitudinal studies [27,28], although whether these are specific to particular phenotypes or not is not yet clear. More recent work has confirmed the presence of cross-reactive immunity between variants with specific binding phenotypes, for example in IgM rosetting [29], as well as broad cross-reactivity to particular epitope regions, such as domain cassette 4, which is associated with ICAM binding [25], where both binding phenotypes are associated with severe infection pathology. Furthermore, structural conservation

in EPCR-binding PfEMP1 has been shown despite large sequence diversity [30], although the contribution this makes to cross-reactive antibody-mediated immunity, especially over the course of an infection, is yet to be established.

We tested two PfEMP1-binding related hypotheses for their effect on the evolution of intragenomic repertoire partitioning: (i) phenotype-specific differences in the degree by which the proteins, or particular binding loci, can diversify without losing functionality, leading to different pool sizes of antigenic variants and (ii) phenotype-specific differences in the degree of cross-inhibition. Both could, in fact, be argued to be part of the same process, such that a constraint on protein diversification might entail an increase in cross-recognition by antibodies. In terms of their role in gene repertoire evolution, we found that both had the same effect and caused a shift in repertoire structuring towards biased phenotype distributions, similar to those observed in *P. falciparum*. Based on features of short sequence tags within the Duffy-binding-like α domain, Bull *et al.* [31] reported similar frequency distributions of groups of *var* genes within the genome as within the population, despite them being very different in size. Similarly, it is well known that the UpsA group of *var* genes exhibits less global diversity than other groups and also comprises a relatively small proportion of the *var* gene repertoire. The most extreme example might be that of *var2csa*, whose presence in every repertoire and unusual sequence conservation compared with other *var* genes makes it an attractive vaccine target [32]. Further empirical investigations are necessary to test whether these observations could, in fact, be due to our hypothesized functional constraints or differences in immunogenic cross-inhibition, however.

We assumed that variant genes are being expressed sequentially over the course of an infection, at least to a level sufficiently high to trigger an immune response. Under this assumption, each additionally expressed gene contributes positively to infection length and thus transmission probability. In the most extreme case, this led to every variant having an equal contribution to the length of infection and thus within-host fitness, which is highly unlikely to be the case in real malaria infection. On the other hand, this particular scenario could not explain the observed repertoire structures. Once we considered cross-immunity, this assumption was partially relaxed and we showed that as long as individual variants would contribute positively to within-host fitness, i.e. as long as infection length was a monotonically increasing function of antigens expressed during infection, the model robustly predicted evolution towards phenotypically diverse repertoire structures. Furthermore, it is also worth mentioning that the sequential dominance by a single or only a few antigenic variants is thought to underlie the characteristic waves of parasitaemia in *P. falciparum* malaria [33–35]. In that respect, it is also interesting to note that transient cross-reactive immune responses against groups of antigenic variants sharing particular PfEMP1 epitopes have previously been proposed to orchestrate the parasite population to exhibit the observed single-variant dominance during infection [36]. This adds further to the notion that cross-reactive immunity could play a highly significant role for parasite evolution and malaria epidemiology in general.

Another interesting consequence of transient cross-reactive immunity is that it offers a simple explanation as to why *var* gene repertoires have evolved to be limited to ≈60 variants. Under variant-specific immunity alone, we would expect either selection for vast repertoires of antigenically distinct variants or the evolution of a mechanism by which diversity can easily be generated during the course of infection, both of which are common strategies for other antigenically variable organisms [37]. Although high rates of mitotic recombination in culture adapted *P. falciparum* isolates have recently been reported [38], it is not yet clear if and to what degree this could contribute to immune evasion during infection or whether its role is more limited towards the generation of antigenic diversity at the population level. Therefore, the relatively limited size of the antigenic repertoire of *P. falciparum* has been a bit of a conundrum, which temporary cross-immunity could possibly resolve, although this is purely speculative at this stage and requires further investigation.

In summary, our results support the hypothesis that phenotypically diverse *var* gene repertories are maintained by a combination of between-strain and within-host immune selection. A similar mechanism has recently been put forward to underlie the sequence and domain evolution of individual *var* genes and gene repertoires [22], which strongly suggests that evolutionary trade-offs, by which the parasite has to optimize between the within-host and within-population-level fitness, are important determinants in *P. falciparum* evolution.

References

1. Su XZ, Heatwole VM, Wertheimer SP, Guinet F, Herrfeldt JA, Peterson DS, Ravetch JA, Wellems TE. 1995 The large diverse gene family *var* encodes proteins involved in cytoadherence and antigenic variation of *Plasmodium falciparum*-infected erythrocytes. *Cell* **82**, 89–100. (doi:10.1016/0092-8674(95)90055-1)

2. Baruch DI, Pasloske BL, Singh HB, Bi X, Ma XC, Feldman M, Taraschi TF, Howard RJ. 1995 Cloning the *P. falciparum* gene encoding PfEMP1, a malarial variant antigen and adherence receptor on the surface of parasitized human erythrocytes. *Cell* **82**, 77–87. (doi:10.1016/0092-8674(95)90054-3)

3. Smith JD *et al.* 1995 Switches in expression of *Plasmodium falciparum var* genes correlate with changes in antigenic and cytoadherent phenotypes of infected erythrocytes. *Cell* **82**, 101–110. (doi:10.1016/0092-8674(95)90056-X)

4. Rowe JA, Moulds JM, Newbold CI, Miller LH. 1997 *P. falciparum* rosetting mediated by a parasite-variant erythrocyte membrane protein and complement-receptor 1. *Nature* **388**, 292–295. (doi:10.1038/40888)

5. Roberts DJ, Craig AG, Berendt AR, Pinches R, Nash G, Marsh K, Newbold CI. 1992 Rapid switching to multiple antigenic and adhesive phenotypes in malaria. *Nature* **357**, 689–692. (doi:10.1038/357689a0)

6. Kyes SA, Kraemer SM, Smith JD. 2007 Antigenic variation in *Plasmodium falciparum*: gene organization and regulation of the *var* multigene family. *Eukaryotic Cell* **6**, 1511–1520. (doi:10.1128/EC.00173-07)

