A speed limit on serial strain replacement from original antigenic sin

Lauren McGough[a,1] [ID] and Sarah Cobey[a]

Many pathogens evolve to escape immunity, yet it remains difficult to predict whether immune pressure will lead to diversification, serial replacement of one variant by another, or more complex patterns. Pathogen strain dynamics are mediated by cross-protective immunity, whereby exposure to one strain partially protects against infection by antigenically diverged strains. There is growing evidence that this protection is influenced by early exposures, a phenomenon referred to as original antigenic sin (OAS) or imprinting. In this paper, we derive constraints on the emergence of the pattern of successive strain replacements demonstrated by influenza, SARS-CoV-2, seasonal coronaviruses, and other pathogens. We find that OAS implies that the limited diversity found with successive strain replacement can only be maintained if R_0 is less than a threshold set by the characteristic antigenic distances for cross-protection and for the creation of new immune memory. This bound implies a "speed limit" on the evolution of new strains and a minimum variance of the distribution of infecting strains in antigenic space at any time. To carry out this analysis, we develop a theoretical model of pathogen evolution in antigenic space that implements OAS by decoupling the antigenic distances required for protection from infection and strain-specific memory creation. Our results demonstrate that OAS can play an integral role in the emergence of strain structure from host immune dynamics, preventing highly transmissible pathogens from maintaining serial strain replacement without diversification.

traveling waves | blunting | transmissibility

Antigenically variable pathogens consist of immunologically distinct strains whose spatiotemporal dynamics depend on the pathogen's capacity to diversify and on the extent to which host immune responses elicited by one strain protect against infection with other strains. This interdependence gives rise to qualitatively distinct patterns of diversity or strain structures (1–6). Examples include pathogens with high antigenic diversity at the host population scale [e.g., *Neisseria meningitidis* (7–10), enteroviruses (11)] and others with lower antigenic diversity at any given time but fast turnover [e.g., influenza A in humans (12, 13), seasonal coronaviruses (14), SARS-CoV-2 (15–17), and others (18)].

Successive strain replacement typifies the dynamics of many common respiratory pathogens, yet it only arises in transmission models in some conditions (1, 2, 6, 12, 19–30). Large individual-based models have found that low antigenic diversity and fast turnover result from low mutation rates and strong cross-immunity between strains (6, 31). Short-term strain-transcending immunity (12, 27) and punctuated antigenic changes (19, 32) have been hypothesized to be essential for serial antigenic replacement. Intrinsic transmissibility also plays a role, with higher R_0 promoting frequent emergence of antigenically novel strains (30, 31). From a theoretical standpoint, successive strain replacement corresponds to epidemic dynamics that admit traveling wave solutions, in which infections create immunity that pushes existing strains in a single direction in phenotypic space (24, 33–36). These dynamics agree with representations of H3N2 antigenic evolution, with strains following a low-dimensional trajectory in a higher-dimensional antigenic space (13, 37).

A common assumption in these investigations is that individuals acquire immunity specific to any strain that infects them, regardless of infection history (3, 20, 22, 23, 33). Decades of experimental and observational evidence of original antigenic sin (OAS) have demonstrated that secondary responses expand memory B cell responses that target previously encountered epitopes, limiting the generation of responses to new sites (38–49). This is especially apparent in responses to less immunogenic influenza vaccines (42, 50, 51). Proposed mechanisms for this limited response to new sites include clearance of antigen by preexisting immunity (52, 53) and the reduced requirements for memory compared to naive B cell activation (e.g., refs. 54, 55, and 56). A straightforward corollary

Significance

Evolving pathogens continually generate genetic diversity through random mutations. Sometimes, these mutations create new strains or variants that partially escape immune memory, allowing them to spread more easily. Why some pathogens are dominated by one variant at a time while others exhibit ongoing diversification or cycling is incompletely understood. Research on this problem has ignored the immune system's strong bias to the familiar, which makes protection dependent on the order in which strains are encountered. This work demonstrates that this order dependence makes serial strain replacement even less expected and available only to relatively slow pathogens, and raises the question of whether influenza, SARS-CoV-2, and other viruses that currently follow this pattern may diversify in the future.

Author affiliations: [a]Department of Ecology and Evolution The University of Chicago, Chicago, IL 60637

[1]To whom correspondence may be addressed. Email: mcgough@uchicago.edu.

is that no protection can be derived from nonexistent responses to new sites, and thus individuals infected or vaccinated with the same strain may differ in their protection against other strains, depending on which epitopes they target (57).

OAS might sometimes boost antibody responses to a site that are cross-reactive, blunting new responses to that site, but are not necessarily so protective. Decoy nonneutralizing epitopes are a feature of some pathogens (58, 59), but in influenza, this discrepancy between reactivity and protection appears to arise occasionally from OAS. For instance, middle-aged adults infected with H3N2 sometimes boost antibody titers that bind well by enzyme linked immunoabsorbent assay (ELISA) but show negligible neutralization activity against the infecting strain (60). Ferrets infected with one influenza A subtype boost their stalk antibodies to that subtype on later infection with another subtype, blunting new responses (61). The same pattern appears in children infected with H1N1 before H3N2 (61). These recalled antibodies appear to be very poor binders to the second subtype, and although the in vivo consequences of hundreds-fold reductions in affinity are uncertain (62), it is plausible they might be accompanied by reduced protection, as suggested by experimental infections in mice (63). Thus, OAS might influence the protection derived from secondary exposures in two ways: first, by preventing any response to new sites, and second, by recalling memory responses to new but similar sites that poorly bind and weakly neutralize the new site. Either mechanism can in theory prevent the creation of protective memory to the infecting strain.

It is not known how OAS might affect pathogen evolution. Transmission models involving OAS have assumed that only one strain circulates at any time and always mutates between seasons (64, 65), and thus have not addressed how OAS impacts antigenic evolution. These transmission models predict that OAS impacts population immunity, leading to "immune blind spots" in different age groups (65), which might relate to recent evidence of ~24-y cycles in the induction of strong antibody responses to H3N2 (66).

We investigate the consequences of OAS for pathogen evolution, and specifically for the conditions that permit successive strain replacement. Our mathematical models of host-pathogen coevolution incorporate OAS via the mechanisms described above. Both models allow for infection without creation of strain-specific memory through tunable parameters that allow cross-reactive responses to blunt the generation of new ones without conferring protection at that epitope. One model assumes a cross-reactive response at one epitope is sufficient to block the induction of new memory to the rest, and the other assumes that any epitope that is sufficiently diverged will elicit a new response. Nature is likely somewhere in between, depending on the type and number of secondary exposures (53, 63, 67, 68). We show that these contrasting assumptions do not change our main result: OAS implies an upper bound on the basic reproduction number R_0 of pathogens that can exhibit successive strain replacement, and this bound implies limits on the speed of evolution, the standing antigenic diversity, and the time to the most recent common ancestor in this regime. OAS thus narrows the conditions in which serial strain replacement is likely to occur.

Model

We separately consider the dynamics of protection and memory creation in individuals before describing pathogen dynamics in the host population.

Host-Scale Protection. Pathogen strains and strain-specific memories are parameterized as real vectors in an abstract, continuous d-dimensional antigenic space in which the distance between two points represents the extent to which a strain-specific memory at one point protects against infection with a strain at the other point, as detailed mathematically below. The model assumes that each of the d dimensions of antigenic space corresponds to a physical region bound by antibodies, and the footprints of antibodies binding this region are entirely contained within it. We refer to these antibody-binding regions as epitopes. A similar model with each axis corresponding to one epitope was used in ref. 69.

Individuals infected with a strain of the pathogen are also infectious. An infected individual can transmit either the infecting strain or a mutant strain. Infected individuals expose on average R_0 other individuals during their infection, where R_0 is the basic reproduction number and does not vary by strain. The duration of infection also does not vary by strain.

Individuals are born without immunological memory. A naive individual always becomes infected when exposed to a strain and forms strain-specific memory on infection. A nonnaive individual who becomes infected might develop memory specific to that strain. The set $\mathcal{S}$ of strains $\vec{s}$ to which an individual has strain-specific memory at time t is their memory profile at t. An individual has at most one specific memory targeting (i.e., derived from infection with) each strain. The number of strains in $\mathcal{S}$ equals m. In general, the number of strains in memory, m, would be expected to vary among individuals, though the population-scale model will assume that a fixed value of m is a reasonable approximation for computing protection.

