



OPEN

Dynamic causal modelling of immune heterogeneity

Thomas Parr¹✉, Anjali Bhat¹, Peter Zeidman¹, Aimee Goel², Alexander J. Billig³, Rosalyn Moran⁴ & Karl J. Friston¹

An interesting inference drawn by some COVID-19 epidemiological models is that there exists a proportion of the population who are not susceptible to infection—even at the start of the current pandemic. This paper introduces a model of the immune response to a virus. This is based upon the same sort of mean-field dynamics as used in epidemiology. However, in place of the location, clinical status, and other attributes of people in an epidemiological model, we consider the state of a virus, B and T-lymphocytes, and the antibodies they generate. Our aim is to formalise some key hypotheses as to the mechanism of resistance. We present a series of simple simulations illustrating changes to the dynamics of the immune response under these hypotheses. These include attenuated viral cell entry, pre-existing cross-reactive humoral (antibody-mediated) immunity, and enhanced T-cell dependent immunity. Finally, we illustrate the potential application of this sort of model by illustrating variational inversion (using simulated data) of this model to illustrate its use in testing hypotheses. In principle, this furnishes a fast and efficient immunological assay—based on sequential serology—that provides a (1) quantitative measure of latent immunological responses and (2) a Bayes optimal classification of the different kinds of immunological response (c.f., glucose tolerance tests used to test for insulin resistance). This may be especially useful in assessing SARS-CoV-2 vaccines.

Recently, we produced a series of technical reports applying dynamic causal modelling (DCM) to epidemiological data to illustrate how this might be applied to understanding the COVID-19 pandemic^{1–4}. This is an approach developed in neuroimaging research that attempts to find the best explanations for observed timeseries data⁵. An interesting outcome of this work is that the best explanations of data from the current pandemic appeal to there being sizeable portions of the population who either are not susceptible to the virus, or who have such a mild course of illness that they are unlikely to pass it on. The term ‘immunological dark matter’ has found some media purchase⁶ in describing this finding. However, the idea behind this phrase is simple. The implication is that one or more unmeasured variables confer a degree of non-susceptibility or resistance to severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). This may be partly due to demographic factors, such as geographical isolation or shielding. In addition, it is likely that heterogeneity in the immune response people mount to the virus contributes to this^{7–11}.

This raises the question as to what the differences in immunity may be and how we could test this. In this paper, we formalise some of the hypotheses raised as to the immunological factors conferring resistance. We have no pretensions of answering this question—which almost certainly has several different concurrently correct answers—but hope to illustrate how the variational methods used in the epidemiological DCM could be repurposed to test hypotheses about the causes of resistance. This may be particularly salient in the setting of vaccine trials, in which a great deal of immunological data may be collected following a timed exposure to viral antigens—and in which the immunity conferred by a vaccine may transcend purely antibody mediated immunity^{12,13}, which appears to have a relatively low prevalence¹⁴. As such, the model presented here is designed to include antibody mediated immunity, but also to allow for alternative (e.g., T-cell mediated) forms of resistance. The use of variational inference procedures in this setting allows for fast and efficient inferences at the individual level, quantification of the uncertainty surrounding these estimates, and hypothesis testing as to the best explanation for individual serology results.

The DCM approach relies upon the notion of a generative or forward model that generates measurable data. The idea here is that, by adjusting the parameters of this model, we can find the parameter configurations that minimise any discrepancy between synthetic and measured data, subject to prior beliefs about the plausibility of those configurations. Practically, this proceeds using variational Bayesian inference^{15,16}. We make use of a

¹Wellcome Centre for Human Neuroimaging, Queen Square Institute of Neurology, London, UK. ²Royal Stoke University Hospital, Stoke-on-Trent, UK. ³UCL Ear Institute, University College London, London, UK. ⁴Centre for Neuroimaging Science, Department of Neuroimaging, IoPPN, King's College London, London, UK. ✉email: thomas.parr.12@ucl.ac.uk

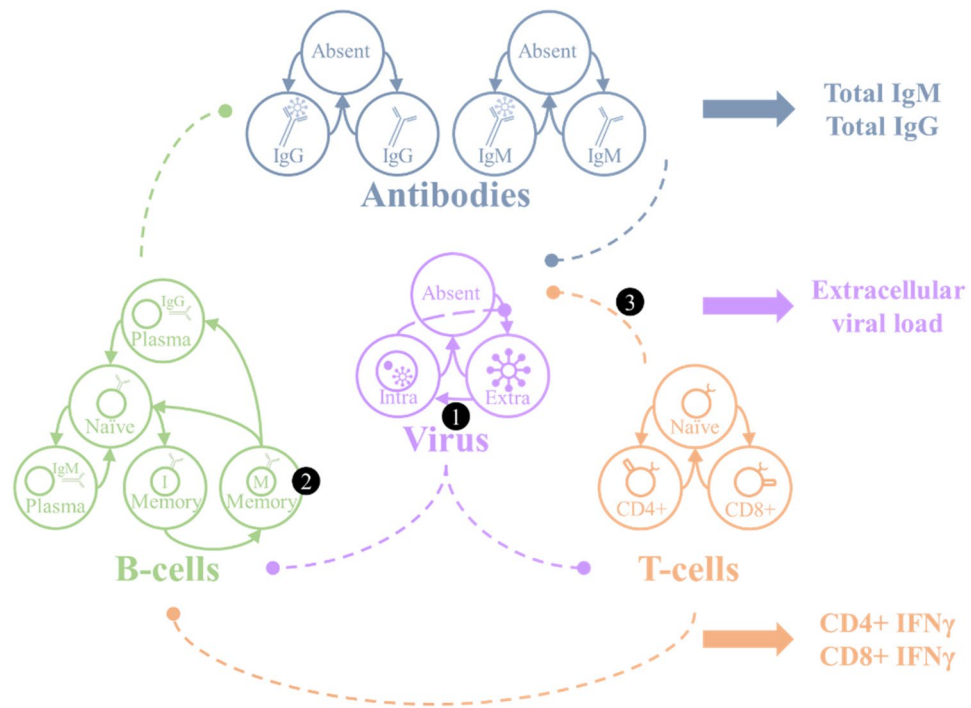


Figure 1. The generative model. This schematic sets out the five factors of the generative model and the data modalities they generate. There are two antibody factors (IgG and IgM). For these factors, the absent antibody can become either neutralising or non-neutralising, and each of these can decay back to being absent. Similarly, the T-cell factor allows for transitions between being naïve to either CD4+ or CD8+ cells. These decay back to naïve, under the assumption that the production of naïve T-cells occurs at approximately the same rate as the death rate of specialised T-cells to maintain a roughly constant T-cell population. The viral dynamics include an absent state, which transitions to being extracellular depending upon the amount of intracellular virus. This represents the intracellular replication and shedding of virus into the extracellular compartment. The extracellular virus then moves into the intracellular compartment—where further replication and shedding can take place. The B-cell factor is slightly more complex as this includes five different levels. Naïve B-cells become either IgM plasma cells or immature memory cells. Immature (I) memory cells develop into mature (M) memory cells, which themselves can become IgG plasma cells. Mature memory cells and plasma cells decay into naïve cells over time. The dashed connections between factors with round arrows express the directed influences between factors, as described in more detail in the main text. The influence of one factor over another is to increase or decrease the rates of transition between selected levels of that factor. The three numbers in black circles indicate the parameters we will manipulate to augment or attenuate the immune response.

variational Laplace procedure¹⁷ that maximises a lower bound on the marginal likelihood (or model evidence) of data. The bound is known variously as an evidence lower bound (ELBO) or negative free energy. The benefit of the forward modelling approach is that it appeals directly to the mechanisms that generate data. We are interested in measures of immunity that include the viral load, antibody levels, and T-cell responses (as indicated by interferon- γ secretion by CD4 and CD8 positive T-cells). This approach allows us to make inferences about immunity that can arise from interacting and multifaceted immune responses.

In the following sections, we first set out a generative model of immune responses. Like many epidemiological models¹⁸—and existing immunological models¹⁹—this is a compartment model. This means each component of the immune system exists in one of several alternative states that they may transition between. As in the location, infection, symptom, testing (LIST) DCM¹, our generative model appeals to different factors, corresponding to different arms of the immune system. After introducing this model, we set out formalisations of hypotheses as to the causes of resistance. We then illustrate the fitting of this model to (synthetic) data highlighting how these alternative hypotheses could be tested for a specific resistant individual. Finally, we discuss the potential implications of this approach—including the importance of not overinterpreting negative antibody tests²⁰.

A mean-field model of the immune response

The approach we pursue here is based upon a mean-field approximation that separates an immune response into five different factors or modules^{21–23}. These are the antibodies (which subdivide into IgG and IgM types), B-cells, T-cells, and the virus itself. Figure 1 summarises the key aspects of the model dynamics. Each factor has a number of levels, with some probability assigned to each level. The dynamics of the model are formulated in terms of the move of probability mass between these levels. For example, antibody synthesis is the process of moving probability mass from the ‘absent’ level to the neutralising or non-neutralising antibody states. The

dashed arrows between factors in Fig. 1 show the direction of influence of factors on one another. We will unpack the model by going through these, starting from the reciprocal influence of the virus on T-cells, followed by T-cell mediation of the B-cell response. We then outline the interaction between B-cells and the virus, B-cells and antibodies, and antibodies and virus.

The influence of the virus on the T-cells is to increase the rate at which naïve T-cells specialise²⁴. Some proportion (θ_{CD4}) of naïve cells transition to being CD4+ cells each hour, scaled by the probability of the virus being extracellular. Similarly, some proportion (θ_{CD8}) of naïve cells transition to being CD8+ cells each hour, scaled by the probability of the virus being intracellular. Our definition of the term ‘naïve’ T-cells is deliberately broad, to account for any T-cell that has the potential to become an activated CD4+ or CD8+ cell selective for a viral antigen. This might include CD4+ or CD8+ cells that have not become activated or started to secrete cytokines like interferon- γ (INF- γ), in addition to cells that have yet to express CD4 or CD8 surface markers²⁵.

The influence of the T-cells on the viral factor are as follows. In proportion to the number of CD4+ T-cells, the extracellular virus moves to the absent state. This stands in for the effect of CD4+ cell mediated activation of parts of the innate immune system—for example, the activation of macrophages that phagocytose and kill pathogens²⁶. Responses mediated by these cells are known to have a role in influenza virus and coronavirus infections^{27,28}. The CD8+ cell causes a move of intracellular virus to the absent state. This represents the cytotoxic activity of these cells, prompting apoptosis of infected cells²⁹. The proportion moving to the absent state is given by the probability of a T-cell being a CD8+ cell multiplied by $0.5(\theta_{TCP} + \theta_{dpi})$, where the first parameter deals with the effect of the cytotoxic (CD8+) cell on the pathogen, and the second is the decay of the intracellular pathogen due to T-cell independent mechanisms. The latter is invariant to the T-cell state. We have not specified the mechanism of decay of intracellular and extracellular virus here, such that these variables could encompass innate immunity, non-functional viral mutations, and host cell death^{30,31}.

In addition to their direct effects on the virus, the T-cells promote a B-cell response³². Specifically, in the presence of CD4+ cells, naïve B-cells become (immature) memory cells³³ or IgM-secreting plasma cells. The proportion of B-cells specialising is proportional to the number of CD4+ cells, with proportionality constant θ_{BC} . The proportion of specialising B-cells becoming memory cells is parameterised as θ_{BCm} . The decay of the (IgM and IgG) plasma cells is parameterised as θ_{dpB} and is independent of T-cells. Similarly, there is a constant move of immature memory cells into the mature memory cell state (θ_{mba}) and a decay of mature memory cells (θ_{dmB}) – neither of which depends upon the T-cells or presence of pathogen. The viral influence on B-cells is to cause the mature memory cells to transition to the IgG plasma cell state (parameterised by θ_{Bmp})³⁴.