7. Kraemer SM, Smith JD. 2006 A family affair: *var* genes, PfEMP1 binding, and malaria disease. *Curr. Opin. Microbiol.* **9**, 374–380. (doi:10.1016/j.mib.2006.06.006)

8. Salanti A, Staalsoe T, Lavstsen T, Jensen AT, Sowa MP, Arnot DE, Hviid L, Theander TG. 2003 Selective upregulation of a single distinctly structured *var* gene in chondroitin sulphate A-adhering *Plasmodium falciparum* involved in pregnancy-associated malaria. *Mol. Microbiol.* **49**, 179–191. (doi:10.1046/j.1365-2958.2003.03570.x)

9. Giha HA *et al.* 1999 Nine-year longitudinal study of antibodies to variant antigens on the surface of *Plasmodium falciparum*-infected erythrocytes. *Infect. Immun.* **67**, 4092–4098.

10. Bull PC, Lowe BS, Kortok M, Molyneux CS, Newbold CI, Marsh K. 1998 Parasite antigens on the infected red cell surface are targets for naturally acquired immunity to malaria. *Nat. Med.* **4**, 358–360. (doi:10.1038/nm0398-358)

11. Bull PC, Marsh K. 2002 The role of antibodies to *Plasmodium falciparum*-infected erythrocyte surface antigens in naturally acquired immunity to malaria. *Trends Microbiol.* **10**, 55–58. (doi:10.1016/S0966-842X(01)02278-8)

12. Hviid L. 2005 Naturally acquired immunity to *Plasmodium falciparum* malaria in Africa. *Acta Trop.* **95**, 270–275. (doi:10.1016/j.actatropica.2005.06.012)

13. Giha HA, Staalsoe T, Dodoo D, Roper C, Satti GM, Arnot DE, Hviid L, Theander TG. 2000 Antibodies to variable *Plasmodium falciparum*-infected erythrocyte surface antigens are associated with protection from novel malaria infections. *Immunol. Lett.* **71**, 117–126. (doi:10.1016/S0165-2478(99)00173-X)

14. Nielsen MA, Staalsoe T, Kurtzhals JA, Goka BQ, Dodoo D, Alifrangis M, Theander TG, Akanmori BD, Hviid L. 2002 *Plasmodium falciparum* variant surface antigen expression varies between isolates causing severe and nonsevere malaria and is modified by acquired immunity. *J. Immunol.* **168**, 3444–3450. (doi:10.4049/jimmunol.168.7.3444)

15. Warimwe GM, Recker M, Kiragu EW, Buckee CO, Wambua J, Musyoki JN, Marsh K, Bull PC. 2013 *Plasmodium falciparum var* gene expression homogeneity as a marker of the host–parasite relationship under different levels of naturally acquired immunity to malaria. *PLoS ONE* **8**, e70467. (doi:10.1371/journal.pone.0070467)

16. Kraemer SM *et al.* 2007 Patterns of gene recombination shape *var* gene repertoires in *Plasmodium falciparum*: comparisons of geographically diverse isolates. *BMC Genomics* **8**, 45. (doi:10.1186/1471-2164-8-45)

17. Rask TS, Hansen DA, Theander TG, Gorm Pedersen A, Lavstsen T. 2010 *Plasmodium falciparum* erythrocyte membrane protein 1 diversity in seven genomes—divide and conquer. *PLoS Comput. Biol.* **6**, 667–674.

18. Kraemer SM, Smith JD. 2003 Evidence for the importance of genetic structuring to the structural and functional specialization of the *Plasmodium falciparum var* gene family. *Mol. Microbiol.* **50**, 1527–1538. (doi:10.1046/j.1365-2958.2003.03814.x)

19. Janes JH, Wang CP, Levin-Edens E, Vigan-Womas I, Guillotte M, Melcher M, Mercereau-Puijalon O, Smith JD. 2011 Investigating the host binding signature on the *Plasmodium falciparum* PfEMP1 protein family. *PLoS Pathog.* **7**, e1002032. (doi:10.1371/journal.ppat.1002032)

20. Gupta S, Maiden MC, Feavers IM, Nee S, May RM, Anderson RM. 1996 The maintenance of strain structure in populations of recombining infectious agents. *Nat. Med.* **2**, 437–442. (doi:10.1038/nm0496-437)

21. Recker M, Arinaminpathy N, Buckee CO. 2008 The effects of a partitioned *var* gene repertoire of *Plasmodium falciparum* on antigenic diversity and the acquisition of clinical immunity. *Malar. J.* **7**, 18. (doi:10.1186/1475-2875-7-18)

22. Buckee CO, Recker M. 2012 Evolution of the multi-domain structures of virulence genes in the human malaria parasite, *Plasmodium falciparum*. *PLoS Comput. Biol.* **8**, e1002451. (doi:10.1371/journal.pcbi.1002451)

23. Gupta S, Anderson RM. 1999 Population structure of pathogens: the role of immune selection. *Parasitol. Today (Regul. Ed).* **15**, 497–501. (doi:10.1016/S0169-4758(99)01559-8)

24. Oleinikov AV, Amos E, Frye IT, Rossnagle E, Mutabingwa TK, Fried M, Duffy PE. 2009 High throughput functional assays of the variant antigen PfEMP1 reveal a single domain in the 3D7 *Plasmodium falciparum* genome that binds ICAM1 with high affinity and is targeted by naturally acquired neutralizing antibodies. *PLoS Pathog.* **5**, e1000386. (doi:10.1371/journal.ppat.1000386)

25. Bengtsson A *et al.* 2013 A novel domain cassette identifies *Plasmodium falciparum* PfEMP1 proteins binding ICAM-1 and is a target of cross-reactive, adhesion-inhibitory antibodies. *J. Immunol.* **190**, 240–249. (doi:10.4049/jimmunol.1202578)

26. Otto TD *et al.* 2014 Genome sequencing of chimpanzee malaria parasites reveals possible pathways of adaptation to human hosts. *Nat. Commun.* **5**, 4754. (doi:10.1038/ncomms5754)