Memory decreases the probability of infection on exposure. Individuals cannot be reinfected by strains in their memory profile. Memory is cross-protective, meaning that strain-specific memory decreases the probability of infection with similar strains (Fig. 1). Cross-protection declines exponentially with the Euclidean antigenic distance between the strain in memory $\vec{s}$ and the challenge strain $\vec{s}'$,

$$\Pr(\text{memory at } \vec{s} \text{ protects host} \mid \text{host is exposed to } \vec{s}')$$
$$= e^{-||\vec{s}-\vec{s}'||/r}, \qquad [1]$$

with $||\vec{s} - \vec{s}'|| = \left(\sum_{i=1}^{d}(s_i - s_i')^2\right)^{1/2}$. The quantity r is the cross-protection distance. It is an independent parameter that determines how antigenic change affects the protectiveness of
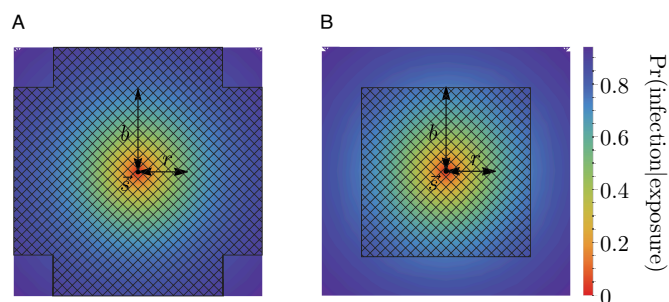
Fig. 1. Cross-protection and blunting from a single memory in 2-dimensional antigenic space. A strain-specific memory $\vec{s}$ at the center protects against infection by nearby strains according to Eq. **1**. The distance over which protection drops by a factor of $1/e$ is r; the blunting distance is b; hatched regions represent excluded regions in the (*A*) every- and (*B*) any-epitope models. Only infections with strains outside these regions would generate new memory. Here, $b = 1.7$ and $r = 1.25$.

existing memory to related strains. If r is high, cross-protection is high, and large mutations in antigenic space are required for protection to drop.

Eq. **1** defines the probability that a single strain-specific memory protects against infection by another strain. Each strain $\vec{s}$ in an individual's memory profile $\mathcal{S}$ contributes to protection against a challenge strain $\vec{s}\,'$. The probability of infection for an individual with memory profile $\mathcal{S}$ exposed to a challenge strain $\vec{s}\,'$ is given by the product of the probabilities of each memory separately failing to protect,

$$\text{Pr(infection at } \vec{s}\,' \,|\, \text{exposure to } \vec{s}\,') = \prod_{\vec{s} \in \mathcal{S}} (1 - e^{-||\vec{s}-\vec{s}\,'||/r}). \quad [2]$$

Implicit in Eq. **2** is the assumption that protection against infection depends on the unordered set of strains in the memory profile, yet we have claimed that OAS induces a dependence on the order of exposure. Order dependence arises from the dynamics of memory creation.

Creation of New Memory. Infections and vaccinations do not always generate new specific memory to the challenge strain. Whether an individual generates a strain-specific memory depends on their existing memory profile. We model OAS via blunting, and consider two models, the "every-epitope" model and the "any-epitope" model, corresponding to the limits of biologically plausible scenarios. In each, a blunting distance, b, sets a threshold antigenic distance for developing new memory.

In the every-epitope model, the challenge strain $\vec{s}\,'$ must escape immunity at every epitope to create strain-specific memory (Fig. 1); escape at only some epitopes can be overcome by adaptive immune responses, such as antibodies, binding conserved regions (69). More precisely, an individual challenged with a strain $\vec{s}\,'$ creates strain-specific memory if $\vec{s}\,'$ differs from every strain $\vec{s}$ in their memory profile $\mathcal{S}$ by an antigenic distance of at least b in each of its d antigenic space coordinates,

every-epitope condition for memory creation:
$$\min_{\vec{s} \in \mathcal{S}} \min_{i=1,\dots,d} |s_i' - s_i| \geq b. \quad [3]$$

In the any-epitope model, it is enough for the challenge strain $\vec{s}\,'$ to escape immunity in a single epitope to create strain-specific memory (Fig. 1) (3). An individual challenged with a strain $\vec{s}\,'$ creates strain-specific memory if $\vec{s}\,'$ differs from every strain $\vec{s}$ in their memory profile $\mathcal{S}$ by at least b in at least one epitope (which may not be the same epitope across $\mathcal{S}$),

any-epitope condition for memory creation:
$$\min_{\vec{s} \in \mathcal{S}} \max_{i=1,\dots,d} |s_i' - s_i| \geq b. \quad [4]$$

These models can be conceptualized as different types of exposures. The every-epitope model could represent an exposure to a small dose of antigen, for example via natural infection, where the presence of any neutralizing memory can limit pathogen replication, reducing antigen availability enough to prevent the generation of new memory. The any-epitope model could represent an exposure to a high dose of antigen or an immunogenic vaccine that forces a response to diverged epitopes even in the presence of memory that would ordinarily suppress a new response. This effect has been shown in animal models since the earliest observations of OAS (46), and more recently in molecular detail in mice (48, 68).

Throughout the remaining analysis, we assume that immune memory can cross-react with a larger set of strains than it can protect against ($b > r$). That is, it can suppress the formation of new memory against some infecting strains. While it is not biologically impossible that r would be greater than b, this would require enough antigen to provoke a memory response without causing an infection. This might occur through vaccination but will not be explored here.

Population-Scale Infection-Immunity Coupling. At the population scale, our model is equivalent to the model of pathogen-immune coevolution developed in ref. 36 with the added assumptions that each dimension corresponds to an epitope in the proper coordinates and that each epitope antigenically evolves at a rate that does not depend on time. The population-scale description is formulated in terms of two densities in antigenic space: the total number of individuals infected by each strain $\vec{s}$ at time t and the total number of individuals with strain-specific memory to $\vec{s}$ at time t, denoted $n(\vec{s}, t)$ and $h(\vec{s}, t)$, respectively. In transitioning to a density-based approach, we trade complete information about each individual's memory profile for tractability. This is similar to the choice of a previous cohort-level description for studying OAS (65), although the motivation is different as we do not incorporate age structure.

Strain fitness is determined by the number of new infections per infectiousness period, which depends on individuals' immunity. The expected protection against infection upon exposure to strain $\vec{s}$ conferred by a single memory chosen uniformly at random from $h(\vec{s}\,', t)$ is denoted $c(\vec{s}, t)$. Using the definition in Eq. **1**, this is equal to

$$c(\vec{s}, t) = \frac{1}{mN_h} \int_{\vec{s}\,' \text{ in antigenic space}} d\vec{s}\,' \, h(\vec{s}\,', t) e^{-||\vec{s}-\vec{s}\,'||/r}, \quad [5]$$

where N_h is the size of the population and the factor mN_h adjusts for the normalization of $h(\vec{s}\,', t)$. In analogy with Eq. **2**, we have

$$\text{Pr(infection at } \vec{s} \,|\, \text{exposure at } \vec{s}) \approx (1 - c(\vec{s}, t))^m. \quad [6]$$

Eqs. **5** and **6** rely on an assumption that the average protection against infection across individuals is well approximated by the average protection in a system where the $h(\vec{s}, t)$ memories are distributed randomly among individuals, with each individual having the same number of memories, m (see *SI Appendix* for further discussion on this point). Defining R_0, the basic reproduction number of the pathogen, to be the expected number of people an infected individual encounters during their infection and would infect (assuming complete susceptibility), the fitness of strain $\vec{s}$ at time t is

$$f(\vec{s}, t) = \log \left(R_0 \left(1 - c(\vec{s}, t) \right)^m \right), \quad [7]$$

as in ref. 36.