The B-cell factor exerts an influence over the antibody factors, with IgM plasma cells leading to a move from absent to neutralising or non-neutralising proportional to θ_{pAb} (where the proportion that are neutralising is given by θ_{nAb}) in the IgM antibody factor, and IgG plasma cells leading to the equivalent dynamics in the IgG antibody factor. Regardless of the B-cell factor, antibodies decay to the absent state with rate parameterised by θ_{dAb} for IgM antibodies and the slower θ_{dAbG} for IgG antibodies³⁵. Finally, the antibodies exert an influence over the viral population, with extracellular virus moving to the absent state in proportion to the antibodies in the neutralising state. The overall structure of these dynamics is given in Eq. (1), where the $\Delta\tau$ used for the simulations is 1 hour.

$$\begin{aligned} \mathbf{p}(\tau + \Delta\tau)^i &= \mathbf{T}(\theta)^i \left(\mathbf{p}(\tau)^V \otimes \mathbf{p}(\tau)^{TC} \otimes \mathbf{p}(\tau)^{BC} \otimes \mathbf{p}(\tau)^{IgM} \otimes \mathbf{p}(\tau)^{IgG} \right) \\ i &\in \{V, TC, BC, IgM, IgG\} \\ \mathbf{p}(0)^V &= [1 - e^{-8}\theta_m \quad e^{-8}\theta_m \quad 0]^T \\ \mathbf{p}(0)^{TC} &= [1 \quad 0 \quad 0]^T \\ \mathbf{p}(0)^{BC} &= [1 - e^{-4}\theta_n \quad 0 \quad 0 \quad e^{-4}\theta_n \quad 0]^T \\ \mathbf{p}(0)^{IgM} &= [1 \quad 0 \quad 0]^T \\ \mathbf{p}(0)^{IgG} &= [1 \quad 0 \quad 0]^T \end{aligned} \quad (1)$$

Each \mathbf{p} is a vector of probabilities, with superscripts indexing the factor that these probabilities pertain to. \mathbf{T} is a transition matrix that updates these probabilities for a given time interval. The order of states shown in these probability vectors are as follows. The viral states are *absent*, *extracellular*, and *intracellular*. The T-cell states are *naïve*, *CD4+*, and *CD8+*. The B-cell states are *naïve*, *plasma (IgM)*, *immature memory*, *mature memory*, *plasma (IgG)*. Both sets of antibody states are *absent*, *non-neutralising*, *neutralising*. The probability vectors have dimensions that match the number of levels for each factor shown in Fig. 1. The transition matrix for each factor then has dimensions equal to the number of levels of the associated factor (rows) and the number of combinations of levels across all factors (columns). Note that a very small amount of virus can be present in the extracellular space at the start, and that a small proportion of B-cells start in the mature memory B-cell state. These initial quantities are kept close to zero by using small initial values (e^{-4} and e^{-8}). However, because these values are multiplied by scale parameter estimates, the initial quantities of virus and memory B-cell are estimated from data.

For a generative model to be useful, it should be able to generate data of the sort that can be measured. The model in Fig. 1 generates antibody levels, extracellular viral load, and the proportion of activated CD4+^{36,37} and CD8+ cells⁹. These are generated by summing the proportions of neutralising and non-neutralising antibodies, and taking the proportions of the other variables, and then scaling these by the maximum values we expect them to take. The scaling factors include a fixed and a random component. The fixed component determines the maximum value when the random component takes its prior expected value of 1. The random variables are θ -parameters as detailed in Table 1, meaning these could be estimated from measured data. These measurements are relative to some baseline which is likely to be zero for viral load but may be higher for the other variables.

Parameter	Symbol	Expectation	Variance
Scaling for initial immunity	$\ln \theta_n$	-2	1
Scaling for initial viral load	$\ln \theta_m$	0	1/256
Decay of IgM antibodies	$\ln \theta_{dAb}$	-6	1/1024
Decay of IgG antibodies	$\ln \theta_{dAbG}$	-8	1/1024
Production of antibodies by plasma cells	$\ln \theta_{pAb}$	0	1/1024
Proportion of antibodies that are neutralising	$\ln \theta_{nAb}$	-1	1/1024
Decay of plasma B-cells	$\ln \theta_{dpB}$	-2	1/1024
Decay of memory B-cells	$\ln \theta_{dmB}$	-32	1/1024
Activation of memory B-cells	$\ln \theta_{mba}$	-4	1/1024
Proportion of mature memory cells becoming plasma cells in presence of infection	$\ln \theta_{Bmp}$	-1/2	1/1024
Specialisation of naïve B-cells in the presence of CD4+ T-cells	$\ln \theta_{BC}$	-3	1/1024
Proportion of B-cells specialising as memory cells	$\ln \theta_{BCm}$	-1	1/1024
Decay of CD4+ cells	$\ln \theta_{dT4}$	-4	1/1024
Decay of CD8+ cells	$\ln \theta_{dT8}$	-5	1/1024
Production of CD4+ cells in presence of extracellular pathogen	$\ln \theta_{CD4}$	-3	1/1024
Production of CD8+ cells in presence of intracellular pathogen	$\ln \theta_{CD8}$	-3	1/1024
Neutralisation of intracellular pathogen by T-cell mediated apoptosis	$\ln \theta_{TCP}$	-2	7/32
Neutralisation of extracellular pathogen by antibody independent CD4+ mechanisms	$\ln \theta_{TAP}$	-8	1/1024
Viral entry into cells	$\ln \theta_{int}$	-2	1/16
Viral shedding into extracellular space (absorbing replication rate)	$\ln \theta_{ext}$	-1	1/1024
Decay of extracellular pathogen (independent of adaptive immunity)	$\ln \theta_{dpe}$	-2	1/1024
Decay of intracellular pathogen (independent of adaptive immunity)	$\ln \theta_{dpi}$	-2	1/1024
Scaling of antibody titre measures	$\ln \theta_{Abm}$	0	1/256
Scaling of viral load measures	$\ln \theta_{VLm}$	0	1/256
Scaling of IFN- γ T-cell assays	$\ln \theta_{IFN}$	0	1/256

Table 1. Priors for normal immune response (log parameters).

IFN- γ secretion is generally accepted as a measure of activation of CD4+ and CD8+ cells³⁸. Evidence from H7N9 influenza virus infections suggests that recovery from the viral infection is positively associated with IFN- γ secretion by CD4+ and CD8+ cells³⁸; emphasising the importance of these ‘activated’ cells as measures of the immune response. Put formally, the process generating data is as follows:

$$\begin{aligned}
 \mathbf{y}(\tau)^{IgM} &= 2^7 \left(\mathbf{p}(\tau)_2^{IgM} + \mathbf{p}(\tau)_3^{IgM} \right) \theta_{Abm} \\
 \mathbf{y}(\tau)^{IgG} &= 2^7 \left(\mathbf{p}(\tau)_2^{IgG} + \mathbf{p}(\tau)_3^{IgG} \right) \theta_{Abm} \\
 \mathbf{y}(\tau)^{VL} &= 10^4 \mathbf{p}(\tau)_2^V \theta_{VLm} \\
 \mathbf{y}(\tau)^{CD4} &= \mathbf{p}(\tau)_2^{TC} \theta_{IFN} \\
 \mathbf{y}(\tau)^{CD8} &= \mathbf{p}(\tau)_3^{TC} \theta_{IFN}
 \end{aligned} \tag{2}$$

Each of these data modalities (\mathbf{y}) are multiplied by a constant and a scaling parameter, which may be estimated during model inversion. By inversion, we mean the process of inferring the parameters that caused the data. We do not assume data are generated every hour (the rate of update of latent dynamics). Instead, we assume measurements are only available every 24 h. This is perhaps a little optimistic, but the model is formulated such that the measurement timings can be arbitrarily set, to account for when measurements are available. Table 1 provides the parameters used in the simulations that follow. It is important to emphasise that the priors chosen here are largely heuristic. These will vary with infection between different organisms, and between hosts. This means the simulations that follow should not be taken as making any comment on any specific infection. Ideally these would be estimated from empirical data. In the absence of these data, we have selected parameters such that the relative timing of each response is broadly consistent with that observed in biology, and that is consistent with a self-limiting infection that is largely cleared within about three weeks. We discuss this consistency in relation to the time course of measurements during SARS-CoV-2 infections below. While this allows us to illustrate general principles, these illustrations should be taken as qualitative as opposed to quantitative.

While the variances in Table 1 may seem very small, it is important to remember that we are working with log parameters—so small changes can result in large changes once exponentiated. We will see the large confidence intervals in subsequent figures reflecting this. Secondly, many of the parameters we deal with play the role of probabilities, which means tight priors are required so that these parameters do not exceed a sum total of one. We use these tighter priors for simulation purposes; we relax the variances later, when demonstrating model fitting. It is worth drawing attention to the variances for the initial immunity, the viral cell entry, and the T-cell

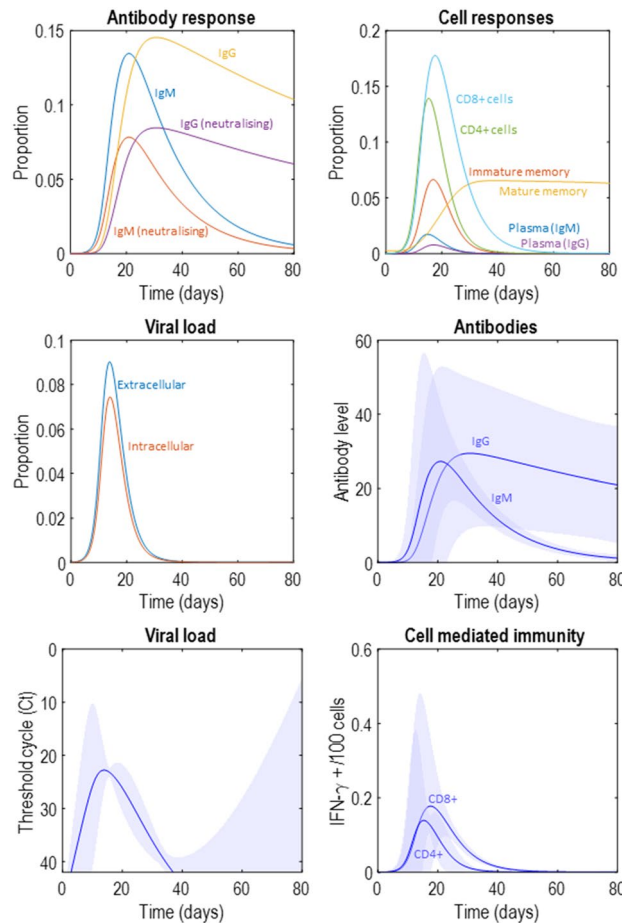


Figure 2. A synthetic immune response. The upper left plot shows the latent variables in terms of the different types of antibody. This distinguishes between the faster IgM response, and the more prolonged IgG response, and between the neutralising and non-neutralising variants of each. The upper right plot shows the cellular latent variables, with an initial increase in the activated T-cell proportions that leads to a later increase in plasma and memory cells, with a persistent memory response developing. The middle left plot shows the viral states, subdivided into intracellular and extracellular antibody components. The middle right plot shows the total antibody level, including both neutralising and non-neutralising. The lower left plot shows the viral load in terms of the ‘threshold cycle’ (Ct) which indicates the number of cycles of polymerase chain reaction (PCR) required before the viral nucleic acids are detectable. A greater number of cycles indicates a greater dilution (i.e., a smaller concentration). To generate these data, we took the negative $\log^{39,40}$ of $\gamma(\tau)^{V_L}$, scaled it by a factor of 4, and added a constant 50—under the assumption that more than 40 cycles implies a negligible viral concentration, based on fluorescence amplification results for a typical qPCR run⁴¹. These scale factors and constants could in principle be fit to empirical data. Here they are used only for the plotting and have no influence on model inversion. Finally, the lower right plot shows the proportion of CD4+ and CD8+ cells releasing IFN- γ . In this and subsequent figures, the intervals are 90% credible intervals for the predictive distribution. These are calculated as in¹, using a first order approximation to the variance based upon the chain rule. The uncertainty in the predicted data inherits from uncertainty about the parameters generating those data. Here, the uncertainty about the parameters under prior beliefs is used. However, these could be replaced by posterior beliefs when fitting to empirical data. One thing that is important to note is that day 1 is not the first day of symptoms. It represents the day of initial exposure to the virus.

mediated apoptosis parameters. The variances for these are chosen to be very wide as these are the key parameters we will be varying later. Specifically, these are chosen such that the modifications we will make to these parameters amount to a change of approximately one standard deviation each. This accounts for the variances of 1, 7/32, and 1/16, which are much larger than those for other priors. The reason for this will become more obvious in the section “[Immunological phenotyping](#)”. Briefly, they are chosen to avoid biasing our model fitting towards one phenotype or another.