27. Kinyanjui SM, Bull P, Newbold CI, Marsh K. 2003 Kinetics of antibody responses to *Plasmodium falciparum*-infected erythrocyte variant surface antigens. *J. Infect. Dis.* **187**, 667–674. (doi:10.1086/373994)

28. Bull PC, Lowe BS, Kaleli N, Njuga F, Kortok M, Ross A, Ndungu F, Snow RW, Marsh K. 2002 *Plasmodium falciparum* infections are associated with agglutinating antibodies to parasite-infected erythrocyte surface antigens among healthy Kenyan children. *J. Infect. Dis.* **185**, 1688–1691. (doi:10.1086/340420)

29. Ghumra A *et al.* 2012 Induction of strain-transcending antibodies against group A PfEMP1 surface antigens from virulent malaria parasites. *PLoS Pathog.* **8**, e1002665. (doi:10.1371/journal.ppat.1002665)

30. Lau CK *et al.* 2015 Structural conservation despite huge sequence diversity allows EPCR binding by the PfEMP1 family implicated in severe childhood malaria. *Cell Host Microbe.* **17**, 118–129. (doi:10.1016/j.chom.2014.11.007)

31. Bull PC, Berriman M, Kyes S, Quail MA, Hall N, Kortok MM, Marsh K, Newbold CI. 2005 *Plasmodium falciparum* variant surface antigen expression patterns during malaria. *PLoS Pathog.* **1**, e26. (doi:10.1371/journal.ppat.0010026)

32. Hviid L. 2010 The role of *Plasmodium falciparum* variant surface antigens in protective immunity and vaccine development. *Hum. Vaccin.* **6**, 84–89. (doi:10.4161/hv.6.1.9602)

33. Peters J, Fowler E, Gatton M, Chen N, Saul A, Cheng Q. 2002 High diversity and rapid changeover of expressed *var* genes during the acute phase of *Plasmodium falciparum* infections in human volunteers. *Proc. Natl Acad. Sci. USA* **99**, 10 689–10 694. (doi:10.1073/pnas.162349899)

34. Borst P, Bitter W, McCulloch R, Van Leeuwen F, Rudenko G. 1995 Antigenic variation in malaria. *Cell* **82**, 1–4. (doi:10.1016/0092-8674(95)90044-6)

35. Miller LH, Good MF, Milon G. 1994 Malaria pathogenesis. *Science* **264**, 1878–1883. (doi:10.1126/science.8009217)

36. Recker M, Nee S, Bull PC, Kinyanjui S, Marsh K, Newbold C, Gupta S. 2004 Transient cross reactive immune responses can orchestrate antigenic variation in malaria. *Nature* **429**, 555–558. (doi:10.1038/nature02486)

37. Deitsch KW, Lukehart SA, Stringer JR. 2009 Common strategies for antigenic variation by bacterial, fungal and protozoan pathogens. *Nat. Rev. Microbiol.* **7**, 493–503. (doi:10.1038/nrmicro2145)

38. Claessens A, Hamilton WL, Kekre M, Otto TD, Faizullabhoy A, Rayner JC, Kwiatkowski D. 2014 Generation of antigenic diversity in *Plasmodium falciparum* by structured rearrangement of *Var* genes during mitosis. *PLoS Genet.* **10**, e1004812. (doi:10.1371/journal.pgen.1004812)

Bibliography

Aguiar, Joao C et al. (1992). "Agglutination of Plasmodium falciparum-infected erythrocytes from east and west African isolates by human sera from distant geographic regions". In: *The American journal of tropical medicine and hygiene* 47.5, pp. 621–632.

Albrecht, Letusa et al. (2006). "Extense variant gene family repertoire overlap in Western Amazon Plasmodium falciparum isolates". In: *Molecular and biochemical parasitology* 150.2, pp. 157–165.

Albrecht, Letusa et al. (2010). "The South American Plasmodium falciparum var gene repertoire is limited, highly shared and possibly lacks several antigenic types". In: *Gene* 453.1, pp. 37–44.

Artzy-Randrup, Yael et al. (2012). "Population structuring of multi-copy, antigen-encoding genes in Plasmodium falciparum". In: *Elife* 1, e00093.

Baird, J Kevin (1995). "Host age as a determinant of naturally acquired immunity to Plasmodium falciparum". In: *Parasitology today* 11.3, pp. 105–111.

Barry, Alyssa E et al. (2007). "Population genomics of the immune evasion (var) genes of Plasmodium falciparum". In: *PLoS pathogens* 3.3, e34.

Baruch, Dror I et al. (1997). "Identification of a region of PfEMP1 that mediates adherence of Plasmodium falciparum infected erythrocytes to CD36: conserved function with variant sequence". In: *Blood* 90.9, pp. 3766–3775.

Beadle, Christine et al. (1995). "Impact of transmission intensity and age on Plasmodium falciparum density and associated fever: implications for malaria vaccine trial design". In: *Journal of Infectious Diseases* 172.4, pp. 1047–1054.

Beck, H-P et al. (1997). "Analysis of multiple Plasmodium falciparum infections in Tanzanian children during the phase III trial of the malaria vaccine SPf66". In: *Journal of Infectious Diseases* 175.4, pp. 921–926.

Beeson, James G et al. (2001). "Parasite adhesion and immune evasion in placental malaria". In: *Trends in parasitology* 17.7, pp. 331–337.

Beier, John C, Gerry F Killeen, and John I Githure (1999). "entomologic inoculation rates and Plasmodium falciparum malaria prevalence in Africa." In: *The American Journal of Tropical Medicine and Hygiene* 61.1, pp. 109–113.

Bekessy, A, L Molineaux, and J Storey (1976). "Estimation of incidence and recovery rates of Plasmodium falciparum parasitaemia from longitudinal data". In: *Bulletin of the World Health Organization* 54.6, p. 685.

Bellan, Steve E (2010). "The importance of age dependent mortality and the extrinsic incubation period in models of mosquito-borne disease transmission and control". In: *PLoS One* 5.4, e10165.