We model the successive strain replacement regime using the traveling wave model described in ref. 36. Traveling wave solutions arise from stochastic coupled dynamical equations between the infection density and the memory density that take into account fitness differences among strains, mutations, noise in the transmission process, and memory addition and deletion:

$$\partial_t n(\vec{s}, t) = f(\vec{s}, t) n(\vec{s}, t) + D \vec{\nabla}^2 n(\vec{s}, t) + \sqrt{n(\vec{s}, t)} \, \eta(\vec{s}, t), \quad [8]$$

$$\partial_t h(\vec{s}, t) = n(\vec{s}, t) - \frac{N_I(t)}{mN_h} h(\vec{s}, t). \quad [9]$$

The first equation (Eq. **8**) states that as the pathogen spreads, each strain grows or shrinks according to its fitness (Eq. **7**); mutations are frequent and have small effect, such that the diffusion approximation is appropriate with diffusion constant D; and pathogen spread is stochastic, with Gaussian process noise, where $\eta(\vec{s}, t)$ is a unit Gaussian ($\langle \eta(\vec{s}, t) \rangle = 0$, $\langle \eta(\vec{s}, t)\eta(\vec{s}', t') \rangle = \delta(\vec{s} - \vec{s}')\delta(t - t')$), and the prefactor $\sqrt{n(\vec{s}, t)}$ sets the mean and SD. The second equation (Eq. **9**) states that the change in memory at a given time is determined by memory addition at the locations of infections and by memory deletion uniformly at random: each person acquires a new strain-specific memory at $\vec{s}$ when they are infected by strain $\vec{s}$, and at that time, they lose one memory from their memory profile $\mathcal{S}$, such that each person maintains a fixed number of m memories over time. Time is modeled in units of the duration of infection.

Successive Strain Replacement Dynamics. Successive strain replacement dynamics are conceptualized as an infection distribution that takes the form of a localized clump that maintains its shape as it moves through antigenic space at a constant speed and direction and is Gaussian in the direction of motion to reflect the assumption that the mutation process is diffusive. The localization implies that strains are immunologically well defined, with low antigenic diversity around any point at any time. The constant speed and direction amount to assuming that the effect of stochasticity in orthogonal directions is negligible. While the speed at which the wave travels is a derived quantity determined by the transmission and immunity parameters governing the dynamics, its direction of travel is determined by the relative rates at which different epitopes antigenically evolve, an independent input to the system which we parameterize by the angles between the direction the wave is traveling and the axes for individual epitopes, $\theta_1, \ldots, \theta_d$ (Fig. 2*C*).
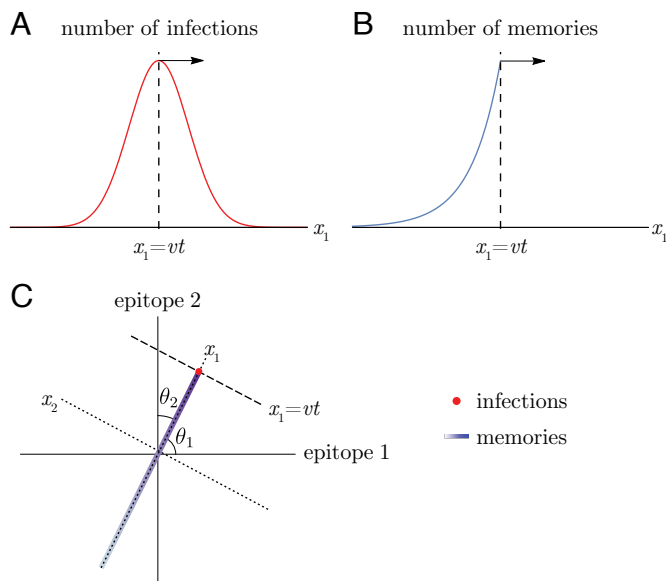
Fig. 2. Successive strain replacement as traveling waves in antigenic space. (*A*) Strain replacement is modeled as a localized infection distribution moving along a one-dimensional trajectory in *d*-dimensional antigenic space with constant velocity $\vec{v}$. (*B*) The memory distribution is an exponentially decaying distribution with wavefront at $x_1 = vt$, localized in orthogonal directions. (*C*) The axes aligned with the direction of motion of the infection and memory distributions are not the same as the epitope-aligned axes. The localized infection peak is visualized as a point at the location of the wavefront.

Successive strain dynamics are most simply expressed in coordinates where one axis is aligned with the direction of motion, with the remaining axes aligned orthogonally. We use $\vec{x}$ (rather than $\vec{s}$) to indicate "wave-adapted" antigenic space coordinates where the x_1 axis lies along the axis of antigenic evolution.

As in ref. 36, the infection distribution takes the form

$$n(\vec{x}, t) = \frac{N_I}{(2\pi\sigma^2)^{1/2}} e^{-\frac{(x_1 - vt)^2}{\sigma^2}} \rho(x_2, \ldots, x_d), \qquad [\mathbf{10}]$$

where N_I is the number of infected individuals, v is the wave velocity, σ is the SD of the Gaussian in the direction of motion, and ρ is a normalized distribution. Both the Gaussian and ρ are taken to be localized to a region of linear dimension much less than the cross-protection distance r ($\sigma \ll r$) (Fig. 2*A*). We are interested in the case where the cross-protection distance is less than the blunting distance, and thus these are also localized to a region of linear dimension much less than the blunting distance b ($\sigma \ll b$).

When the pathogen undergoes serial strain replacement, represented by the infection distribution in Eq. **10**, the population develops a memory distribution in the shape of a wavefront with maximum at the peak of the infection distribution that falls off exponentially fast in the direction of earlier infections (Fig. 2*B*). There is no specific memory ahead of the peak of infections. The wavefront travels at the same speed, v, as the infection distribution, and in the same direction. Since in this model each person has m memories in their memory profile, the total number of memories in the population (the integral over the entire memory distribution) equals the size of the population, N_h, times m.

As derived in ref. 36, the speed of antigenic evolution, v; SD of the infection distribution, σ; and number of infected individuals, N_I, are not independently specified. Given the dynamics specified by Eqs. **8** and **9**, they are determined by the cross-protection distance, r; the number of individuals in the population, N_h; the diffusion coefficient characterizing the effect of mutations, D; the number of memories per person, m; and the basic reproduction number, R_0. See *Materials and Methods* for the relevant expressions.

Results

We first describe static properties that the population memory and protection distributions must satisfy to be consistent with the two blunting models, regardless of the transmission dynamics that give rise to them. We then derive the consequences for pathogen evolution: to obtain serial strain replacement, the speed of evolution, v; the strain diversity, σ; and the time to the most recent common ancestor, t_{MRCA}, are constrained by an upper bound on the basic reproduction number, R_0. This upper bound is determined by the characteristic distances defining memory creation and cross-protection (b and r, respectively), and by how changes in individual epitopes impact protection (the blunting model).

Blunting Neighborhoods. In both the every-epitope model and the any-epitope model, blunting implies a geometric condition that no individual's memory profile contains specific memory to strains that are too close to one another. The precise definition of "too close" differs between the models.

The every-epitope assumption (Eq. **3**) implies that for every pair of strains $\vec{s}, \vec{s}'$ in a host's memory profile $\mathcal{S}$, the difference
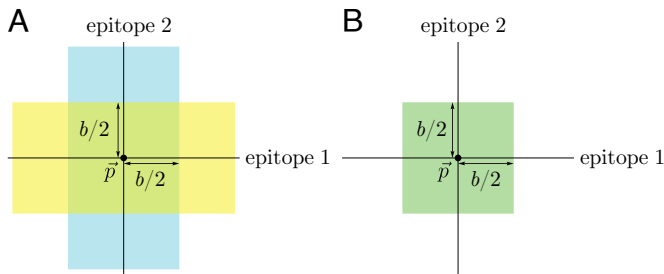
Fig. 3. Blunting neighborhoods about reference point $\vec{p}$ for two blunting models, shown in $2d$ antigenic space. If an individual has memory to a strain in a blunting neighborhood of a given reference point and is later infected by a different strain in that neighborhood, no new memory will be created. (*A*) Blunting neighborhoods for epitope 1 (yellow) and epitope 2 (blue) about $\vec{p}$ in the every-epitope model. (*B*) Blunting neighborhood about $\vec{p}$ in the any-epitope model.