Figure 2 illustrates an immune response generated from the prior expectations in Table 1. Here we see the rapid increase in viral load as the initial extracellular virus enters the cells enabling intracellular replication and shedding into the extracellular space. As the extracellular viral load increases, this initiates a helper (CD4+) T-cell response that increases the number of plasma B-cells and causes a rapid increase in antibody levels. Later peaks

are seen in the populations of cytotoxic (CD8+) T-cells and memory B-cells. As the viral load is suppressed, the immune cell populations decay, with the exception of the mature memory cells. This is because, in the presence of the virus, mature memory cells rapidly transition into plasma cells. This means they are the only cell type to decrease in number during the initial phase of the infection but increase later on from the immature memory cells acquired during the infection. In the antibody response plot, note the early IgM response, followed by a slower but more sustained IgG response. This reflects the characteristic role of these isotypes in fast and long-lasting immunity, respectively.

Figure 2 also illustrates some face validity of this approach in relation to the emerging data on the immune response in COVID-19⁴². The time-course and magnitude of the antibody response is broadly consistent with data suggesting peak antibody titres in the region of about 2^{643} (measured as chemiluminescence levels relative to some cut-off), consistent with the upper bound approaching 60 in the antibody plot in Fig. 2, and an increase in IgG relative to IgM over time, both plateauing at around day 20 following symptom onset. This is slightly earlier than in our simulation but may reflect the time it takes for the viral load to reach a level that causes symptoms. For ease of comparison with empirical reports we plot the viral load (y^{VL}), generated according to Eq. (2), after transforming into the cycle threshold for detection using a polymerase chain reaction (PCR). See the legend of Fig. 2 for details. The viral load appears to decline from a cycle threshold (Ct) of about 30–40 over the space of about 10 days⁴⁴, consistent with the lower left plot in Fig. 2. In the lower right plot, the percentage of CD4+ and CD8+ cells producing IFN- γ peaks at around 0.2 per hundred cells (with upper limit around 0.5), consistent with the kinds of ranges reported in some empirical studies⁴⁵. While these comparisons are heuristic and will almost certainly vary with laboratory equipment and local measurement procedures, they illustrate face validity in the sense that the data generated broadly resemble measurable data.

Routes to resistance

Figure 3 provides an overview of the mechanisms for resistance that we address here. This is not exhaustive but highlights three important ways in which the immune system might vary between individuals. We will take each of these in turn and demonstrate the influence they have on the dynamics of the immune response under our model. Figure 4 offers a comparison between the dynamics of latent states under each manipulation for ease of comparison between the time course of antibody, cellular, and viral responses. The same plots will be shown in subsequent figures, which we will unpack in detail.

The first route to resistance is a reduced capacity for viral entry into cells. In the context of SARS-CoV-2, this could be due to polymorphisms in the cell surface ACE-receptors targeted by the virus⁴⁶—or to age-dependent differences in expression⁴⁷. In terms of the model outlined above, this is an attenuation of the rate at which extracellular virus becomes intracellular. Specifically, we changed the expectation of $\ln\theta_{im}$ from -2 to -2.2 . Figure 5 shows a simulation of this scenario. Given that viral reproduction occurs in the intracellular compartment, slower entry into this compartment leads to a failure to increase the viral population. This leads to a viral load that is lower at its peak than in Fig. 2—possibly suggesting people with this phenotype are less infectious (note the differences in y-axis scales between Figs. 2 and 5). However, the most obvious difference between Figs. 2 and 5 is the magnitude of the immune response required to clear the virus. In Fig. 5, we see a much smaller response, and might anticipate a failure to detect antibodies post-infection in such individuals. In addition, a smaller immune response might mean fewer symptoms during the course of the infection.

The second route is pre-existing cross-reactivity of antibodies, which we simulate through assuming a larger population of memory cells capable of rapid specialisation into plasma cells that target the virus. By cross-reactivity, we mean that antibodies developed to some previous infection may also be effective against the current infection^{10,48}. Specifically, we formalise this by increasing $\ln\theta_n$ from -2 to 2 . Figure 6 illustrates the rapid drop in mature memory cell populations as they become IgG plasma cells, and the resulting increase in IgG antibodies. Note that the plasma IgG cells never reach the same level as the original memory cell population. This is because their rate of decay is considerably faster than that of memory cells. The key differences between this simulation and that reported in Fig. 2 are the earlier and more dramatic IgG response, which now precedes the attenuated IgM response, and the decrease in peak viral load. Longitudinal studies of these antibodies in COVID-19 show this pattern in patients with low severity disease, compared to an earlier IgM response (as in Fig. 2) in more severe disease⁴⁹. While this formulation of cross-reactive immunity seems to accommodate findings in SARS-CoV-2 infection, the opposite effect—an ineffective memory-driven response—may occur in other infections. The ‘original antigenic sin’ hypothesis⁵⁰ explains why this might be. If the pathogen has antigens that are sufficiently different to the previous infection, the secondary IgG response may fail to neutralise the second infection. However, it may still attenuate the IgM response, precluding the development of a primary response—as if to a novel pathogen. This could be modelled by attenuating the rate of decline in the pathogen in the presence of IgG, but not IgM.

Third, the response could be mediated by antibody independent T-cell mechanisms^{9,51} that suppress the virus before the viral load is sufficient to promote plasma cell proliferation. A similar hypothesis could be advanced for a purely innate immune response, by increasing the rate of decay of the virus independently of the states of the other three factors. Cell-mediated immunity is well recognised to vary with age^{52,53} and could account for the relative resistance of younger members of the population. Suggestive evidence in favour of the relevance of CD8+ dependent mechanisms comes from genomic studies that have identified loci associated with lung CD8+ cells as differing between those with severe COVID-19 and the control population (people without a known diagnosis of COVID-19)⁵⁴. Some T-cells in the respiratory system are thought to have a role in immunological memory, which may be independent of B-cell memory systems^{55,56}. This is one possible interpretation for a more vigorous T-cell response in some people, whose T-cells have encountered similar antigens in the past. There are also emerging data suggesting SARS-CoV-2 itself induces a memory T-cell response⁴⁵, in addition to there being

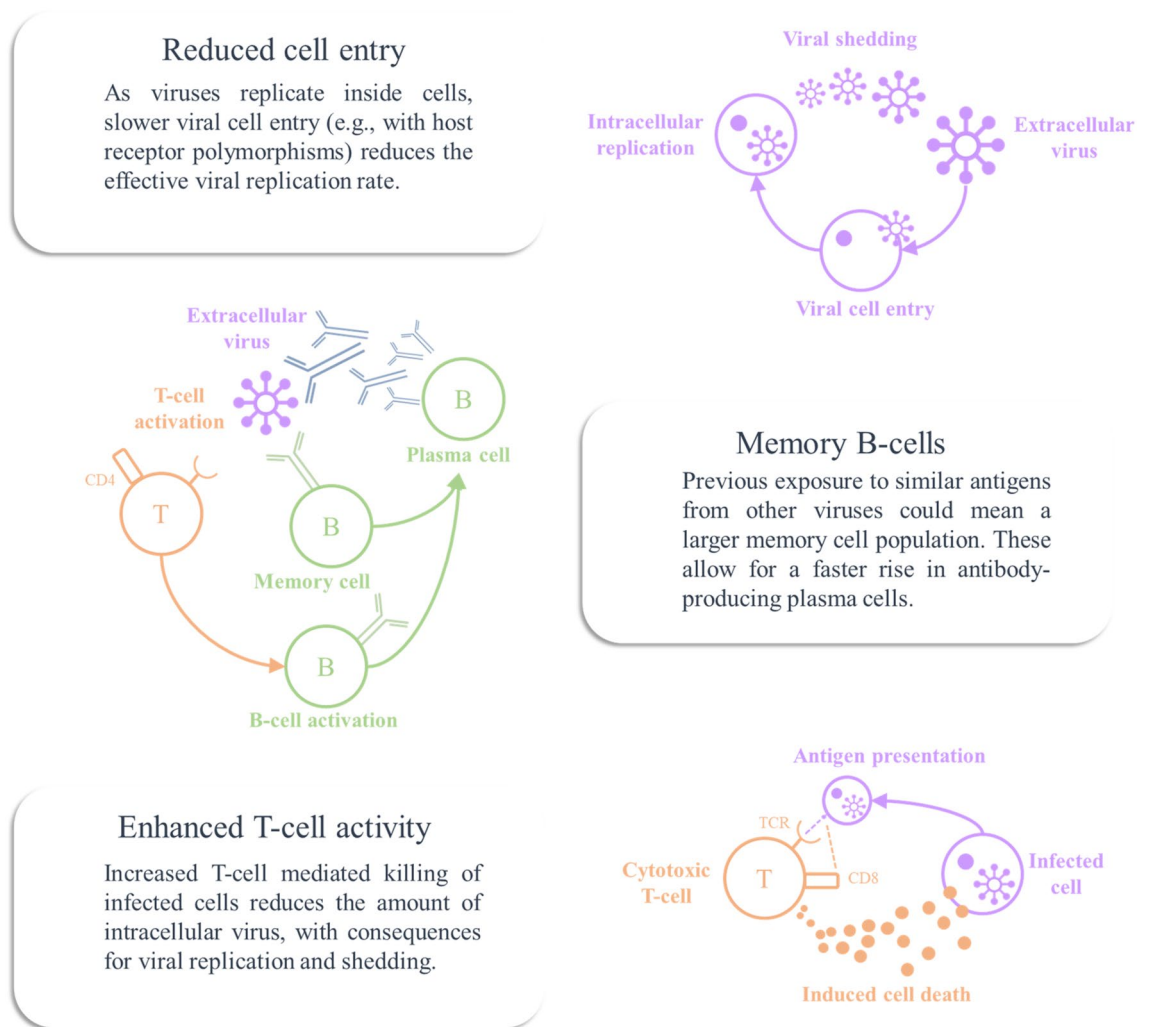


Figure 3. Mechanisms of resistance. In what follows, we will illustrate various modifications to the model simulated in Fig. 2, by changing model priors. The graphics in this figure provide an overview of the three manipulations we will pursue in terms of the associated biological mechanisms. These mechanisms are reduced viral entry into host cells, pre-existing memory B-cells, and enhanced T-cell dependent killing of infected cells.