Bopp, Selina ER et al. (2013). "Mitotic evolution of Plasmodium falciparum shows a stable core genome but recombination in antigen families". In: *PLoS genetics* 9.2, e1003293.

Borst, Piet et al. (1995). "Antigenic variation in malaria". In: *Cell* 82.1, pp. 1–4.

Boudin, Christian et al. (1993). "High human malarial infectivity to laboratory-bred Anopheles gambiae in a village in Burkina Faso". In: *The American journal of tropical medicine and hygiene* 48.5, pp. 700–706.

Boudin, Christian et al. (2004). "Plasmodium falciparum transmission blocking immunity under conditions of low and high endemicity in Cameroon". In: *Parasite immunology* 26.2, pp. 105–110.

Bousema, Teun and Chris Drakeley (2011). "Epidemiology and infectivity of Plasmodium falciparum and Plasmodium vivax gametocytes in relation to malaria control and elimination". In: *Clinical microbiology reviews* 24.2, pp. 377–410.

Bousema, Teun et al. (2010). "Revisiting the circulation time of Plasmodium falciparum gametocytes: molecular detection methods to estimate the duration of gametocyte carriage and the effect of gametocytocidal drugs". In: *Malaria journal* 9.1, p. 136.

Bousema, Teun et al. (2011). "Human immune responses that reduce the transmission of Plasmodium falciparum in African populations". In: *International journal for parasitology* 41.3, pp. 293–300.

Bruce, MC et al. (2000). "Age- and species-specific duration of infection in asymptomatic malaria infections in Papua New Guinea". In: *Parasitology* 121.3, pp. 247–256.

Buckee, Caroline O, Peter C Bull, and Sunetra Gupta (2009). "Inferring malaria parasite population structure from serological networks". In: *Proceedings of the Royal Society of London B: Biological Sciences* 276.1656, pp. 477–485.

Buckee, Caroline O and Mario Recker (2012). "Evolution of the multi-domain structures of virulence genes in the human malaria parasite, Plasmodium falciparum". In: *PLoS computational biology* 8.4, e1002451.

Bull, PC et al. (2000). "Plasmodium falciparum-infected erythrocytes: agglutination by diverse Kenyan plasma is associated with severe disease and young host age". In: *The Journal of infectious diseases* 182.1, pp. 252–259.

Bull, Peter C and Abdirahman I Abdi (2016). "The role of PfEMP1 as targets of natu-
rally acquired immunity to childhood malaria: prospects for a vaccine". In: *Para-
sitology* 143.2, pp. 171–186.

Bull, Peter C et al. (1998). "Parasite antigens on the infected red cell surface are tar-
gets for naturally acquired immunity to malaria". In: *Nature medicine* 4.3, pp. 358–
360.

Bull, Peter C et al. (1999). "Antibody recognition of Plasmodium falciparum ery-
throcyte surface antigens in Kenya: evidence for rare and prevalent variants". In:
Infection and Immunity 67.2, pp. 733–739.

Bull, Peter C et al. (2002). "Plasmodium falciparum infections are associated with
agglutinating antibodies to parasite-infected erythrocyte surface antigens among
healthy Kenyan children". In: *Journal of Infectious Diseases* 185.11, pp. 1688–1691.

Bull, Peter C et al. (2005). "Plasmodium falciparum variant surface antigen expres-
sion patterns during malaria". In: *PLoS pathogens* 1.3, e26.

Bull, Peter C et al. (2008). "Plasmodium falciparum antigenic variation. Mapping
mosaic var gene sequences onto a network of shared, highly polymorphic se-
quence blocks". In: *Molecular microbiology* 68.6, pp. 1519–1534.

Charlwood, JD et al. (2003). "Mating does not affect the biting behaviour of Anophe-
les gambiae from the islands of Sao Tome and Principe, West Africa". In: *Annals
of Tropical Medicine & Parasitology* 97.7, pp. 751–756.

Chattopadhyay, Rana et al. (2003). "Plasmodium falciparum infection elicits both
variant-specific and cross-reactive antibodies against variant surface antigens".
In: *Infection and immunity* 71.2, pp. 597–604.

Chen, Donald S et al. (2011). "A molecular epidemiological study of var gene di-
versity to characterize the reservoir of Plasmodium falciparum in humans in
Africa". In: *PLoS One* 6.2, e16629.

Chen, Qijun, Martha Schlichtherle, and Mats Wahlgren (2000). "Molecular aspects of severe malaria". In: *Clinical microbiology reviews* 13.3, pp. 439–450.

Claessens, Antoine et al. (2012). "A subset of group A-like var genes encodes the malaria parasite ligands for binding to human brain endothelial cells". In: *Proceedings of the National Academy of Sciences* 109.26, E1772–E1781.

Claessens, Antoine et al. (2014). "Generation of antigenic diversity in Plasmodium falciparum by structured rearrangement of var genes during mitosis". In: *PLoS genetics* 10.12, e1004812.

Clausen, Thomas M et al. (2012). "Structural and functional insight into how the Plasmodium falciparum VAR2CSA protein mediates binding to chondroitin sulfate A in placental malaria". In: *Journal of Biological Chemistry* 287.28, pp. 23332–23345.

Coleman, PG, BD Perry, and MEJ Woolhouse (2001). "Endemic stability - a veterinary idea applied to human public health". In: *The Lancet* 357.9264, pp. 1284–1286.

Coleman, Russell E et al. (2004). "Infectivity of asymptomatic Plasmodium-infected human populations to Anopheles dirus mosquitoes in western Thailand". In: *Journal of medical entomology* 41.2, pp. 201–208.

Conway, David J et al. (1999). "High recombination rate in natural populations of Plasmodium falciparum". In: *Proceedings of the National Academy of Sciences* 96.8, pp. 4506–4511.

Craig, Alister and Artur Scherf (2001). "Molecules on the surface of the Plasmodium falciparum infected erythrocyte and their role in malaria pathogenesis and immune evasion". In: *Molecular and biochemical parasitology* 115.2, pp. 129–143.