$|s_i - s'_i|$ is at least b in each dimension $i = 1, \ldots, d$. This means that given any reference point $\vec{p}$ in antigenic space, for any dimension i, there exists at most one strain in $\mathcal{S}$ that is a distance less than or equal to $\frac{b}{2}$ from $\vec{p}$ in that dimension. The set of all such strains defines d stripe-shaped regions in which there can be at most one memory: the set of points $\vec{p}\,'$ with $|p_i - p'_i| < \frac{b}{2}$ for at least one value of i (Fig. 3*A*). Each stripe has width b parallel to the ith axis and extends infinitely in the remaining $(d - 1)$ axes (derivation in *SI Appendix*). We refer to these stripes as the blunting neighborhoods of $\vec{p}$ in the every-epitope model and denote them $^{\mathrm{E}}\langle\langle\vec{p}\rangle\rangle^i_b$.

The any-epitope assumption (Eq. **4**) implies that for every pair of strains $\vec{s}, \vec{s}\,'$ in a host's memory profile $\mathcal{S}$, the difference $|s_i - s'_i|$ is at least b in at least one dimension $i = 1, \ldots, d$. This means that given any reference point $\vec{p}$ in antigenic space, there exists at most one strain in $\mathcal{S}$ that is within the distance $\frac{b}{2}$ from $\vec{p}$ in all dimensions, defining a d-dimensional cube with side length b centered at $\vec{p}$ (Fig. 3*B*) (derivation in *SI Appendix*). We refer to this cube as the blunting neighborhood of $\vec{p}$ in the any-epitope model and denote it $^{\mathrm{A}}\langle\langle\vec{p}\rangle\rangle_b$.

Order-Dependent Protection. Both proposed blunting models are consistent with a major observation motivating our study of OAS: they imply that individuals' immunity depends not only on the set of strains to which they have been exposed but also on the order of exposure (Fig. 4) (60, 70).

Consider the simplest case of a pathogen whose evolution takes place in one-dimensional antigenic space. In one dimension, the every-epitope model and the any-epitope model are equivalent. Two strains are present: one, $\vec{s}_i$, at the origin and the other, $\vec{s}_j$, at $b/2$. Consider two different hosts, h_1 and h_2, both of whom have been infected by strains $\vec{s}_i$ and $\vec{s}_j$ but in different orders: h_1 was
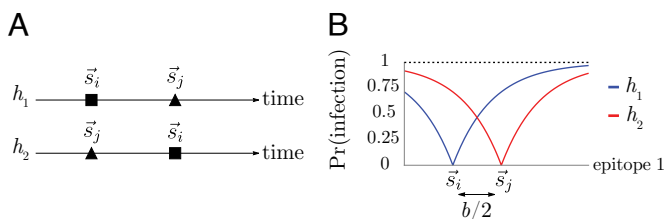
infected by $\vec{s}_i$ followed by $\vec{s}_j$, and h_2 was infected by $\vec{s}_j$ followed by $\vec{s}_i$ (Fig. 4*A*). Host h_1 develops specific memory to strain $\vec{s}_i$ but not to strain $\vec{s}_j$, whereas host h_2 develops specific memory to strain $\vec{s}_j$ but not $\vec{s}_i$ (Fig. 4*B*). Given these memory profiles, the probability that each host is infected by any strain $\vec{s}\,'$ is given by

$$\mathrm{Pr}(\text{infection by } \vec{s}\,'|\text{exposure at } \vec{s}\,', h_1) = (1 - e^{-|s'|/r}), \qquad \textbf{[11]}$$
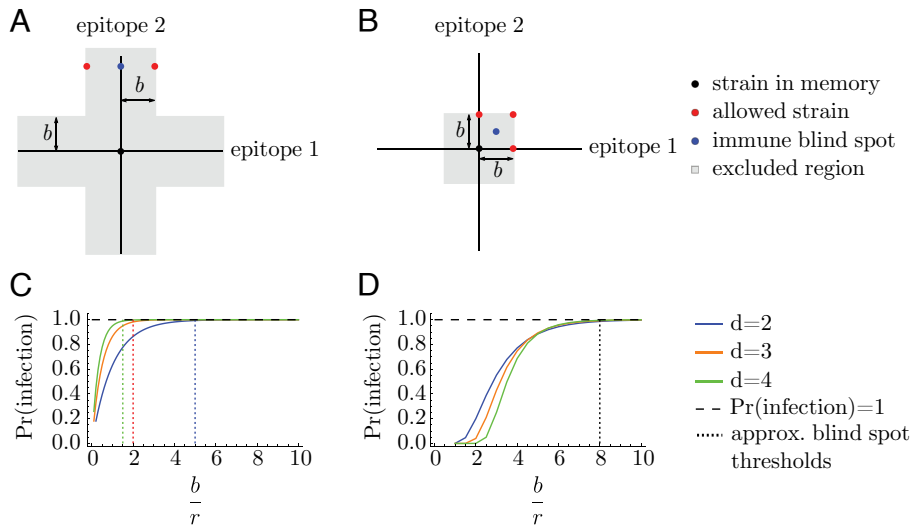
$$\mathrm{Pr}(\text{infection by } \vec{s}\,'|\text{exposure at } \vec{s}\,', h_2) = (1 - e^{-|\frac{b}{2} - s'|/r}). \qquad \textbf{[12]}$$

For example, if host h_1 is exposed to the strain $\vec{s}\,'$ at b, they will be infected with probability $(1 - e^{-b/r})$, whereas if host h_2 is exposed to the same $\vec{s}\,' = b$, they will be infected with probability $(1 - e^{-b/2r})$, a smaller value. Although h_1 and h_2 have the same set of exposures in their history, blunting dynamics lead to h_1 having less protection to $\vec{s}\,'$ than h_2 (Fig. 4*B*). Both blunting models therefore predict order-dependent protection to antigenically diverged strains.

In higher dimensions, the mechanism by which the every- and the any-epitope models exhibit order-dependent protection is essentially the same as in one dimension, but because blunting neighborhoods are infinitely extended in the every-epitope model and localized in the any-epitope model, they differ in their implications for heterogeneity of protection in the population. In the every-epitope model, order dependence can induce strong heterogeneity of protection, as a pathogen with some slowly evolving (i.e., conserved) epitopes and at least one fast-evolving epitope travels far before exiting blunting neighborhoods in the conserved direction. Since individuals do not create new memory until the pathogen exits the blunting region, and protection decreases exponentially with distance between the memory and challenge strains, experienced individuals will be less protected to circulating strains than individuals with more recent primary infections. By contrast, in the any-epitope model, order dependence can have less impact over time, as there is a maximum distance the pathogen travels before leaving blunting neighborhoods. This means that two individuals' protection cannot diverge arbitrarily as the pathogen evolves, since new infections will trigger memory creation more frequently even in experienced individuals. This distance is maximized when all epitopes evolve at similar rates.

Immune Blind Spots. Epidemiological observations indicate that individuals do not always develop specific memory against a strain $\vec{s}\,'$ even after repeated exposures (60, 71). Those strains occur in "immune blind spots" (65), which we show here create blind spots in protection as well. In both the every-epitope model and the any-epitope model, immune blind spots arise when b is large compared to r: the blunting conditions guarantee the existence of points that are too antigenically similar to existing memory to trigger creation of new memory upon exposure but are nonetheless far enough to escape cross-protection from infection. Intuitively, this is because OAS prevents memory profiles from becoming too dense, and thus limits the antigenic "coverage" of protection.

To demonstrate the existence of immune blind spots, we will consider a scenario for each model. In the every-epitope model, any strain $\vec{s}$ in the memory profile defines an infinitely extended d-dimensional cross-shaped region where no additional memory can be created (Fig. 5*A*). For example, any strain $\vec{s}\,'$ sharing at least one epitope with $\vec{s}$ can never be added to the memory profile. Suppose strain $\vec{s}\,'$ has all but one epitope conserved with $\vec{s}$, and the remaining epitope is so far diverged as to escape all existing memory (i.e., the distance in that direction is much greater than

Fig. 4. Protection is order-dependent with blunting. (*A*) Infection history of hosts h_1 and h_2. (*B*) Memory profiles and protection distributions of hosts h_1 and h_2. Here, $r = \frac{b}{2}$.