T-cells reactive to viral antigens in people who have yet to be exposed⁹. Although we have formalised this in terms of CD8+ responses (by increasing $\ln\theta_{TCP}$ from -2 to $-1/8$), there is also evidence in favour of a role for CD4+ responses in rapid suppression of influenza viruses in some individuals—something that has been proposed as a mechanism for the difference in disease severity between successive infections in individuals²⁸. Figure 7 shows the influence of increasing the effectiveness of T-cell dependent killing on the immune response, which appears to reduce viral load and dampen and prolong antibody and cell mediated responses.

Some of these phenotypes are inconsistent with most SARS-CoV-2 trajectories. For instance, in the plot of the viral load with enhanced T-cell activity, the heavy tail deviates from the suppression of the virus in most people 10–14 days after developing symptoms—noting that there is considerable variability in these estimates^{57,58}. If this phenotype exists at all, it is likely that such people are in the minority. Interestingly, there is emerging evidence in favour of an association between the T-cell response and positive PCR testing weeks following recovery from infection⁵⁹, as this model would predict.

Clearly there are many other phenotypes we could have considered⁶⁰. These could include mixtures of the above, or more extreme values of the parameters. For example, if we reduce the viral cell entry sufficiently, the viral load never rises at all, but immediately decays to zero—with a negligible immune response. In contrast, if we sufficiently attenuate the immune response (e.g., by reducing activation of T and B cells), the viral load increases to its carrying capacity and stays there. There are a range of other intermediate scenarios, like a small but persistent viral load that is never fully suppressed (perhaps more consistent with some hepatitis viruses⁶¹), or even second waves (without reinfection) if the antibodies decay sufficiently quickly. These parameter settings might be more appropriate in modelling viruses with latent phases that reactivate when immunity decays. For example, varicella zoster virus causes chickenpox before being suppressed, but can re-emerge later as shingles⁶².

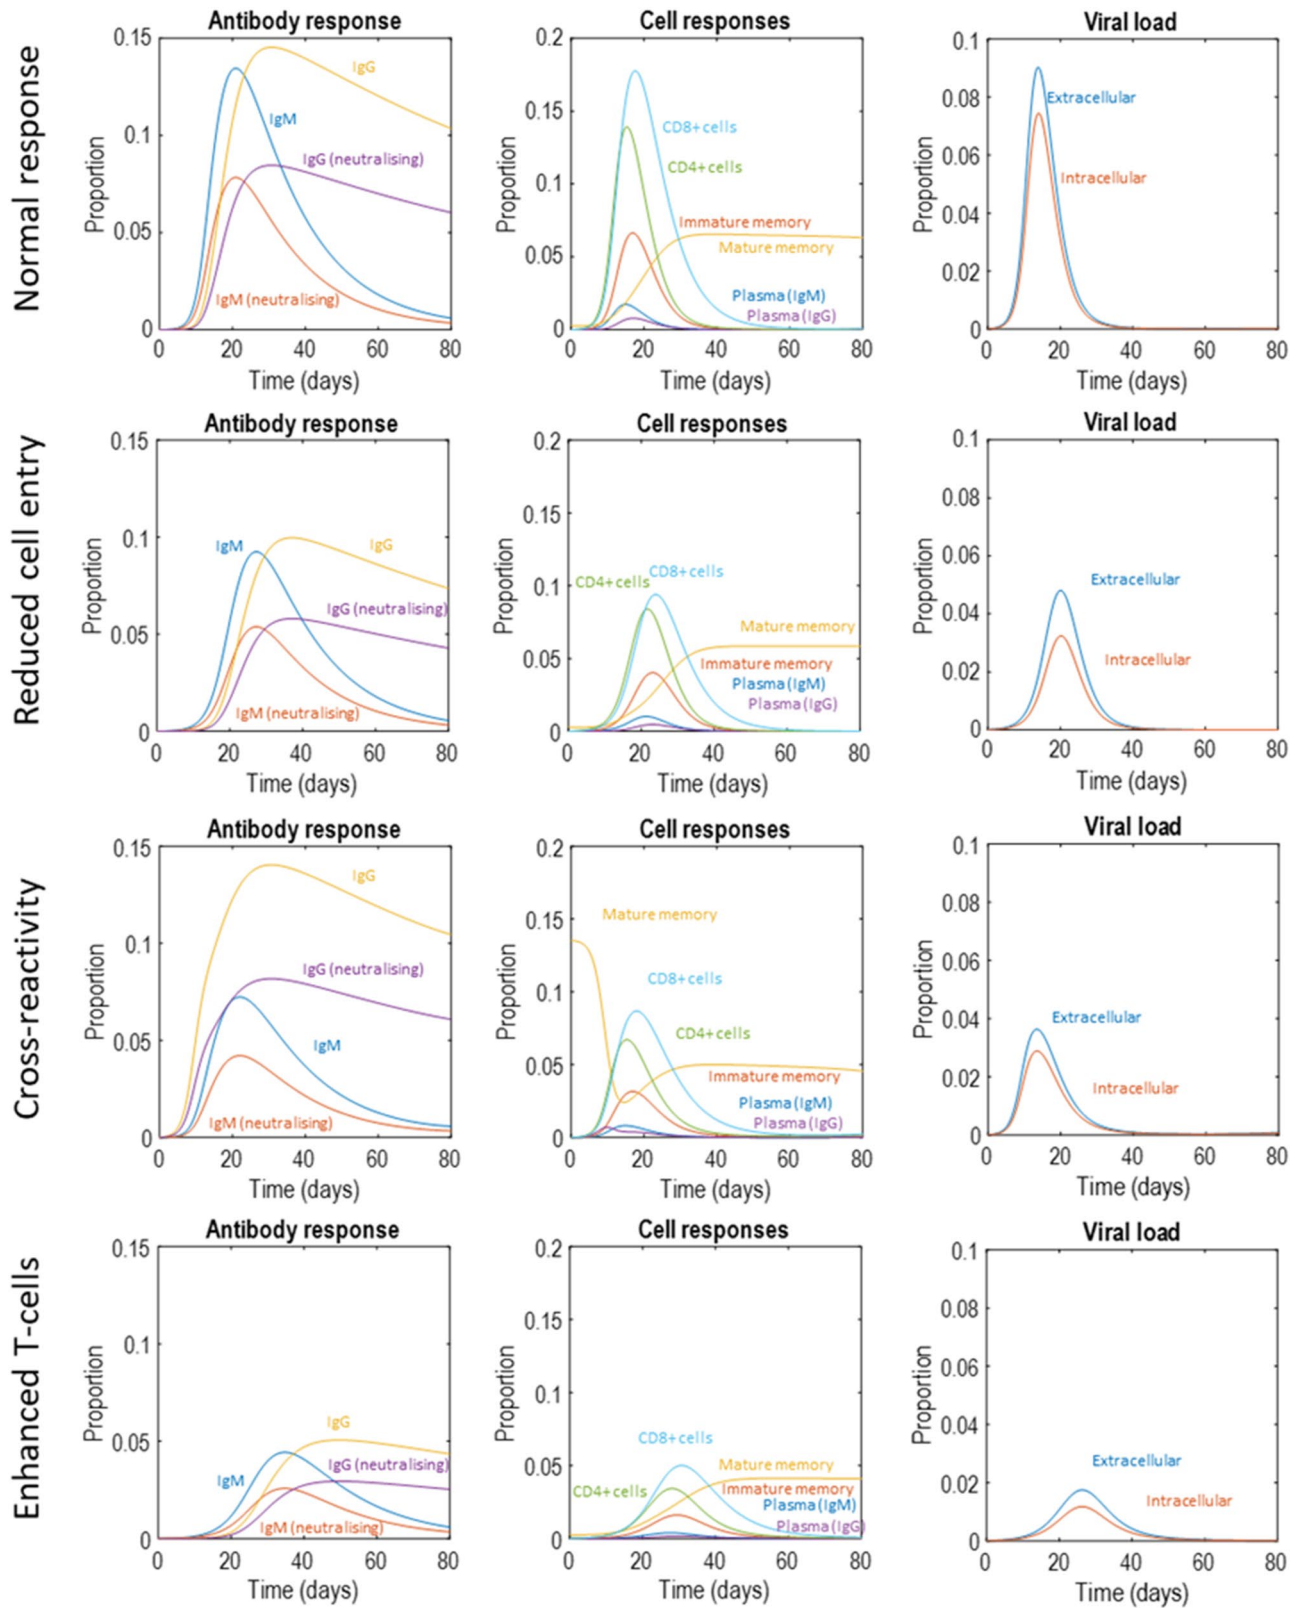


Figure 4. Latent variables. This figure takes the latent variable plots from Fig. 2 and supplements them with those that will appear in Figs. 5, 6 and 7, standardising the y-axes for ease of comparison. The format is as described in Fig. 2. We will unpack the details of these plots as they are introduced in subsequent figures. However, the highlights are (1) the attenuation of peak viral load in the second to fourth row of plots, relative to the first row, (2) the smaller antibody response required to suppress the infection in the second and fourth rows of plots, (3) the earlier peak in IgG plasma cells and antibodies in the third row of plots. The key message from this figure is that there are several different dimensions along which immune responses may vary, with consequences for viral suppression.

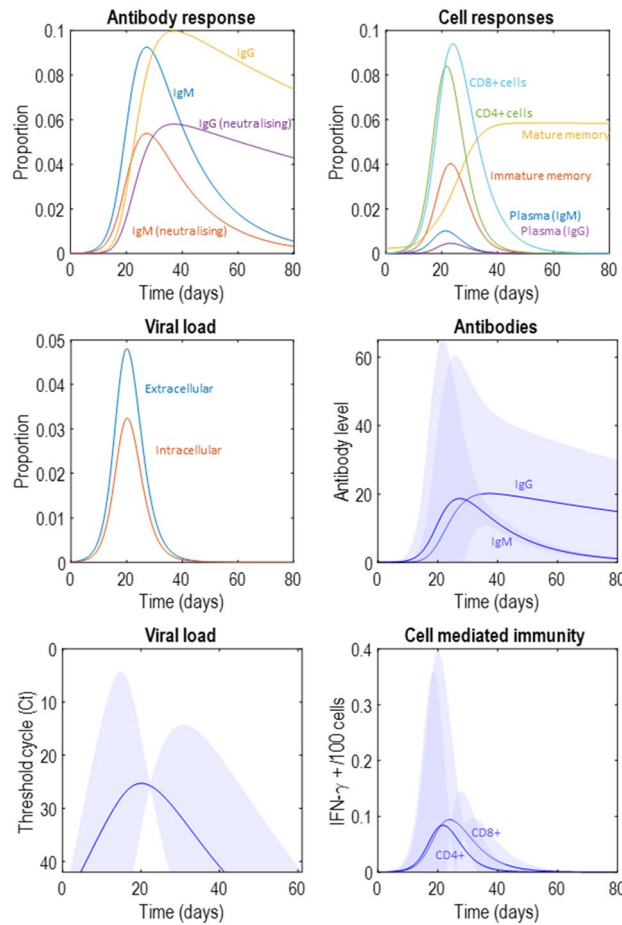


Figure 5. Attenuating viral entry into cells. This figure uses the same layout as Fig. 2. Here, the proportion of extracellular virus entering the intracellular compartment per hour has been decreased, remembering that the production of extracellular virus depends upon there being intracellular virus that can use the cellular machinery to manufacture and secrete new virus particles. Decreasing the rate of cell entry therefore slows the growth of the viral population, leading to a smaller peak with a smaller immune response required to suppress the infection. Note the decrease in antibody titre, which might mean that people exhibiting this immune phenotype may not have detectable antibodies in serological testing post-infection.