Day, Karen P et al. (2017). "Evidence of strain structure in Plasmodium falciparum var gene repertoires in children from Gabon, West Africa". In: *Proceedings of the National Academy of Sciences*, E4103–E4111.

Dietz, K, L Molineaux, and A Thomas (1974). "A malaria model tested in the African savannah". In: *Bulletin of the World Health Organization* 50.3-4, p. 347.

Doolan, Denise L, Carlota Dobaño, and J Kevin Baird (2009). "Acquired immunity to malaria". In: *Clinical microbiology reviews* 22.1, pp. 13–36.

Drakeley, Chris et al. (2006a). "The epidemiology of Plasmodium falciparum gametocytes: weapons of mass dispersion". In: *Trends in parasitology* 22.9, pp. 424–430.

Drakeley, CJ et al. (2000). "Estimates of the infectious reservoir of Plasmodium falciparum malaria in The Gambia and in Tanzania". In: *Transactions of the Royal Society of tropical Medicine and Hygiene* 94.5, pp. 472–476.

Drakeley, CJ et al. (2006b). "Transmission-reducing immunity is inversely related to age in Plasmodium falciparum gametocyte carriers". In: *Parasite immunology* 28.5, pp. 185–190.

Druilhe, P and JL Perignon (1997). "A hypothesis about the chronicity of malaria infection". In: *Parasitology Today* 13.9, pp. 353–357.

Ebenezer, Amawulu, Aline Edith M Noutcha, and Samuel N Okiwelu (2016). "Relationship of annual entomological inoculation rates to malaria transmission indices, Bayelsa State, Nigeria". In: *Journal of vector borne diseases* 53.1, p. 46.

Eckhoff, Philip A (2012). "Malaria parasite diversity and transmission intensity affect development of parasitological immunity in a mathematical model". In: *Malaria journal* 11.1, p. 419.

Eichner, Martin et al. (2001). "Genesis, sequestration and survival of Plasmodium falciparum gametocytes: parameter estimates from fitting a model to malariatherapy data". In: *Transactions of the Royal Society of Tropical Medicine and Hygiene* 95.5, pp. 497–501.

Felger, Ingrid et al. (2012). "The dynamics of natural Plasmodium falciparum infections". In: *PloS one* 7.9, e45542.

Ferguson, Heather M and Andrew F Read (2004). "Mosquito appetite for blood is stimulated by Plasmodium chabaudi infections in themselves and their vertebrate hosts". In: *Malaria journal* 3.1, p. 12.

Forsyth, Karen P et al. (1989). "Diversity of antigens expressed on the surface of erythrocytes infected with mature Plasmodium falciparum parasites in Papua New Guinea". In: *The American journal of tropical medicine and hygiene* 41.3, pp. 259–265.

Fowkes, Freya JI et al. (2010). "The relationship between anti-merozoite antibodies and incidence of Plasmodium falciparum malaria: A systematic review and meta-analysis". In: *PLoS medicine* 7.1, e1000218.

Frank, Matthias et al. (2007). "Variable switching rates of malaria virulence genes are associated with chromosomal position". In: *Molecular microbiology* 64.6, pp. 1486–1498.

Freitas-Junior, Lúcio H et al. (2000). "Frequent ectopic recombination of virulence factor genes in telomeric chromosome clusters of P. falciparum". In: *Nature* 407.6807, p. 1018.

Gamain, Benoit, Louis H Miller, and Dror I Baruch (2001). "The surface variant antigens of Plasmodium falciparum contain cross-reactive epitopes". In: *Proceedings of the National Academy of Sciences* 98.5, pp. 2664–2669.

Gardner, Malcolm J et al. (2002). "Genome sequence of the human malaria parasite Plasmodium falciparum". In: *Nature* 419.6906.

Ghani, Azra C et al. (2009). "Loss of population levels of immunity to malaria as a result of exposure-reducing interventions: consequences for interpretation of disease trends". In: *PLoS One* 4.2, e4383.

Giha, HA et al. (1999a). "Overlapping antigenic repertoires of variant antigens expressed on the surface of erythrocytes infected by Plasmodium falciparum". In: *Parasitology* 119.1, pp. 7–17.

Giha, Haider A et al. (1999b). "Nine-year longitudinal study of antibodies to variant antigens on the surface of Plasmodium falciparum-infected erythrocytes". In: *Infection and immunity* 67.8, pp. 4092–4098.

Githeko, AK et al. (1992). "The reservoir of Plasmodium falciparum malaria in a holoendemic area of western Kenya". In: *Transactions of the Royal Society of Tropical Medicine and Hygiene* 86.4, pp. 355–358.

Graves, Patricia M et al. (1988). "Measurement of malarial infectivity of human populations to mosquitoes in the Madang area, Papua New Guinea". In: *Parasitology* 96.2, pp. 251–263.

Guerra, Carlos A et al. (2008). "The limits and intensity of Plasmodium falciparum transmission: implications for malaria control and elimination worldwide". In: *PLoS medicine* 5.2, e38.

Gupta, S and RM Anderson (1999). "Population structure of pathogens: the role of immune selection". In: *Parasitology Today* 15.12, pp. 497–501.

Gupta, S and KP Day (1994). "A strain theory of malaria transmission". In: *Parasitology Today* 10.12, pp. 476–481.

Gupta, Sunetra, Neil Ferguson, and Roy Anderson (1998). "Chaos, persistence, and evolution of strain structure in antigenically diverse infectious agents". In: *Science* 280.5365, pp. 912–915.

Gupta, Sunetra, Jonathan Swinton, and Roy M Anderson (1994). "Theoretical studies of the effects of heterogeneity in the parasite population on the transmission dynamics of malaria". In: *Proceedings of the Royal Society of London B: Biological Sciences* 256.1347, pp. 231–238.

Gupta, Sunetra et al. (1994). "Parasite virulence and disease patterns in Plasmodium falciparum malaria". In: *Proceedings of the National Academy of Sciences* 91.9, pp. 3715–3719.

Gupta, Sunetra et al. (1996). "The maintenance of strain structure in populations of recombining infectious agents". In: *Nature medicine* 2.4, pp. 437–442.