Fig. 5. Immune blind spots arise when b is large compared to r. Given a memory at the origin, the "allowed strains" are examples of the closest possible strains that could be added to the memory profile. The protection is necessarily low at the immunity blind spot strains regardless of infection history in (*A*) the every-epitope model and (*B*) the any-epitope model as long as $\frac{b}{r}$ is sufficiently large (and, in (*A*), the distance from the memory strain to the immune blind spot is greater than both the blunting distance and the cross-protection distance). (*C* and *D*) show the probability of infection at the blind spot points considered here for (*C*) the every-epitope model and (*D*) the any-epitope model. Vertical dotted lines indicate approximate values of b/r above which infection is essentially guaranteed [decreasing with dimension in (*C*), dimension-independent in (*D*)].

both the cross-protection distance and the blunting distance). Since cross-protection to $\vec{s}'$ depends on the Euclidean distance from $\vec{s}$, the diverged epitope prevents $\vec{s}$ from protecting against $\vec{s}'$ despite shared epitopes. In fact, if the blunting distance is much greater than the cross-protection distance, an individual with $\vec{s}$ in their memory profile will never develop protection to $\vec{s}'$. This is because future exposures will create memory only to strains diverged from $\vec{s}$ by at least the blunting distance in every direction, and if $b \gg r$, such strains cannot protect against strains sharing epitopes with $\vec{s}$, including $\vec{s}'$ (Fig. 5*A*). When the blunting and cross-protection distances are comparable, a strain diverged from $\vec{s}$ (and therefore $\vec{s}'$) by exactly the blunting distance in the conserved directions can partially cross-protect against $\vec{s}'$. This means that the immune blind spot disappears if the blunting distance b is sufficiently close to the cross-protection distance r. An explicit calculation of the loss of protection at $\vec{s}'$ as a function of the ratio of blunting to cross-protection (*Materials and Methods*) shows that the immune blind spot appears at modest values of b/r, and the approximate minimum threshold for the appearance of immune blind spots shrinks as the number of epitopes (d) increases (Fig. 5*C*).

Similarly, in the any-epitope model, any strain $\vec{s}'$ that is diverged by less than the blunting length from $\vec{s}$ in at least one epitope cannot be added to the memory profile, and individuals will maintain some susceptibility to that strain as long as $b > r$. One example is a strain $\vec{s}'$ which is diverged from $\vec{s}$ by antigenic distance of half the blunting length ($b/2$) in every epitope ($s'_i = s_i + \frac{b}{2}$, for example) (Fig. 5*B*). The closest strains to $\vec{s}'$ that could be added to the memory profile are strains that are diverged from the reference strain $\vec{s}$ in at least one epitope by exactly the blunting distance ($+b$). The strain $\vec{s}'$ escapes cross-protection if $\frac{b}{r}$ is very large, but may be protected from these nearby strains if the cross-protection distance is comparable to the blunting distance (Fig. 5*D*). As long as $\frac{b}{r} \gtrsim 8$, infection upon exposure to strain $\vec{s}'$ is essentially guaranteed, and so $\vec{s}'$ falls in an immune blind spot (Fig. 5*D*) (*Materials and Methods*).

Population-Scale Memory Constraint. Both blunting models imply maximum total numbers of specific memories the population maintains in certain regions of antigenic space. This follows from the existence of blunting neighborhoods. Each host has at most one memory within the blunting neighborhood(s) $^{\mathrm{E,A}}\langle\langle\vec{p}\rangle\rangle_b^{(i)}$ of any reference point $\vec{p}$ in antigenic space. (The i index applies to the every-epitope model only.) It follows by taking the sum of memories at each point over all hosts that in any collection of N_h nonnaive individuals with collective memory distribution $h(\vec{s}, t)$, the total number of memories in each blunting neighborhood of $\vec{p}$ cannot exceed N_h. In symbols, for both blunting models, we have

$$\int_{\vec{s}' \in \,^{\mathrm{E,A}}\langle\langle\vec{p}\rangle\rangle_b^{(i)}} d\vec{s}' \, h(\vec{s}', t) \leq N_h. \quad \text{[13]}$$

The every-epitope model and the any-epitope model differ in the regions over which the integral is taken, with the every-epitope model integrating over d infinitely extended stripes that include $\vec{p}$, and the any-epitope model integrating over a single d-cube centered at $\vec{p}$ (Fig. 3).

We refer to Eq. **13** as the OAS constraint. It holds in any model of OAS with blunting.

Limits on Successive Strain Replacement Dynamics. We find that the conditions for serial strain replacement are harder to obtain in models that assume blunting (OAS). Only pathogens with a sufficiently low R_0 compared to the ratio of the cross-reactivity distance to the blunting distance (formally, $e^{b/r}$) can exhibit successive strain replacement dynamics in populations where individuals accumulate many immune memories (i.e., in the large-m limit).[*] This bound implies a maximum speed of antigenic evolution that depends on the blunting model.

---

[*]See *Materials and Methods*, Eq. **22**, for the finite-m correction, which quickly approaches the large-m behavior.

In the dynamics described by Marchi et al. (36), the population accumulates memories in a small interval close to the wavefront. This localized buildup of memory is necessary for maintaining the fitness gradient required to force the infection distribution to move at a constant speed while maintaining its shape. We have shown, however, that the OAS constraint prevents arbitrarily localized buildup of memory. Since the shape of blunting neighborhoods depends on the blunting model, the amount of memory close to the wavefront is also sensitive to specific assumptions about blunting. Specifically, OAS prevents memory accumulation in an interval whose length is determined by the projection of the direction the wave travels in the fastest (every-epitope model) or slowest (any-epitope model) direction (*Materials and Methods*). We refer to the angle between the fastest (slowest) direction in the every- (any-) epitope model as θ^*.

The OAS constraint requires that

$$\log R_0 \leq \frac{r \cos \theta^*}{b}, \qquad [14]$$

in the large-m limit (Fig. 6; derivation in *Materials and Methods*). Since $\cos \theta_i$ measures how quickly the ith epitope evolves compared to the other $(d-1)$ epitopes, this equation says that successive strain replacement is possible only if R_0 is not too large compared to the inverse of the blunting distance b in units of the cross-protection distance r, modulated by the projection onto the most important epitope for that model.

In the traveling wave model, the cross-protection distance, r; the population size, N_h; the diffusion constant, D; and the basic reproduction number, R_0 govern pathogen evolution, including the speed v at which the pathogen evolves, the SD σ of infections in the direction of evolution, and the time to the most recent common ancestor, t_{MRCA}, of cocirculating strains (36). The inequality in Eq. 14 implies that blunting leads to (N_h- and D-dependent) maximum v, minimum σ, and minimum t_{MRCA} in the context of successive strain replacement (Fig. 7). These limiting values are achieved for pathogens that saturate the inequality (*Materials and Methods*). For example, if the blunting length b is equal to twice the cross-protection length r and there are two epitopes evolving at the same rate ($\theta^* = \frac{\pi}{4}$ for both blunting models), then $r \cos \theta^*/b \approx 0.35$, and the maximum R_0 consistent with successive strain replacement dynamics is ~ 1.42. In this case, the minimum time required for the viral population (peak of the infection distribution) to travel the cross-reactivity distance varies up to ~ 100 times the duration of infection

(for, e.g., $r = 5$); the minimum SD of the viral population varies between ~ 0.1 and 0.5 cross-protection distances for values of r between 1 and 5; and the minimum time to the most recent common ancestor for two cocirculating strains varies between ~ 100 to 400 times the duration of infection for values of r between 1 and 5.

Discussion

Many factors have been hypothesized to lead to serial strain replacement, but theoretical tests of these hypotheses have tended to assume that individuals uniformly develop protective strain-specific immunity from infection. We have shown that OAS limits the density of population memory in antigenic space, a picture we expect to hold broadly across models of strain evolution. This bound on the number of memories constrains the successive strain replacement regime to low-transmissibility pathogens, with "more OAS"—a greater ability of memory to blunt the formation of new responses—corresponding to a tighter bound. Our results imply that observed patterns of serial replacement in nature may be less stable than thought.