Immunological phenotyping

In this section, we attempt to do two things. The first is to find out whether the three phenotypes identified above (lower proportion of extracellular virus entering cells; higher proportion of mature memory cells on infection; and higher proportion of intracellular virus destroyed through T-cell mediation) may be disambiguated from one another based upon the data that we assume they generate. This rests on the idea that an individual has one, rather than a mixture, of these phenotypes. However, it would be simple to extend this analysis by specifying additional phenotypes that comprise mixtures of those outlined in the section “Routes to resistance”. The second aim in this section is to ask whether we could identify interactions between demographic factors—such as whether or not someone has received a Bacillus Calmette-Guérin (BCG) vaccine—and parameters of the models above—for example, T-cell dependent killing of pathogen⁶³. We will use the BCG hypothesis, that T-cell dependent killing is enhanced by the BCG vaccine, as an example to aid intuition. This is not intended to endorse or refute this hypothesis, but to show how DCM may be used to formulate and test hypotheses of this sort.

For the first of these aims, we took the outputs of the above simulations assuming one measurement per day for the first three weeks, added univariate random fluctuations with a variance of 1/8 to the square roots of the data. We then fit each synthetic dataset—comprising daily measurements of the viral load, antibody levels, and proportion of T-cells releasing of IFN- γ —to the model with the prior expectations in Table 1, but with variances multiplied by 16. Model fitting here means finding approximate posterior probabilities Q that maximise a lower bound on model evidence (the marginal likelihood of the data given the model):

$$Q(\theta|m) = \arg \max_{Q(\theta|m)} \{ \mathbb{E}_{Q(\theta|m)} [\ln P(\theta, \mathbf{y}|m) - \ln Q(\theta|m)] \} \approx P(\theta|\mathbf{y}, m) \tag{3}$$

$$\mathbb{E}_{Q(\theta|m)} [\ln P(\theta, \mathbf{y}|m) - \ln Q(\theta|m)] \leq \ln P(\mathbf{y}|m)$$

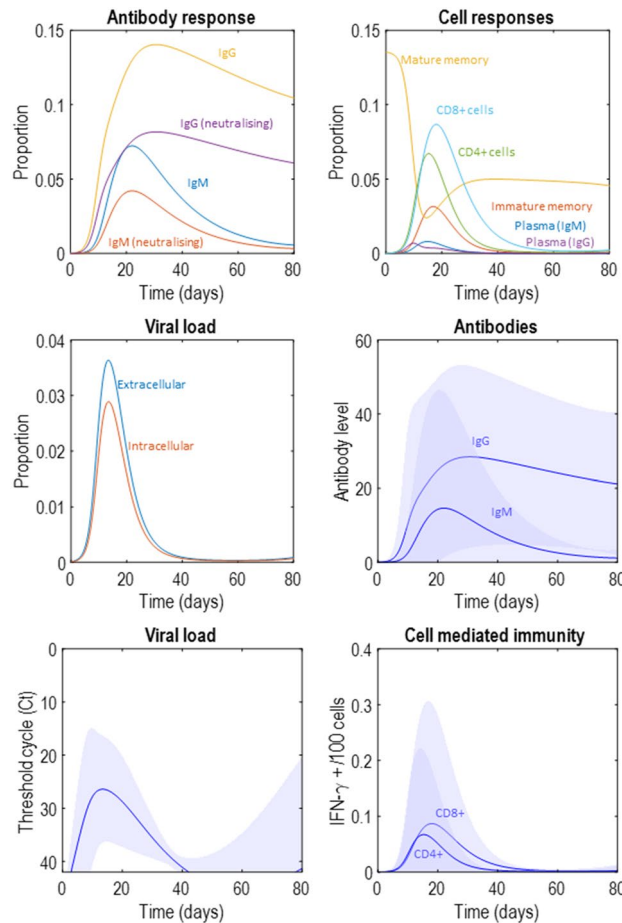


Figure 6. Pre-existing immunity. This figure uses the same layout as Fig. 2. Here, the proportion of mature memory cells at the start of the infection has been increased, to simulate the scenario that immunological memory may have been instantiated by exposure to related viruses in the past. Note the rapid increase in the IgG response, in contrast to the early IgM response we saw in Fig. 2. Like in Fig. 5, the virus is suppressed and never reaches as high a viral load as in Fig. 2. However, here we see a sustained increase in serum immunoglobulin in the latent (upper left) and measurable (middle right) antibody response plots. This pattern of response might be associated with a very mild illness, not prompting testing for the virus (or possibly even an undetectable viral load on being tested) but would be associated with positive serological tests later on.

In Eq. (3), θ are the parameters, \mathbf{y} the data, and m is the model. Maximisation of the model evidence ensures an optimal trade-off between the accuracy with which data can be explained and the simplicity of that explanation—penalising excessively complex explanations. This maximisation was achieved using a standard software routine (`spm_nlsi_GN.m`—see “Software note”). This routine uses the Gauss–Newton method to maximise the evidence lower bound. Practically, this means finding the gradients of the quantity in Eq. (3) with respect to the expectations of the parameters under an approximate posterior density. Using a regularised gradient ascent, the approximate posterior is iteratively improved until convergence. This is known as variational Laplace. We then used Bayesian model reduction (BMR)—a method for efficient post-hoc model comparison^{64,65}—to evaluate the relative evidence for alternative models with the prior expectations set out in Table 1, with expectations modified as described in “Routes to resistance”. The reason for the more uncertain priors used during fitting of the full model is that BMR treats each alternative model as a reduced version of the ‘full’ model. By setting the variances in the full model to be much larger than those in any of the reduced models, we ensured that the expectations of the reduced priors would be plausible under the full prior distribution. The principle that underwrites BMR is that, when only the priors differ between models, the relative evidence of two models is given by the following equation—which only requires inversion of one of the two:

$$\ln P(\mathbf{y}|\tilde{m}) - \ln P(\mathbf{y}|m) = \mathbb{E}_{P(\theta|m,\mathbf{y})} \left[\ln \frac{P(\theta|\tilde{m})}{P(\theta|m)} \right] \quad (4)$$

The graphic on the left of Fig. 8 shows a confusion matrix, whose rows correspond to the model used to fit the data, and whose columns correspond to the model that generated those data. Each element of the matrix is shaded by the posterior probability of the model given the data, with white indicating a probability of one and black zero. Posterior probabilities are obtained through using a softmax (normalised exponential) operator on

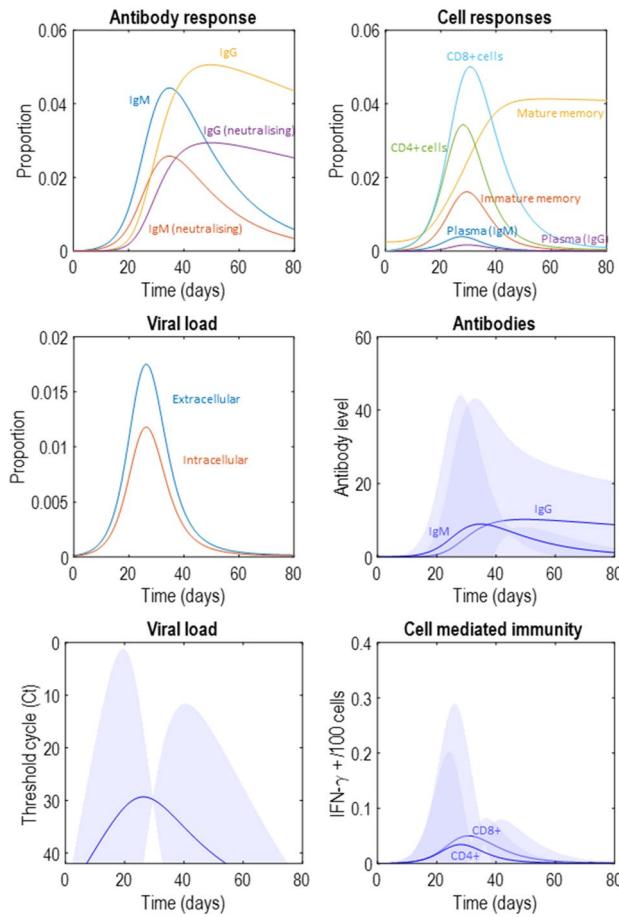


Figure 7. Cell mediated immunity. This figure uses the same layout as Fig. 2. The scenario simulated here is one in which the proportion of intracellular virus destroyed through T-cell mediated mechanisms is increased. This could be interpreted as a memory T-cell mediated response. Here we see a smaller viral load and a flattened antibody response.

the log model evidence given by BMR for each column. The lighter shading along the diagonal elements compared to the off-diagonal elements in the same column implies it is possible to recover the mechanism generating observed timeseries—albeit with varying degrees of confidence. The implication of this is that, in principle, it is possible to use empirical data to assign individuals to the immunological phenotype that best explains their measured serology.

The plots shown on the right of Fig. 8 illustrate an alternative application of this approach to answer a question at the level of between subject effects. Here, we generated 32 immune responses, as if from 32 different subjects. Each subject was assigned to the BCG group or the unvaccinated group with a probability of 1/2. The priors were all as in Table 1, except for the θ_{TCP} parameter whose expectation was sampled from a normal distribution for each individual. The mean of this distribution was set at -2 , with 2 added for those with a BCG. In this synthetic example, this implies the BCG vaccine confers some T-cell mediated immunity. We then fit each of the 32 synthetic datasets using the priors of Table 1, and used a parametric empirical Bayes scheme (`spm_dcm_peb.m`) to estimate the parameters of a second level linear model^{64,67}, which predicted the θ_{TCP} parameter based upon the presence or absence of the BCG. This second level model has the form:

$$\ln \theta_{TCP}^i = \beta_0 + X^i \beta_1 + \omega^i \tag{5}$$

Here, the superscript i indicates the subject, the β_0 is a constant value for those with no BCG, and the β_1 is the effect of having had the BCG. X^i is one if a subject has had the BCG, and zero otherwise. The ω term indicates normally distributed, zero-mean, fluctuations. The plots on the right of Fig. 8 show that a model in which both β parameters are allowed to be non-zero provides the best explanation for the synthetic data. Furthermore, the maximum a posteriori estimates are consistent with the -2 for the absence of a BCG and the positive effect associated with this vaccine used to generate the synthetic data. This illustrates that, if it were true that the BCG vaccine enhances T-cell-mediated immunity, we could recover this information using this sort of analysis. Again, this is not to suggest that this hypothesis is correct (or incorrect). Instead, our aim is to show that it could be demonstrated if it were true.