Gupta, Sunetra et al. (1999). "Immunity to non-cerebral severe malaria is acquired after one or two infections". In: *Nature medicine* 5.3, pp. 340–343.

Hoffman, Stephen L et al. (2002). "Protection of humans against malaria by immunization with radiation-attenuated Plasmodium falciparum sporozoites". In: *The Journal of infectious diseases* 185.8, pp. 1155–1164.

Holding, Thomas and Mario Recker (2015). "Maintenance of phenotypic diversity within a set of virulence encoding genes of the malaria parasite Plasmodium falciparum". In: *Journal of The Royal Society Interface* 12.113, p. 20150848.

Horrocks, Paul et al. (2004). "Variable var transition rates underlie antigenic variation in malaria". In: *Proceedings of the National Academy of Sciences of the United States of America* 101.30, pp. 11129–11134.

Howell, Dasein P-G et al. (2008). "Mapping a common interaction site used by Plasmodium falciparum Duffy binding-like domains to bind diverse host receptors". In: *Molecular microbiology* 67.1, pp. 78–87.

Hviid, Lars (2004). "The immuno-epidemiology of pregnancy-associated Plasmodium falciparum malaria: a variant surface antigen-specific perspective". In: *Parasite immunology* 26.11-12, pp. 477–486.

Hviid, Lars and Trine Staalsoe (2004). "Malaria immunity in infants: a special case of a general phenomenon?" In: *Trends in parasitology* 20.2, pp. 66–72.

Janes, Joel H et al. (2011). "Investigating the host binding signature on the Plasmodium falciparum PfEMP1 protein family". In: *PLoS pathogens* 7.5, e1002032.

Jensen, Anja TR et al. (2004). "Plasmodium falciparum associated with severe childhood malaria preferentially expresses PfEMP1 encoded by group A var genes". In: *Journal of Experimental Medicine* 199.9, pp. 1179–1190.

Jiang, Hongying et al. (2011). "High recombination rates and hotspots in a Plasmodium falciparum genetic cross". In: *Genome biology* 12.4, R33.

Juliano, Jonathan J et al. (2009). "Misclassification of drug failure in Plasmodium falciparum clinical trials in Southeast Asia". In: *The Journal of infectious diseases* 200.4, pp. 624–628.

Juliano, Jonathan J et al. (2010). "Exposing malaria in-host diversity and estimating population diversity by capture-recapture using massively parallel pyrosequencing". In: *Proceedings of the National Academy of Sciences* 107.46, pp. 20138–20143.

Kaestli, Mirjam et al. (2006). "Virulence of malaria is associated with differential expression of Plasmodium falciparum var gene subgroups in a case-control study". In: *The Journal of infectious diseases* 193.11, pp. 1567–1574.

Kitua, Andrew Y et al. (1996). "Plasmodium falciparum malaria in the first year of life in an area of intense and perennial transmission". In: *Tropical Medicine & International Health* 1.4, pp. 475–484.

Klein, EY (2013). "Antimalarial drug resistance: a review of the biology and strategies to delay emergence and spread". In: *International journal of antimicrobial agents* 41.4, pp. 311–317.

Klowden, Marc J and Hans Briegel (1994). "Mosquito gonotrophic cycle and multiple feeding potential: contrasts between Anopheles and Aedes (Diptera: Culicidae)". In: *Journal of Medical Entomology* 31.4, pp. 618–622.

Koella, Jacob C, Flemming L Sörensen, and RA Anderson (1998). "The malaria parasite, Plasmodium falciparum, increases the frequency of multiple feeding of its mosquito vector, Anopheles gambiae". In: *Proceedings of the Royal Society of London B: Biological Sciences* 265.1398, pp. 763–768.

Kraemer, Susan M and Joseph D Smith (2003). "Evidence for the importance of genetic structuring to the structural and functional specialization of the Plasmodium falciparum var gene family". In: *Molecular microbiology* 50.5, pp. 1527–1538.

— (2006). "A family affair: var genes, PfEMP1 binding, and malaria disease". In: *Current opinion in microbiology* 9.4, pp. 374–380.

Kraemer, Susan M et al. (2007). "Patterns of gene recombination shape var gene repertoires in Plasmodium falciparum: comparisons of geographically diverse isolates". In: *BMC genomics* 8.1, p. 45.

Kyes, Sue A, Susan M Kraemer, and Joseph D Smith (2007). "Antigenic variation in Plasmodium falciparum: gene organization and regulation of the var multigene family". In: *Eukaryotic Cell* 6.9, pp. 1511–1520.

Kyriacou, Helen M et al. (2006). "Differential var gene transcription in Plasmodium falciparum isolates from patients with cerebral malaria compared to hyperparasitaemia". In: *Molecular and biochemical parasitology* 150.2, pp. 211–218.

Langhorne, Jean et al. (2008). "Immunity to malaria: more questions than answers". In: *Nature immunology* 9.7, pp. 725–732.

Lavstsen, Thomas et al. (2003). "Sub-grouping of Plasmodium falciparum 3D7 var genes based on sequence analysis of coding and non-coding regions". In: *Malaria Journal* 2.1, p. 27.

Macdonald, G (1956). "Epidemiological basis of malaria control". In: *Bulletin of the World Health Organization* 15.3-5, p. 613.

Macdonald, George et al. (1957). "The epidemiology and control of malaria." In: *The Epidemiology and Control of Malaria.*

Marsh, K and S Kinyanjui (2006). "Immune effector mechanisms in malaria". In: *Parasite immunology* 28.1-2, pp. 51–60.

Marsh, Kevin and Russell J Howard (1986). "Antigens induced on erythrocytes by P. falciparum: expression of diverse and conserved determinants". In: *Science* 231.4734, pp. 150–153.

Mayor, Alfredo et al. (2003). "Plasmodium falciparum multiple infections in Mozambique, its relation to other malariological indices and to prospective risk of malaria morbidity". In: *Tropical Medicine & International Health* 8.1, pp. 3–11.