This work represents a different approach compared to past studies of successive strain replacement. Previous work on waves in antigenic space (36) found that traveling waves solve the model for any set of parameters. Other investigations into the existence of traveling waves in pathogen evolution (24, 72–74) focused on phenomena such as disease treatment (73, 74) without including OAS. These models found traveling wave solutions for all $R_0 > 1$ with a minimum speed of evolution, as opposed to our result, which implies that OAS imposes a maximum speed of evolution. The previous studies that implemented OAS assumed successive strain replacement and examined its impact in age-structured models (64, 65). Our results are complementary, in that we identify where successive strain replacement is likely to arise in the first place.

By preventing arbitrary buildup of population memory in localized regions of antigenic space, OAS effectively limits the steepness of the fitness landscape that drives selection to escape immunity, and thus restricts the space of possible infection dynamics. Traveling waves are one example of a dynamical pattern that is subject to this restriction, as they require a dense buildup of memory near the wavefront. The necessary density depends on the characteristic cross-protection distance r and the basic reproduction number R_0, with smaller r and larger R_0 resulting in more memory near the wavefront. Since the blunting
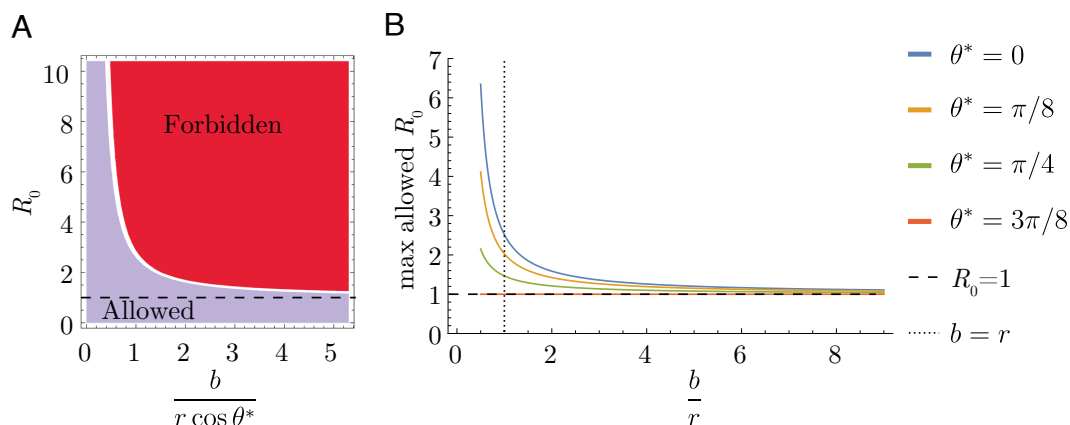
Fig. 6. Transmission limit implied by the OAS constraint. (*A*) Strain replacement imposes a maximum value on R_0. (*B*) Maximum value for R_0 as a function of $\frac{b}{r}$ for different values of θ^*.
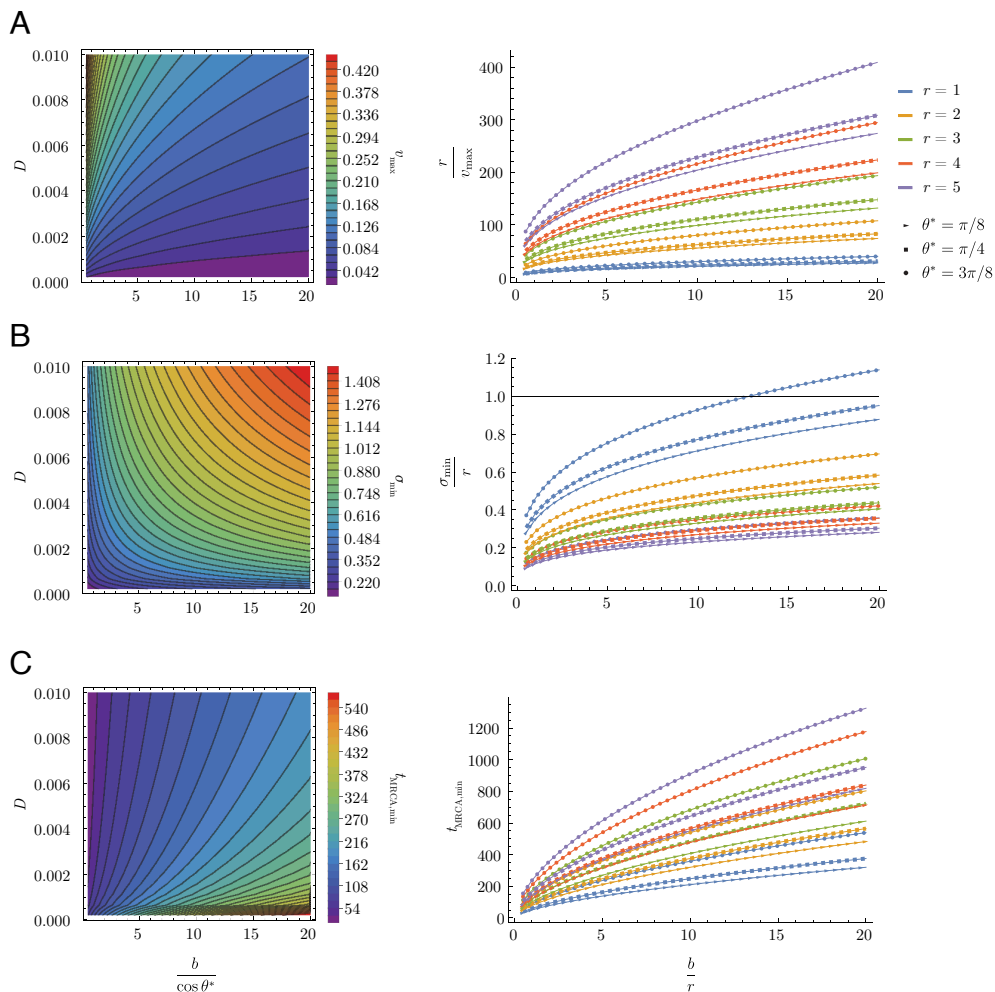
Fig. 7. Dynamical limits implied by the OAS constraint. (*A*) Maximum speed of evolution v (antigenic distance per infection), which depends on the diffusion constant D and the blunting distance modulated by the projection onto the fastest (slowest) evolving epitope for every (any) epitope model (*Left*), and the time for the infection peak to travel the cross-protection distance as a function of $\frac{b}{r}$ for different values of r and θ^* at saturating value of v_{max} (*Right*). (*B*) Minimum standing diversity σ (antigenic distance) as a function of the diffusion coefficient and blunting distance modulated by the projection onto the fastest (slowest) evolving epitope for the every (any) epitope model (*Left*), and the minimum standing diversity in units of cross-protection distance as a function of $\frac{b}{r}$ for different values of r and θ^*. Parameter combinations lying near or above the black solid line at $\sigma_{min}/r = 1$ cannot lead to serial strain replacement, as the traveling wave solution requires $\sigma/r \ll 1$. (*C*) Minimum time to the most recent common ancestor, t_{MRCA} (measured in duration of infection), as a function of the diffusion coefficient D and the blunting distance modulated by the projection onto the fastest (slowest) evolving epitope for the every (any) epitope model (*Left*), shown as a function of b/r for different values of r and θ^* (*Right*). Here, $N_h = 10^9$ (*Left* and *Right*) and $D = 0.004$ (*Right*).

distance b controls the number of memories that are "captured" in the blunting neighborhood at the wavefront, traveling wave dynamics are only consistent for small enough b in relation to r^{-1} and R_0. If the OAS bound is not satisfied, then the wave solution is unstable to the frequent generation of escape mutants. Depending on the diffusion coefficient, this could lead to the wave spreading or to continual generation and disappearance of multiple peaks, similar to the mound and comb patterns described in a similar model of virus-immune coevolution (75).

The implementation of the OAS bound in terms of blunting neighborhoods also demonstrates an intimate connection between OAS and the rules of memory creation in the presence of multiple evolving epitopes. The geometry of blunting neighborhoods runs parallel to the epitope-aligned axes. By contrast, as long as more than one epitope is evolving, the traveling wave does not travel parallel to any axis. Consequently, the number of memories captured in the blunting neighborhood near the wavefront depends on the relative speeds at which different epitopes evolve. Because the shape of blunting neighborhoods

depends on the interactions among epitopes, the slowest evolving epitope determines the upper bound on R_0 for the every-epitope model, and the fastest evolving epitope determines the bound for the any-epitope model.