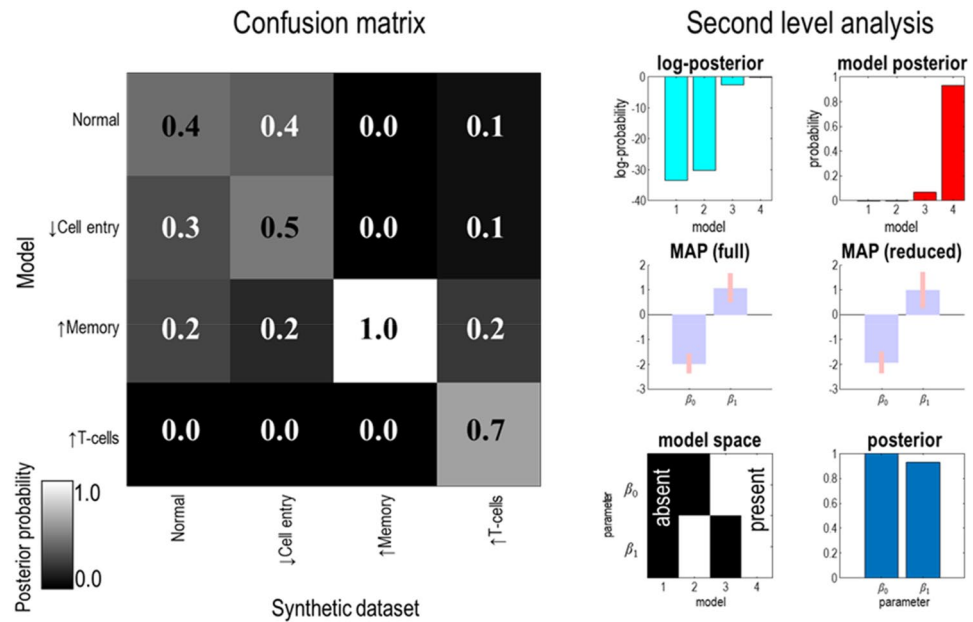


Figure 8. Immunological phenotyping. This figure illustrates two ways in which this modelling approach may be used. The left panel shows a confusion matrix that plots the posterior probability of each model given each dataset. White indicates a probability of one, while black indicates a probability of zero, with grey being intermediate values. The diagonal elements of this matrix represent the posterior probability of a model given the data that it was used to generate, and these probabilities are reassuringly higher than those for other models. The models are labelled according to the manipulation applied to the parameters and correspond to the rows of Fig. 4. The idea here is that alternative hypotheses may be posed to data from an individual's immune response, and that we hope to be able to recover which model best represents that individual's immune response. As the priors over each model are uniform, the posterior probabilities (rounded to one decimal place) overlaid on each square are proportional to the marginal likelihood for each model. They can therefore be used to compute Bayes factors (i.e., marginal likelihood ratios); for comparing one model with another. There are various heuristic criteria for what constitutes a definitive model comparison. One option is to use a Bayes factor of 20 in analogy with the 0.05 threshold common in frequentist statistics. However, this is not an all-or-nothing threshold, and 3.2 is often quoted as indicating substantial, if not definitive, evidence⁶⁶. Here we see that the model that includes memory cells can be definitively disambiguated from all alternatives, that all models can be disambiguated from the enhanced T-cell model, and that other pairwise model comparisons have varying degrees of confidence. The second use of this approach is shown on the right, where a synthetic group level analysis has been used to test hypotheses about the role of demographic factors (e.g., whether or not someone has received a BCG vaccine) on parameters of the model (e.g., the cytotoxic T-cell response). The plots labelled *log-posterior* and *model posterior* show the probability associated with group level models that include various combinations of constant (β_0) and linear (β_1) effects on the θ_{TCP} parameter. The constant effect is the expected value for all those without a BCG vaccine, while the linear effect is the amount added to the log expected parameter when someone has had this vaccine. The presence or absence of these parameters (where absence means set to zero) for each model is shown in the *model space* plot with white indicating presence and black absence. The *posterior* plot shows the probability of each parameter being present. The *MAP* plot shows maximum a posteriori estimates for the model including all parameters (with 90% credible intervals) while the *MAP (reduced)* plot shows these when averaged under the relative probability of each reduced model.

Discussion

How does this sort of modelling relate to epidemiology? Drawing from theoretical neurobiology, we could think of immunological and epidemiological models as simply different levels of parcellation—just as we have the option of parcellating the brain on the level of regions or on the level of networks, depending upon the question at hand. Recent work⁶⁸ in this direction makes use of renormalisation groups⁶⁹ to translate between these scales. In brief, this works by identifying functional units at small spatial scales (e.g., individuals at different stages of viral exposure) and extracting low dimensional summaries of their dynamics based upon the system's Jacobian. These are typically along the direction of the eigenvectors with the smallest negative eigenvalues—effectively omitting dynamics that dissipate very quickly⁷⁰. This procedure may be applied recursively to move to high level summaries in a population (e.g., the proportion of individuals at each stage). A key feature of these analyses is that increases in spatial scales are accompanied by slower dynamics. Consistent with this, the course of infection for an individual—lasting for a period of a few days to weeks—is much shorter than that of an epidemic (a few months) or a pandemic (potentially years). This method affords the possibility of using inferences drawn at the level of individual immune responses to nuance models used at an epidemiological level.

It is important to emphasise that none of the hypotheses considered above are ‘new’. All have been advanced in the context of the current pandemic or in other infections. What is offered here is a formalisation of these hypotheses, and a generative model that could be used to evaluate them. In addition, generative models may be used to simulate interventions (e.g., pharmacological treatments) to try to understand what might happen before moving to more expensive or risky empirical tests. The generative modelling approach has been employed in the past to assess models of the immune response to dengue virus⁷¹. The authors used temporally dense measurements of viral load and antibody titre to fit a ‘target-cell’ model—highlighting the utility of dynamic modelling of immunity. The approach pursued in this paper is closely related to the target-cell approach to modelling viral infections but differs in a few important ways⁷². Target-cell models—in their simplest manifestation—describe the dynamics of an infection in much the same way Susceptible-Exposed-Infectious-Recovered (SEIR) models describe an epidemic. In place of the number of individuals at various stages of being infected by a pathogen, target-cell models deal with the number of cells infected. Our approach here differs from this in that it factorises the immune response into several interacting modules. This affords the opportunity for richer generative models which, using variational Bayes and Bayesian model reduction, may be inverted relatively quickly.

It is interesting that the mechanisms of resistance illustrated here tend to flatten, as opposed to fully suppress, the viral load curve. This could mean one of (or some combination of) two things. A lower load may make someone less able to transmit the infection. Alternatively, a prolonged period of having a non-negligible viral load may increase the period for which they are infectious. Each of these has important consequences in SEIR (or LIST) epidemiological models. The former favours resistance both to the infection itself, and to passing it on. The latter might be more consistent with the notion of a super-spreader, who has an insufficient viral load to cause symptoms that would prompt testing or isolation but is infectious for longer.

Clearly a major limitation of the work presented here is that it is based upon purely synthetic data. As such, this should be viewed as a proof-of-principle whose priors will require refinement as data from longitudinal studies⁴⁵ and vaccine trials become available⁷³. Related to this are the temporally dense (daily) measurements assumed in the model fitting. In the presence of more precise prior beliefs derived from such studies, it should be possible to achieve posterior estimates with greater certainty and less data per individual. Another limitation is the relative coarseness of the immune model. There are many nuances that could be incorporated into a model of this sort. Relevant to the current SARS-CoV-2 pandemic is the direct influence of the virus on T-cell populations. Some stages of the SARS-CoV-2 infection are associated with a decrease in the functional CD8+ cell population¹¹ and lymphopenia⁷⁴. There is precedent for thinking about targeting of lymphocytes by viral pathogens, with the most obvious example being the reduction in CD4 counts mediated by human immunodeficiency virus (HIV)⁷⁵. However, the pattern in COVID-19 seems to be more a skew in the sorts of cells naïve T-cells differentiate into⁷⁶—a viral effect on T-cell transition probabilities that could be explicitly parameterised using this approach.

There is a trade-off in our decision to use normally distributed priors for the log parameters. The limitation of this choice is that a small but finite probability is afforded to elements of the probability transition matrices being greater than one. This is mitigated here by using small prior variances—expressing the prior belief that probabilities do not exceed one. An alternative parameterisation could have used Dirichlet distributions as priors for the columns of the transition matrices, guaranteeing summation to one. The downside of a Dirichlet prior is that it requires specification of a concentration parameter for every element in the transition matrix, as opposed to a pair of (mean and variance) parameters for each allowable transition. Either of these distributions would have been a reasonable choice for the priors. Practically, the flexibility afforded by using normal distributions, and the ensuing variational Laplace scheme—that has been validated in a wide range of applications^{5,1,77–79}—speaks to the use of lognormal priors.

One possible extension to this work would be to consider separate anatomical compartments. Here, the same mean-field model could be replicated for different parts of the body. This would be similar to the connection of multiple neural mass models in neuroscience—to model effective connectivity between different brain regions⁸⁰—or to connectivity between multiple LIST models as was used to model SARS-CoV-2 propagation between parts of the United States³. This would allow for modelling of the transfer of virus between different anatomical sites (e.g., between the nasopharynx and the lungs) in addition to transfer of immune cells (e.g., between lymphoid organs and the lungs)⁸¹. Each site, with its distinct tissue makeup, will require modifications to the model to account for the dominant cell and antibody type. For example, the IgA isotype is known to play an important role in mucosa-associated lymphoid tissue (MALT), in the lung and nasopharynx^{82–85}. This diversity could be incorporated into a mean-field model as follows: one could add a factor representing IgA antibodies—or an additional cell type—and specify tissue-specific priors. For instance, in MALT tissue, the prior belief about the rate at which IgA is synthesised in the presence of plasma cells would be higher than in lymph nodes. Including this anatomical dimension may also be of use in modelling the effect of COVID-19 on specific body systems, including the increasingly recognised neurological manifestations, enabling measurements in the cerebrospinal fluid to inform the same parameters as serological measures. However, it is worth noting that there is little evidence so far for cerebrospinal viral load playing a direct role in neurological disease⁸⁶.

Finally, having discussed the translation of analysis tools from neuroscience to immunology, it is worth briefly considering other points of contact that could be exploited in future work. One of these is the importance of immunological processes in neurological disease. Another is the relationship between the brain and immune system in health. Regarding the former, there is an opportunity to combine neurobiological and immunological models. Recently, DCM for neuroimaging has been extended to accommodate fusion of several different imaging modalities⁸⁷. This raises the possibility of fusing immunological and neurobiological DCMs, to model conditions—like multiple sclerosis^{88,89} or autoimmune encephalitis⁹⁰—in which the immune and nervous systems interact. Here the aim would be to generate both neurophysiological and immunological measurements. Given

serial electroencephalograms and serological tests, it might be possible to measure the influence of immune system compartments on neural connectivity over time.

The relationship between the nervous and immune systems in health is another important topic. Here there is another opportunity to translate between theoretical neurobiology and immunology. A prominent approach in the former is active inference, which frames the internal dynamics of the nervous system as solving an inference problem based upon sensory data—where these inferences direct behaviour through effector systems⁹¹. From this perspective, it suffices to define the problem being solved. The optimal solution emerges from identification of this problem. This has been successful in accounting for a range of neurophysiological and behavioural observations^{92–94}. Like the nervous system, the immune system exhibits complex internal dynamics⁸¹, and interfaces with foreign agents via (molecular) receptors and effectors. It follows that the immune system may be amenable to similar approaches, which could help uncover the principles from which the interactions simulated above could emerge. Given extensive interfaces between neural and immune tissue—both directly^{95–99} and via the neuroendocrine system¹⁰⁰—it may be possible to take this further, through identifying the problem that these systems jointly solve. We hope to explore this avenue of theoretical neuroimmunology in future work.

Conclusion

Ultimately, the simulations above are not designed to draw any specific conclusions about the nature of ‘immunological dark matter’, but to show how dynamic causal modelling could be applied to understand which mechanisms of resistance (i.e., attenuation of viral load) might be in play. However, one conclusion that can be drawn from this is the importance of not over-interpreting serological results. If alternative immunological phenotypes like those simulated above exist in populations affected by COVID-19, serological measures alone are likely to underestimate the effective immunity (or resistance) in the population. We have illustrated the use of a dynamic model of immune responses to supplement empirical and conceptual analyses of responses to viral pathogens or vaccines. In addition, we found that—using synthetic data—it is possible to recover the models that generate immune responses and to test hypotheses about the relationship between demographic factors and parameters of these models.

Software note. The simulations presented here may be reproduced and customised using the **DEM_Immune.m** demo which will be available as part of the next public release of SPM12. This academic freeware may be downloaded from <https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>.