Miller, Louis H et al. (2002). "The pathogenic basis of malaria". In: *Nature* 415.6872, pp. 673–679.

Mo, Min et al. (2008). "The C-terminal segment of the cysteine-rich interdomain of Plasmodium falciparum erythrocyte membrane protein 1 determines CD36 binding and elicits antibodies that inhibit adhesion of parasite-infected erythrocytes". In: *Infection and immunity* 76.5, pp. 1837–1847.

Molineaux, Louis, Gabriele Gramiccia, World Health Organization, et al. (1980). "The Garki project: research on the epidemiology and control of malaria in the Sudan savanna of West Africa". In:

Morlais, Isabelle et al. (2015). "Plasmodium falciparum mating patterns and mosquito infectivity of natural isolates of gametocytes". In: *PloS one* 10.4, e0123777.

Mu, Jianbing et al. (2005). "Recombination hotspots and population structure in Plasmodium falciparum". In: *PLoS biology* 3.10, e335.

Newbold, CI et al. (1992). "Plasmodium falciparum: the human agglutinating antibody response to the infected red cell surface is predominantly variant specific". In: *Experimental parasitology* 75.3, pp. 281–292.

Nielsen, Morten A et al. (2002). "Plasmodium falciparum variant surface antigen expression varies between isolates causing severe and nonsevere malaria and is modified by acquired immunity". In: *The Journal of Immunology* 168.7, pp. 3444–3450.

Nielsen, Morten A et al. (2004). "Geographical and temporal conservation of antibody recognition of Plasmodium falciparum variant surface antigens". In: *Infection and immunity* 72.6, pp. 3531–3535.

Ntoumi, Francine et al. (1995). "Age-dependent carriage of multiple Plasmodium falciparum merozoite surface antigen-2 alleles in asymptomatic malaria infections". In: *The American journal of tropical medicine and hygiene* 52.1, pp. 81–88.

Nwakanma, Davis et al. (2008). "High gametocyte complexity and mosquito infectivity of Plasmodium falciparum in the Gambia". In: *International journal for parasitology* 38.2, pp. 219–227.

Ofori, Michael F et al. (2002). "Malaria-induced acquisition of antibodies to Plasmodium falciparum variant surface antigens". In: *Infection and immunity* 70.6, pp. 2982–2988.

Okiro, Emelda A et al. (2009). "Age patterns of severe paediatric malaria and their relationship to Plasmodium falciparum transmission intensity". In: *Malaria journal* 8.1, p. 4.

Peyerl-Hoffmann, G et al. (2001). "Genetic diversity of Plasmodium falciparum and its relationship to parasite density in an area with different malaria endemicities in West Uganda". In: *Tropical Medicine & International Health* 6.8, pp. 607–613.

Rask, Thomas S et al. (2010). "Plasmodium falciparum erythrocyte membrane protein 1 diversity in seven genomes - divide and conquer". In: *PLoS computational biology* 6.9, e1000933.

Recker, Mario, Nimalan Arinaminpathy, and Caroline O Buckee (2008). "The effects of a partitioned var gene repertoire of Plasmodium falciparum on antigenic diversity and the acquisition of clinical immunity". In: *Malaria journal* 7.1, p. 18.

Recker, Mario and Sunetra Gupta (2005). "A model for pathogen population structure with cross-protection depending on the extent of overlap in antigenic variant repertoires". In: *Journal of theoretical biology* 232.3, pp. 363–373.

Recker, Mario et al. (2004). "Transient cross-reactive immune responses can orchestrate antigenic variation in malaria". In: *Nature* 429.6991, p. 555.

Recker, Mario et al. (2011). "Antigenic variation in Plasmodium falciparum malaria involves a highly structured switching pattern". In: *PLoS Pathogens* 7.3, e1001306.

Reyburn, Hugh et al. (2005). "Association of transmission intensity and age with clinical manifestations and case fatality of severe Plasmodium falciparum malaria". In: *Jama* 293.12, pp. 1461–1470.

Rich, Stephen M, Marcelo Urbano Ferreira, and Francisco J Ayala (2000). "The origin of antigenic diversity in Plasmodium falciparum". In: *Parasitology Today* 16.9, pp. 390–396.

Rich, Stephen M, Richard R Hudson, and Francisco J Ayala (1997). "Plasmodium falciparum antigenic diversity: evidence of clonal population structure". In: *Proceedings of the National Academy of Sciences* 94.24, pp. 13040–13045.

Robinson, Bridget A, Teresa L Welch, and Joseph D Smith (2003). "Widespread functional specialization of Plasmodium falciparum erythrocyte membrane protein 1 family members to bind CD36 analysed across a parasite genome". In: *Molecular microbiology* 47.5, pp. 1265–1278.

Ross, Amanda, Gerry Killeen, and Thomas Smith (2006). "Relationships between host infectivity to mosquitoes and asexual parasite density in Plasmodium falciparum". In: *The American journal of tropical medicine and hygiene* 75.2_suppl, pp. 32–37.

Ross, Ronald (1905). "The Logical Basis of the Sanitary Policy of Mosquito Reduction". In: *Science* 22, pp. 689–699.

— (1908). "The prevention of malaria in British possessions, Egypt, and parts of America". In: *Journal of the Royal Army Medical Corps* 10.3, pp. 263–280.

— (1911a). "Some quantitative studies in epidemiology". In: *Nature* 87.2188, pp. 466–467.

— (1911b). *The prevention of malaria*. John Murray; London.

Rowe, Alexandra et al. (1995). "Plasmodium falciparum rosetting is associated with malaria severity in Kenya." In: *Infection and immunity* 63.6, pp. 2323–2326.

Rowe, J Alexandra et al. (2009). "Adhesion of Plasmodium falciparum-infected erythrocytes to human cells: molecular mechanisms and therapeutic implications". In: *Expert reviews in molecular medicine* 11.

Salanti, Ali et al. (2003). "Selective upregulation of a single distinctly structured var gene in chondroitin sulphate A-adhering Plasmodium falciparum involved in pregnancy-associated malaria". In: *Molecular microbiology* 49.1, pp. 179–191.