The importance of R_0 and epitopes' rates of evolution suggests that effective public health measures, including nonpharmaceutical interventions and vaccines, might make pathogens more likely to exhibit successive strain replacement. An immunogenic, well-matched vaccine can deliver a high enough or "long enough" dose of antigen to overcome immune memory, forcing responses to new sites and the creation of strain-specific memory (e.g., refs. 48, 53, and 68). In contrast, at small doses of antigen, immune memory might effectively clear antigen unless every epitope is so antigenically diverged that new strain-specific memory can form. This suggests that population immunity dominated by immunogenic vaccines is likely better approximated by the any-epitope model. Since $\cos \theta^*$ is always greater in the any-epitope model than in the every-epitope model, the OAS bound on R_0 is weaker. In other words, we expect pathogens will be more likely

to exhibit successive strain replacement and less likely to diversify when vaccines can induce broader memory than infections.

One limitation of this work is that the population-level OAS constraint is imposed, rather than allowed to emerge from the individual-level model dynamics. Although we expect successive strain replacement models to generically induce memory buildup near the wavefront, we would need to incorporate additional mechanisms into the model for memory dynamics to recapitulate OAS endogenously. We still expect that regimes with a moving wavefront exist, but their functional form may be more complex or even time-dependent. Since all solutions to such a model would satisfy the OAS constraint by design, restrictions on the parameter regimes where successive strain replacement occurs would derive from existence conditions on wavelike solutions to the differential equation rather than consistency conditions on memory at the wavefront. Since in such a model memory could saturate, we hypothesize that a cross-protective, short-term immunity might be required to avoid repeated infections from accumulating, consistent with the conditions in refs. 12 and 27. With dynamics that enforce OAS, evolutionary dynamics could be explored in more regimes, e.g., to identify the conditions under which branching occurs. It would be interesting to compare such a study to refs. 6 and 34, which construct similar phase diagrams in models without OAS.

There is substantial heterogeneity in antibody-mediated protection against influenza and SARS-CoV-2 in the human population, with individuals showing diversity in their antibody landscapes (39, 66, 76, 77) and the particular epitopes they target, including their sensitivity to different escape mutations (71, 78–81). Some of this heterogeneity arises from differences in initial exposures (71), but there is striking diversity within birth cohorts (66, 77). Future versions of this model could include different subpopulations, distinguished by their memory repertoires, that are tracked over time and parameterized empirically.

Several assumptions should be reexamined in applying this theory to data, namely, the constant cross-protection distance r across epitopes, the simple memory creation and cross-protection functions, the single dimension of antigenic space for each epitope, and the artificial dichotomy of the any-epitope and every-epitope models. The correlates of protection for pathogens including influenza and SARS-CoV-2, necessary for understanding the cross-protection distance r, are not well understood and require additional investigation via prospective studies and challenge experiments. Other model assumptions, especially the blunting distance b and any- and every-epitope model forms, require a more quantitative understanding of the dynamics of immune memory. Longitudinal studies that combine diverse exposure histories with detailed immune profiling will help to accelerate understanding on these fronts.

Materials and Methods

Immune Blind Spots Calculations. We derive the existence of immune blind spots in the large-$\frac{b}{r}$ limit, and extend the argument to explain how Fig. 5 C and D were generated. For both the every- and any-epitope models, immune blind spots exist because the every- and any-epitope conditions prevent memories from being too closely spaced. The arguments thus proceed by determining, given a strain $\vec{s}$ in the memory profile, the maximum possible protection that could be achieved at specially chosen points $\vec{s}'$ in the case of a maximally "dense" memory profile near $\vec{s}'$.

In the every-epitope model, if strain $\vec{s}$ is in a host's memory profile $\mathcal{S}$, no strain $\vec{s}'$ that differs from $\vec{s}$ by a distance less than b in at least one dimension can be in the memory profile. Let $\vec{s}'$ be a strain with $s_1' = s_1, s_i' > s_i + b$ for

all $i = 2, \ldots, d$, and $(\sum_{i=2}^{d}(s_i' - s_i)^2)^{1/2} \gg r$, as shown in Fig. 5A in the case where $d = 2$ (so $s_2' > s_2 + b$ and $s_2' \gg s_2 + r$). Note that the host with memory (only) at $\vec{s}$ is essentially not protected at $\vec{s}'$, as its Euclidean distance from $\vec{s}$ is much greater than the cross-protection distance, and the strain $\vec{s}'$ cannot be added to the memory profile $\mathcal{S}$ even if the host is infected by $\vec{s}'$, since $\vec{s}$ and $\vec{s}'$ share the first epitope. The closest strains to $\vec{s}'$ that could be added to $\mathcal{S}$ are the two strains displaced from $\vec{s}'$ by b in the first epitope [with coordinates $(s_1 \pm b, s_2', \ldots, s_d')$], and the strains are mutually exclusive; that is, at most one of these strains can be added without violating the every-epitope condition in epitopes 2 through d. If the blunting distance b is sufficiently large compared to the cross-protection distance r, the maximum protection the individual could have at $\vec{s}'$ is well approximated by the protection conferred by one of these two closest strains, even if the memory profile $\mathcal{S}$ also contains strains further away. The minimum probability of infection upon exposure to $\vec{s}'$ is thus approximately $1 - e^{-b/r}$, which approaches 1 as b/r becomes large. The host thus has an immune blind spot at $\vec{s}'$ (Fig. 5A) in the large-b/r limit.

A similar observation holds for the any-epitope model. An individual with strain $\vec{s}$ in their memory profile $\mathcal{S}$ cannot have specific memory to any strain $\vec{s}'$ whose coordinates are all within a distance b of $\vec{s}$. Let $\vec{s}'$ be the strain with coordinates $s_i' = s_i + b/2$ for $i = 1, \ldots, d$. In addition to $\vec{s}$, there are $(2^d - 1)$ strains that differ from $\vec{s}'$ by a distance of $b/2$ in each epitope: these are strains that are located at either s_j or $s_j + b$ in each direction j (with at least one coordinate differing from $\vec{s}$). Each of these strains is distance $d^{1/2}b/2$ from $\vec{s}'$, and these are the closest strains to $\vec{s}'$ that could be in $\mathcal{S}$. As in the every-epitope case, for sufficiently large b/r, the maximum protection the individual could have at $\vec{s}'$ is well approximated by the protection conferred by these closest strains if memory had been acquired to all of them. The minimum probability that an individual with $\vec{s}$ in their memory profile is infected by $\vec{s}'$ upon exposure therefore equals the probability that neither $\vec{s}$ nor any of the other closest strains is protective, which is equal to $(1 - e^{-d^{1/2}b/2r})^{2^d}$. This is close to 1 when b/r is very large, meaning that the host has an immune blind spot at $\vec{s}'$ (Fig. 5B).

In reality, we don't expect the blunting distance to be many orders of magnitude greater than the cross-protection distance. If the blunting distance is not too much larger than the cross-protection distance, so that the large-b/r limit does not hold, the immune blind spots may cease to be blind spots, as cross-protection from allowed strains (outside the $\vec{s}$ excluded region) can be high enough to provide protection at $\vec{s}'$. In this case, to determine the existence of immune blind spots, it may not be enough to compute only the protection conferred from the closest allowed strains, as the contribution to protection at $\vec{s}'$ from further allowed strains may be nonnegligible. We therefore consider how much protection from further allowed strains can contribute at $\vec{s}'$ in both the every- and any-epitope models in order to calculate the maximum achievable protection at the points that become blind spots in the large-b/r limit considered above.