Received: 29 September 2020; Accepted: 20 May 2021

Published online: 31 May 2021

References

1. Friston, K. *et al.* Dynamic causal modelling of COVID-19. *Wellcome Open Res.* <https://doi.org/10.12688/wellcomeopenres.15881.1> (2020).
2. Friston, K. *et al.* Testing and tracking in the UK: A dynamic causal modelling study. *Wellcome Open Res.* <https://doi.org/10.12688/wellcomeopenres.16004.1> (2020).
3. Friston, K. *et al.* Second waves, social distancing, and the spread of COVID-19 across America. *Wellcome Open Research* <https://doi.org/10.12688/wellcomeopenres.15986.1> (2020).
4. Friston, K. J., Parr, T., Zeidman P *et al.* Effective immunity and second waves: a dynamic causal modelling study[version 2; peer review: 2 approved]. *Wellcome Open Res* 2020, 5:204. <https://doi.org/10.12688/wellcomeopenres.16253.2>
5. Friston, K. J., Harrison, L. & Penny, W. Dynamic causal modelling. *Neuroimage* **19**, 1273–1302 (2003).
6. Spinney, L. COVID-19 expert Karl Friston: ‘Germany may have more immunological “dark matter”’. *The Observer* (2020).
7. Doshi, P. COVID-19: Do many people have pre-existing immunity?. *BMJ* **370**, m3563. <https://doi.org/10.1136/bmj.m3563> (2020).
8. Bunyavanich, S., Do, A. & Vicencio, A. Nasal gene expression of angiotensin-converting enzyme 2 in children and adults. *JAMA* **323**, 2427–2429. <https://doi.org/10.1001/jama.2020.8707> (2020).
9. Grifoni, A. *et al.* Targets of T cell responses to SARS-CoV-2 coronavirus in humans with COVID-19 disease and unexposed individuals. *Cell* **181**, 1489–1501.e1415. <https://doi.org/10.1016/j.cell.2020.05.015> (2020).
10. Ng, K. *et al.* Pre-existing and de novo humoral immunity to SARS-CoV-2 in humans. *bioRxiv* <https://doi.org/10.1101/2020.05.14.095414> (2020).
11. Zheng, M. *et al.* Functional exhaustion of antiviral lymphocytes in COVID-19 patients. *Cell. Mol. Immunol.* **17**, 533–535. <https://doi.org/10.1038/s41423-020-0402-2> (2020).
12. Akondy, R. S. *et al.* The yellow fever virus vaccine induces a broad and polyfunctional human memory CD8⁺ T cell response. *J. Immunol.* **183**, 7919. <https://doi.org/10.4049/jimmunol.0803903> (2009).
13. Gilbert, S. C. T-cell-inducing vaccines—What’s the future. *Immunology* **135**, 19–26. <https://doi.org/10.1111/j.1365-2567.2011.03517.x> (2012).
14. Pollán, M. *et al.* Prevalence of SARS-CoV-2 in Spain (ENE-COVID): A nationwide, population-based seroepidemiological study. *Lancet* [https://doi.org/10.1016/S0140-6736\(20\)31483-5](https://doi.org/10.1016/S0140-6736(20)31483-5) (2020).
15. Beal, M. J. Variational algorithms for approximate Bayesian inference (Doctoral thesis, University College London) (2003).
16. Winn, J. & Bishop, C. M. Variational message passing. *J. Mach. Learn. Res.* **6**, 661–694 (2005).
17. Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J. & Penny, W. Variational free energy and the Laplace approximation. *Neuroimage* **34**, 220–234. <https://doi.org/10.1016/j.neuroimage.2006.08.035> (2007).
18. Wu, J. T., Leung, K. & Leung, G. M. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: A modelling study. *Lancet* **395**, 689–697. [https://doi.org/10.1016/S0140-6736\(20\)30260-9](https://doi.org/10.1016/S0140-6736(20)30260-9) (2020).
19. Baccam, P., Beauchemin, C., Macken, C. A., Hayden, F. G. & Perelson, A. S. Kinetics of influenza a virus infection in humans. *J. Virol.* **80**, 7590. <https://doi.org/10.1128/JVI.01623-05> (2006).
20. Bermingham, W. H., Wilding, T., Beck, S. & Huissoon, A. SARS-CoV-2 serology: Test, test, test, but interpret with caution!. *Clin. Med.* **20**, 365–368. <https://doi.org/10.7861/clinmed.2020-0170> (2020).
21. Weiss, P. L’hypothèse du champ moléculaire et la propriété ferromagnétique. *J. Phys. Theor. Appl.* **6**, 661–690 (1907).

22. Kadanoff, L. P. More is the same; phase transitions and mean field theories. *J. Stat. Phys.* **137**, 777. <https://doi.org/10.1007/s10955-009-9814-1> (2009).
23. Parr, T., Sajid, N. & Friston, K. J. Modules or mean-fields?. *Entropy* **22**, 552 (2020).
24. Korenkov, D., Isakova-Sivak, I. & Rudenko, L. Basics of CD8 T-cell immune responses after influenza infection and vaccination with inactivated or live attenuated influenza vaccine. *Expert Rev. Vaccines* **17**, 977–987. <https://doi.org/10.1080/14760584.2018.1541407> (2018).
25. Adkins, B. *et al.* Early events in T-cell maturation. *Annu. Rev. Immunol.* **5**, 325–365. <https://doi.org/10.1146/annurev.iy.05.040187.001545> (1987).
26. Swain, S. L., McKinstry, K. K. & Strutt, T. M. Expanding roles for CD4⁺ T cells in immunity to viruses. *Nat. Rev. Immunol.* **12**, 136–148. <https://doi.org/10.1038/nri3152> (2012).
27. Channappanavar, R., Zhao, J. & Perlman, S. T cell-mediated immune response to respiratory coronaviruses. *Immunol. Res.* **59**, 118–128. <https://doi.org/10.1007/s12026-014-8534-z> (2014).
28. Kelso, A. CD4⁺ T cells limit the damage in influenza. *Nat. Med.* **18**, 200–202. <https://doi.org/10.1038/nm.2654> (2012).
29. Harty, J. T., Twinnereim, A. R. & White, D. W. CD8⁺ T cell effector mechanisms in resistance to infection. *Annu. Rev. Immunol.* **18**, 275–308. <https://doi.org/10.1146/annurev.immunol.18.1.275> (2000).
30. Koyama, S., Ishii, K. J., Coban, C. & Akira, S. Innate immune response to viral infection. *Cytokine* **43**, 336–341. <https://doi.org/10.1016/j.cyto.2008.07.009> (2008).
31. Thomson, B. J. Viruses and apoptosis. *Int. J. Exp. Pathol.* **82**, 65–76. <https://doi.org/10.1111/j.1365-2613.2001.iep0082-0065-x> (2001).
32. Parker, D. C. T cell-dependent B cell activation. *Annu. Rev. Immunol.* **11**, 331–360. <https://doi.org/10.1146/annurev.iy.11.040193.001555> (1993).
33. Weisel, F. & Shlomchik, M. Memory B cells of mice and humans. *Annu. Rev. Immunol.* **35**, 255–284. <https://doi.org/10.1146/annurev-immunol-041015-05531> (2017).
34. Kurosaki, T., Kometani, K. & Ise, W. Memory B cells. *Nat. Rev. Immunol.* **15**, 149–159. <https://doi.org/10.1038/nri3802> (2015).
35. Lobo, E. D., Hansen, R. J. & Balthasar, J. P. Antibody pharmacokinetics and pharmacodynamics. *J. Pharm. Sci.* **93**, 2645–2668. <https://doi.org/10.1002/jps.20178> (2004).
36. Havenar-Daughton, C. *et al.* Cytokine-independent detection of antigen-specific germinal center T follicular helper cells in immunized nonhuman primates using a live cell activation-induced marker technique. *J. Immunol.* **197**, 994. <https://doi.org/10.4049/jimmunol.1600320> (2016).
37. Dan, J. M. *et al.* A cytokine-independent approach to identify antigen-specific human germinal center T follicular helper cells and rare antigen-specific CD4⁺ T cells in blood. *J. Immunol.* **197**, 983. <https://doi.org/10.4049/jimmunol.1600318> (2016).
38. Wang, Z. *et al.* Recovery from severe H7N9 disease is associated with diverse response mechanisms dominated by CD8⁺ T cells. *Nat. Commun.* **6**, 6833. <https://doi.org/10.1038/ncomms7833> (2015).
39. Karlen, Y., McNair, A., Perseguers, S., Mazza, C. & Mermod, N. Statistical significance of quantitative PCR. *BMC Bioinform.* **8**, 131. <https://doi.org/10.1186/1471-2105-8-131> (2007).
40. Rutledge, R. G. & Côté, C. Mathematics of quantitative kinetic PCR and the application of standard curves. *Nucleic Acids Res.* **31**, e93. <https://doi.org/10.1093/nar/gng093> (2003).
41. McCall, M. N., McMurray, H. R., Land, H. & Almudevar, A. On non-detects in qPCR data. *Bioinformatics* **30**, 2310–2316. <https://doi.org/10.1093/bioinformatics/btu239> (2014).
42. Huang, A. T. *et al.* A systematic review of antibody mediated immunity to coronaviruses: Kinetics, correlates of protection, and association with severity. *Nat. Commun.* **11**, 4704. <https://doi.org/10.1038/s41467-020-18450-4> (2020).
43. Long, Q.-X. *et al.* Antibody responses to SARS-CoV-2 in patients with COVID-19. *Nat. Med.* **26**, 845–848. <https://doi.org/10.1038/s41591-020-0897-1> (2020).
44. He, X. *et al.* Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nat. Med.* **26**, 672–675. <https://doi.org/10.1038/s41591-020-0869-5> (2020).
45. Sekine, T. *et al.* Robust T cell immunity in convalescent individuals with asymptomatic or mild COVID-19. *bioRxiv*. <https://doi.org/10.1101/2020.06.29.174888> (2020).
46. Hoffmann, M. *et al.* SARS-CoV-2 cell entry depends on ACE2 and TMPRSS2 and is blocked by a clinically proven protease inhibitor. *Cell* **181**, 271–280.e278. <https://doi.org/10.1016/j.cell.2020.02.052> (2020).
47. Xie, X., Chen, J., Wang, X., Zhang, F. & Liu, Y. Age- and gender-related difference of ACE2 expression in rat lung. *Life Sci.* **78**, 2166–2171. <https://doi.org/10.1016/j.lfs.2005.09.038> (2006).
48. Chan, K. H. *et al.* Cross-reactive antibodies in convalescent SARS patients' sera against the emerging novel human coronavirus EMC (2012) by both immunofluorescent and neutralizing antibody tests. *J. Infect.* **67**, 130–140. <https://doi.org/10.1016/j.jinf.2013.03.015> (2013).
49. Seow, J., Graham, C., Merrick, B. *et al.* Longitudinal observation and decline of neutralizing antibody responses in the three months following SARS-CoV-2 infection in humans. *Nat. Microbiol.* **5**, 1598–1607. <https://doi.org/10.1038/s41564-020-00813-8> (2020).
50. Vatti, A. *et al.* Original antigenic sin: A comprehensive review. *J. Autoimmun.* **83**, 12–21. <https://doi.org/10.1016/j.jaut.2017.04.008> (2017).
51. Herati, R. S. *et al.* Successive annual influenza vaccination induces a recurrent oligoclonotypic memory response in circulating T follicular helper cells. *Sci. Immunol.* **2**, eaag2152. <https://doi.org/10.1126/sciimmunol.aag2152> (2017).
52. Chakravarti, B. & Abraham, G. N. Aging and T-cell-mediated immunity. *Mech. Ageing Dev.* **108**, 183–206. [https://doi.org/10.1016/S0047-6374\(99\)00009-3](https://doi.org/10.1016/S0047-6374(99)00009-3) (1999).
53. Ponnappan, S. & Ponnappan, U. Aging and immune function: molecular mechanisms to interventions. *Antioxid. Redox Signal.* **14**, 1551–1585. <https://doi.org/10.1089/ars.2010.3228> (2011).
54. Ellinghaus, D. *et al.* Genomewide association study of severe COVID-19 with respiratory failure. *N. Engl. J. Med.* <https://doi.org/10.1056/NEJMoa2020283> (2020).
55. Wein, A. N. *et al.* CXCR6 regulates localization of tissue-resident memory CD8 T cells to the airways. *J. Exp. Med.* **216**, 2748–2762. <https://doi.org/10.1084/jem.20181308> (2019).
56. Zens, K. D. & Farber, D. L. Memory CD4 T cells in influenza. *Curr. Top. Microbiol. Immunol.* **386**, 399–421. https://doi.org/10.1007/82_2014_401 (2015).
57. Cevik, M. *et al.* SARS-CoV-2, SARS-CoV, and MERS-CoV viral load dynamics, duration of viral shedding, and infectiousness: A systematic review and meta-analysis. *Lancet Microbe* **2**, e13–e22. [https://doi.org/10.1016/S2666-5247\(20\)30172-5](https://doi.org/10.1016/S2666-5247(20)30172-5) (2021).
58. van Kampen, J. J. A. *et al.* Duration and key determinants of infectious virus shedding in hospitalized patients with coronavirus disease-2019 (COVID-19). *Nat. Commun.* **12**, 267. <https://doi.org/10.1038/s41467-020-20568-4> (2021).
59. Vibholm, L. K. *et al.* SARS-CoV-2 persistence is associated with antigen-specific CD8 T-cell responses. *EBioMedicine* **64**, 103230. <https://doi.org/10.1016/j.ebiom.2021.103230> (2021).
60. Mathew, D. *et al.* Deep immune profiling of COVID-19 patients reveals distinct immunotypes with therapeutic implications. *Science* **369**, eabc8511. <https://doi.org/10.1126/science.abc8511> (2020).
61. Fallet, B. *et al.* Chronic viral infection promotes efficient germinal center B cell responses. *Cell Rep.* **30**, 1013–1026.e1017. <https://doi.org/10.1016/j.celrep.2019.12.023> (2020).