Salanti, Ali et al. (2004). "Evidence for the involvement of VAR2CSA in pregnancy-associated malaria". In: *Journal of Experimental Medicine* 200.9, pp. 1197–1203.

Sama, Wilson, Gerry Killeen, and Tom Smith (2004). "Estimating the duration of Plasmodium falciparum infection from trials of indoor residual spraying". In: *The American journal of tropical medicine and hygiene* 70.6, pp. 625–634.

Scherf, Artur, Jose Juan Lopez-Rubio, and Loïc Riviere (2008). "Antigenic variation in Plasmodium falciparum". In: *Annual Review of Microbiology* 62, pp. 445–470.

Schneider, Petra et al. (2007). "Submicroscopic Plasmodium falciparum gametocyte densities frequently result in mosquito infection". In: *The American journal of tropical medicine and hygiene* 76.3, pp. 470–474.

Shaukat, Ayesha M, Joel G Breman, and F Ellis McKenzie (2010). "Using the entomological inoculation rate to assess the impact of vector control on malaria parasite transmission and elimination". In: *Malaria journal* 9.1, p. 122.

Sherman, Irwin W, Shigetoshi Eda, and Enrique Winograd (2003). "Cytoadherence and sequestration in Plasmodium falciparum: defining the ties that bind". In: *Microbes and Infection* 5.10, pp. 897–909.

Smalley, ME and RE Sinden (1977). "Plasmodium falciparum gametocytes: their longevity and infectivity". In: *Parasitology* 74.1, pp. 1–8.

Smith, David L et al. (2012). "Ross, Macdonald, and a theory for the dynamics and control of mosquito-transmitted pathogens". In: *PLoS pathogens* 8.4, e1002588.

Smith, DL et al. (2005). "The entomological inoculation rate and Plasmodium falciparum infection in African children". In: *Nature* 438.7067, p. 492.

Smith, Joseph D (2014). "The role of PfEMP1 adhesion domain classification in Plasmodium falciparum pathogenesis research". In: *Molecular and biochemical parasitology* 195.2, pp. 82–87.

Smith, Joseph D et al. (2000). "Classification of adhesive domains in the Plasmodium falciparum erythrocyte membrane protein 1 family". In: *Molecular and biochemical parasitology* 110.2, pp. 293–310.

Smith, T et al. (1999a). "11. Premunition in Plasmodium falciparum infection: insights from the epidemiology of multiple infections". In: *Transactions of the Royal Society of Tropical Medicine and Hygiene* 93.Supplement_1, pp. 59–64.

Smith, T et al. (1999b). "4. Age dependence of the multiplicity of Plasmodium falciparum infections and of other malariological indices in an area of high endemicity". In: *Transactions of the Royal Society of Tropical Medicine and Hygiene* 93, pp. 15–20.

Smith, T et al. (1999c). "Dynamics of multiple Plasmodium falciparum infections in infants in a highly endemic area of Tanzania". In: *Transactions of the Royal Society of Tropical Medicine and Hygiene* 93, pp. 35–39.

Smith, Tom and Penelope Vounatsou (2003). "Estimation of infection and recovery rates for highly polymorphic parasites when detectability is imperfect, using hidden Markov models". In: *Statistics in medicine* 22.10, pp. 1709–1724.

Snow, Robert W et al. (1997). "Relation between severe malaria morbidity in children and level of Plasmodium falciparum transmission in Africa". In: *The Lancet* 349.9066, pp. 1650–1654.

Stockdale, Chris et al. (2008). "Antigenic variation in Trypanosoma brucei: joining the DOTs". In: *PLoS biology* 6.7, e185.

Storn, Rainer and Kenneth Price (1997). "Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces". In: *Journal of global optimization* 11.4, pp. 341–359.

Struchiner, Claudio J, M Elizabeth Halloran, and Andrew Spielman (1989). "Modeling malaria vaccines I: new uses for old ideas". In: *Mathematical Biosciences* 94.1, pp. 87–113.

Taylor, Helen M, Susan A Kyes, and Christopher I Newbold (2000). "Var gene diversity in Plasmodium falciparum is generated by frequent recombination events". In: *Molecular and biochemical parasitology* 110.2, pp. 391–397.

Taylor, LH and AF Read (1997). "Why so few transmission stages? Reproductive restraint by malaria parasites". In: *Parasitology Today* 13.4, pp. 135–140.

Trape, Jean-François and Christophe Rogier (1996). "Combating malaria morbidity and mortality by reducing transmission". In: *Parasitology today* 12.6, pp. 236–240.

Tuikue Ndam, Nicaise et al. (2017). "Parasites Causing Cerebral Falciparum Malaria Bind Multiple Endothelial Receptors and Express EPCR and ICAM-1-Binding PfEMP1". In: *The Journal of Infectious Diseases*.

Udeinya, Iroka J et al. (1983). "Plasmodium falciparum strain-specific antibody blocks binding of infected erythrocytes to amelanotic melanoma cells". In: *Nature* 303.5916, pp. 429–431.

Vigan-Womas, Inès et al. (2008). "An in vivo and in vitro model of Plasmodium falciparum rosetting and autoagglutination mediated by varO, a group A var gene encoding a frequent serotype". In: *Infection and immunity* 76.12, pp. 5565–5580.

Wang, Christian W et al. (2009). "The Plasmodium falciparum var gene transcription strategy at the onset of blood stage infection in a human volunteer". In: *Parasitology international* 58.4, pp. 478–480.

Ward, Christopher P et al. (1999). "Analysis of Plasmodium falciparum PfEMP-1/var genes suggests that recombination rearranges constrained sequences". In: *Molecular and biochemical parasitology* 102.1, pp. 167–177.

Warimwe, George M et al. (2016). "Serological conservation of parasite-infected erythrocytes predicts Plasmodium falciparum erythrocyte membrane protein 1 gene expression but not severity of childhood malaria". In: *Infection and immunity* 84.5, pp. 1331–1335.

World Health Organization (2016). *World Malaria Report 2016*. World Health Organization.