62. Markus, A., Lebenthal-Loinger, I., Yang, I. H., Kinchington, P. R. & Goldstein, R. S. An in vitro model of latency and reactivation of varicella zoster virus in human stem cell-derived neurons. *PLoS Pathog.* **11**, e1004885. <https://doi.org/10.1371/journal.ppat.1004885> (2015).
63. O'Neill, L. A. J. & Netea, M. G. BCG-induced trained immunity: Can it offer protection against COVID-19? *Nat. Rev. Immunol.* **20**, 335–337. <https://doi.org/10.1038/s41577-020-0337-y> (2020).
64. Friston, K. J. *et al.* Bayesian model reduction and empirical Bayes for group (DCM) studies. *Neuroimage* **128**, 413–431. <https://doi.org/10.1016/j.neuroimage.2015.11.015> (2016).
65. Friston, K., Parr, T. & Zeidman, P. Bayesian model reduction. arXiv preprint arXiv:1805.07092 (2018).
66. Kass, R. E. & Raftery, A. E. Bayes factors. *J. Am. Stat. Assoc.* **90**, 773–795. <https://doi.org/10.1080/01621459.1995.10476572> (1995).
67. Kass, R. E. & Steffey, D. Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical Bayes models). *J. Am. Stat. Assoc.* **407**, 717–726 (1989).
68. Friston, K. J. *et al.* Parcels and particles: Markov blankets in the brain. arXiv preprint arXiv:2007.09704 (2020).
69. Schwabl, F. *Statistical Mechanics* 327–404 (Springer Berlin Heidelberg, 2002).
70. Haken, H. *Synergetics: An Introduction. Non-equilibrium Phase Transition and Self-Selforganisation in Physics, Chemistry and Biology* (Springer Verlag, 1983).
71. Clapham, H. E. *et al.* Modelling virus and antibody dynamics during dengue virus infection suggests a role for antibody in virus clearance. *PLoS Comput. Biol.* **12**, e1004951. <https://doi.org/10.1371/journal.pcbi.1004951> (2016).
72. Zitzmann, C. & Kaderali, L. Mathematical analysis of viral replication dynamics and antiviral treatment strategies: From basic models to age-based multi-scale modeling. *Front. Microbiol.* <https://doi.org/10.3389/fmicb.2018.01546> (2018).
73. Folegatti, P. M. *et al.* Safety and immunogenicity of the ChAdOx1 nCoV-19 vaccine against SARS-CoV-2: A preliminary report of a phase 1/2, single-blind, randomised controlled trial. *Lancet* [https://doi.org/10.1016/S0140-6736\(20\)31604-4](https://doi.org/10.1016/S0140-6736(20)31604-4) (2020).
74. Chen, Z. & John Wherry, E. T cell responses in patients with COVID-19. *Nat. Rev. Immunol.* <https://doi.org/10.1038/s41577-020-0402-6> (2020).
75. Okoye, A. A. & Picker, L. J. CD4(+) T-cell depletion in HIV infection: Mechanisms of immunological failure. *Immunol. Rev.* **254**, 54–64. <https://doi.org/10.1111/imr.12066> (2013).
76. De Biasi, S. *et al.* Marked T cell activation, senescence, exhaustion and skewing towards TH17 in patients with COVID-19 pneumonia. *Nat. Commun.* **11**, 3434. <https://doi.org/10.1038/s41467-020-17292-4> (2020).
77. Adams, R. A., Aponte, E., Marshall, L. & Friston, K. J. Active inference and oculomotor pursuit: The dynamic causal modelling of eye movements. *J. Neurosci. Methods* **242**, 1–14. <https://doi.org/10.1016/j.jneumeth.2015.01.003> (2015).
78. Stephan, K. E. *et al.* Nonlinear dynamic causal models for fMRI. *Neuroimage* **42**, 649–662. <https://doi.org/10.1016/j.neuroimage.2008.04.262> (2008).
79. Kiebel, S. J., Garrido, M. I. & Friston, K. J. Dynamic causal modelling of evoked responses: The role of intrinsic connections. *Neuroimage* **36**, 332–345. <https://doi.org/10.1016/j.neuroimage.2007.02.046> (2007).
80. Moran, R., Pinotsis, D. A. & Friston, K. Neural masses and fields in dynamic causal modeling. *Front. Comput. Neurosci.* **7**, 57. <https://doi.org/10.3389/fncom.2013.00057> (2013).
81. De Boer, R. J., Perelson, A. S. & Kevrekidis, I. G. Immune network behavior—I. From stationary states to limit cycle oscillations. *Bull. Math. Biol.* **55**, 745–780. [https://doi.org/10.1016/S0092-8240\(05\)80188-0](https://doi.org/10.1016/S0092-8240(05)80188-0) (1993).
82. Li, Y., Jin, L. & Chen, T. The effects of secretory IgA in the mucosal immune system. *Biomed. Res. Int.* **2020**, 2032057. <https://doi.org/10.1155/2020/2032057> (2020).
83. Boyaka, P. N. Inducing mucosal IgA: A challenge for vaccine adjuvants and delivery systems. *J. Immunol.* **199**, 9–16 (2017).
84. Macpherson, A. J., McCoy, K. D., Johansen, F. E. & Brandtzaeg, P. The immune geography of IgA induction and function. *Mucosal Immunol.* **1**, 11–22. <https://doi.org/10.1038/mi.2007.6> (2008).
85. Ma, H. *et al.* Serum IgA, IgM, and IgG responses in COVID-19. *Cell. Mol. Immunol.* **17**, 773–775. <https://doi.org/10.1038/s41423-020-0474-z> (2020).
86. Paterson, R. W. *et al.* The emerging spectrum of COVID-19 neurology: Clinical, radiological and laboratory findings. *Brain* <https://doi.org/10.1093/brain/awaa240> (2020).
87. Wei, H. *et al.* Bayesian fusion and multimodal DCM for EEG and fMRI. *Neuroimage* **211**, 116595. <https://doi.org/10.1016/j.neuroimage.2020.116595> (2020).
88. Baecher-Allan, C., Kaskow, B. J. & Weiner, H. L. Multiple sclerosis: Mechanisms and immunotherapy. *Neuron* **97**, 742–768. <https://doi.org/10.1016/j.neuron.2018.01.021> (2018).
89. Cooray, G. K., Sundgren, M. & Brismar, T. Mechanism of visual network dysfunction in relapsing-remitting multiple sclerosis and its relation to cognition. *Clin. Neurophysiol.* **131**, 361–367. <https://doi.org/10.1016/j.clinph.2019.10.029> (2020).
90. Rosch, R. E. *et al.* NMDA-receptor antibodies alter cortical microcircuit dynamics. *Proc. Natl. Acad. Sci.* **115**, E9916–E9925. <https://doi.org/10.1073/pnas.1804846115> (2018).
91. Friston, K. J., Daunizeau, J., Kilner, J. & Kiebel, S. J. Action and behavior: A free-energy formulation. *Biol. Cybern.* **102**, 227–260. <https://doi.org/10.1007/s00422-010-0364-z> (2010).
92. Friston, K., Samothrakis, S. & Montague, R. Active inference and agency: Optimal control without cost functions. *Biol. Cybern.* **106**, 523–541. <https://doi.org/10.1007/s00422-012-0512-8> (2012).
93. Friston, K. J., Rosch, R., Parr, T., Price, C. & Bowman, H. Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* **77**, 388–402. <https://doi.org/10.1016/j.neubiorev.2017.04.009> (2017).
94. Moutoussis, M., Trujillo-Barreto, N. J., El-Deredy, W., Dolan, R. J. & Friston, K. J. A formal model of interpersonal inference. *Front. Hum. Neurosci.* **8**, 160. <https://doi.org/10.3389/fnhum.2014.00160> (2014).
95. Kenney, M. J. & Ganta, C. K. Autonomic nervous system and immune system interactions. *Compr. Physiol.* **4**, 1177–1200. <https://doi.org/10.1002/cphy.c130051> (2014).
96. Barrios-Payán, J. *et al.* The contribution of the sympathetic nervous system to the immunopathology of experimental pulmonary tuberculosis. *J. Neuroimmunol.* **298**, 98–105. <https://doi.org/10.1016/j.jneuroim.2016.07.012> (2016).
97. Alaniz, R. C. *et al.* Dopamine β -hydroxylase deficiency impairs cellular immunity. *Proc. Natl. Acad. Sci.* **96**, 2274–2278. <https://doi.org/10.1073/pnas.96.5.2274> (1999).
98. Rosas-Ballina, M. & Tracey, K. J. The neurology of the immune system: Neural reflexes regulate immunity. *Neuron* **64**, 28–32. <https://doi.org/10.1016/j.neuron.2009.09.039> (2009).
99. Tracey, K. J. Reflex control of immunity. *Nat. Rev. Immunol.* **9**, 418–428. <https://doi.org/10.1038/nri2566> (2009).
100. Raony, Í. *et al.* Psycho-neuroendocrine-immune interactions in COVID-19: Potential impacts on mental health. *Front. Immunol.* <https://doi.org/10.3389/fimmu.2020.01170> (2020).

Acknowledgements

The Wellcome Centre for Human Neuroimaging is supported by core funding from Wellcome [203147/Z/16/Z]. AB is supported by a Medical Research Council doctoral studentship [MR/N013867/1]. KJF is a Wellcome Principle Research Fellow [088130/Z/09/Z].

Author contributions

T.P. wrote and performed the simulations. T.P., A.B., P.Z., A.G., A.J.B., R.M., and K.J.F. wrote and edited the manuscript. The model inversions used software routines written by K.J.F.

Competing interests

K.J.F. is a member of the Independent SAGE advisory group in the UK. The other authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to T.P.

Reprints and permissions information is available at www.nature.com/reprints.

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



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

© The Author(s) 2021, corrected publication 2021