In the every-epitope model, the coordinates of the nearest allowed memories are all equal to $s_i' \pm b$; however, at most one such memory can be added. The next-nearest memories are distance $2b$ away in the first coordinate, and distance b in the remaining $(d - 1)$ coordinates. In general, for the kth-nearest memories, the first coordinate is distance $(kb + 1)$ away and the remaining $(d - 1)$ coordinates are distance kb away, and only one such point can be added for each k (without violating the every-epitope condition), no matter the number of dimensions. Each of these points is distance $((d-1)*(kb)^2 + (kb+1)^2)^{1/2}$ from $\vec{s}'$, and so the probability it protects against infection at $\vec{s}'$ equals $p_k = e^{-\frac{1}{r}(d\cdot(kb)^2 + (2kb) + 1)^{1/2}}$. If all of these strains were in the individual's memory profile, the total probability of infection at $\vec{s}'$ would be equal to the product $\Pi_k(1 - p_k)$, a quantity which is very close to 1 for $b/r \gtrsim 5$ ($d = 2$), 2 ($d = 3$), 1.8 ($d = 4$), etc. (Fig. 5C).

In the any-epitope model, blind spots occur in the center of a d-dimensional hypercube, and the densest memory possible profile containing a strain at the origin has strains at each point on a lattice with edges of distance b. The probability of protection given the densest possible distribution of memory is thus the product over lattice points of $(1 - e^{-\Delta_l/r})$, where Δ_l is the distance from the point $\vec{s}'$ to the lattice point l. To compute the probability of infection plotted in Fig. 5D, we explicitly computed the given product over all lattice points

up to a distance sufficiently far from the immune blind spot such that adding further lattice points does not appreciably impact the probability of protection. Note that the number of such lattice points at each distance depends on the number of epitopes d in a complicated way, and so we calculated separately for each d. The associated minimum probability of infection at the blind spot for an individual with memory at $\vec{s}$, computed by subtracting the probability of protection from 1, is shown in Fig. 5D.

Details of Transmission Bound Calculation. The general strategy for determining the transmission bound for either blunting model is to impose the OAS constraint over blunting neighborhoods of the appropriate geometry, find the neighborhood over which the integral of the memory distribution is maximized, carry out the integral, and determine the consistency conditions for the inequality to be satisfied.

In the every-epitope model, the ith blunting neighborhood about a point $\vec{p}$ takes the form ${}^E\langle\langle\vec{p}\rangle\rangle^i_b = \{\vec{s} : p_i - b/2 < s_i < p_i + b/2; -\infty < s_k < \infty \text{ for } k \neq i\}$. The memory distribution in the successive strain regime falls off exponentially into the past from the wavefront $x_1 = vt$. While blunting neighborhoods are infinitely extended rectangles parallel to the axes, the direction in which the wave travels is not generally parallel to the axes, so the amount of memory captured in the blunting neighborhoods depends on the direction the wave is traveling. The blunting neighborhood over which the integral of memory is maximized is parallel to the axis of the slowest-evolving epitope, with the wavefront at one of its corners (*SI Appendix*, Fig. S1 *A and B*). The memory captured therefore corresponds to the interval $vt - \frac{b}{\cos(\theta^*)} < x_1 < vt$ in wave-adapted coordinates. The every-epitope OAS constraint simplifies to

$$\int_{x'_1 = vt - \frac{b}{\cos(\theta^*)}}^{vt} dx'_1 \left(\frac{mN_h}{v\tau} e^{-(vt-x'_1)/v\tau} \right) \leq N_h. \quad [15]$$

In the any-epitope model, the blunting neighborhood about a point $\vec{p}$ takes the form ${}^A\langle\langle\vec{p}\rangle\rangle_b = \{\vec{s} : p_i - b/2 < s_i < p_i + b/2 \text{ for } i = 1, \ldots, d\}$. These blunting neighborhoods are d-dimensional cubes with side length b and sides parallel to the axes. As in the every-epitope model, the amount of memory captured in the blunting neighborhoods depends on the direction the wave is traveling. For the any-epitope geometry, the blunting neighborhood over which the integral of memory is maximized again has the wavefront at one of its corners. This time, however, the length of the interval captured in the neighborhood depends on the direction of the fastest-evolving epitope. The memory captured still corresponds to the interval $vt - \frac{b}{\cos(\theta^*)} < x_1 < vt$ in wave-adapted coordinates (*SI Appendix*, Fig. S1 *C and D*), but now θ^* indicates the fastest-evolving rather than the slowest-evolving direction. The any-epitope OAS constraint simplifies to

$$\int_{x'_1 = vt - \frac{b}{\cos(\theta^*)}}^{vt} dx'_1 \left(\frac{mN_h}{v\tau} e^{-(vt-x'_1)/v\tau} \right) \leq N_h. \quad [16]$$

Both of these integrals are of the form $\int_{x'_1 = vt - \ell}^{vt} dx'_1 \left(\frac{mN_h}{v\tau} e^{-(vt-x'_1)/v\tau} \right)$ with different values of ℓ, so we evaluate this integral and then substitute for ℓ to obtain the relevant inequalities for each blunting model. The inequality simplifies to the condition

$$m(1 - e^{-\ell/v\tau}) \leq 1. \quad [17]$$

This is equivalent to the condition that

$$\frac{\ell}{v\tau} \leq -\log(1 - \frac{1}{m}). \quad [18]$$

The derived quantity $v\tau$ can be expressed in terms of the fundamental parameters that characterize the system by manipulating several expressions derived in ref. 36. The timescale τ equals mN_h/N_l. The incremental fitness gradient along the direction of motion of the wave is analogous to the fitness

effect of mutations, s, and equals $\frac{m}{r}\left(R_0^{1/m} - 1\right)$ in the traveling wave system. Marchi et al. argue that s is related to v via the relation $s = \frac{N_l}{N_h v}$. It follows that

$$\frac{1}{v\tau} = \frac{1}{v}\frac{N_l}{mN_h} = \frac{s}{m} = \frac{1}{r}\left(R_0^{1/m} - 1\right), \quad [19]$$

so that the inequality in Eq. **18** becomes

$$\frac{\ell}{r}\left(R_0^{1/m} - 1\right) \leq -\log(1 - \frac{1}{m}). \quad [20]$$

The limit as m becomes large of $m\left(R_0^{1/m} - 1\right)$ is equal to $\log R_0$, and more generally the following inequality holds for any value of m greater than 1,

$$\log R_0 < m\left(R_0^{1/m} - 1\right). \quad [21]$$

Together, Eqs. **20** and **21** imply that the following inequality among the fundamental parameters that characterize transmission and the immune response must hold,

$$\log R_0 < \frac{r}{\ell} \cdot m \log\left(\frac{m}{m-1}\right). \quad [22]$$

In the large-m limit, Eq. **22** simplifies to

$$\log R_0 < \frac{r}{\ell}. \quad [23]$$

This is the same as the expressions given in the Results, where ℓ equals $\frac{b}{\cos\theta^*}$.

Bounds on Dynamical Quantities in Traveling Wave Model. In this section, we give the expressions for dynamical quantities in the antigenic space model, originally derived in ref. 36, and numerically demonstrate the monotonicity of these quantities in the ratio $\frac{b}{\cos\theta^*}$, implying D-dependent extrema when R_0 equals the limiting value $\frac{r\cos\theta^*}{b}$, as in Eq. **14**.

The quantities v, σ and t_{MRCA} are simply expressed in terms of the fitness increment $s = \frac{m}{r}\left(R_0^{1/m} - 1\right)$:

Wave velocity: $v = \frac{N_l}{N_h s}$, $\quad [24]$

$$= D^{2/3} s^{1/3}\left(24 \log\left(N_l(Ds^2)^{1/3}\right)\right)^{1/3}. \quad [25]$$

Wave width: $\sigma = \sqrt{\frac{v}{s}}. \quad [26]$

Average time to the most recent common ancestor:

$$t_{MRCA} \approx (1.66) \cdot \frac{\sigma^2}{2D}. \quad [27]$$

Note that the wave velocity is defined implicitly through Eqs. **24** and **25**, which determine both v and the number of infected hosts N_l in terms of the independent parameters of the system (m, R_0, r, D, and N_h).

The quantities v (*SI Appendix*, Fig. S2A), σ (*SI Appendix*, Fig. S2B), and t_{MRCA} (*SI Appendix*, Fig. S2C) are monotonic in the fitness gradient s. The bound on R_0 for serial strain replacement, Eq. **14**, is equivalent to the statement that $s \leq \frac{\cos\theta^*}{b}$ in the large-m limit, since $s = \frac{m}{r}\left(R_0^{1/m} - 1\right)$ approaches $\frac{\log R_0}{r}$ and $\log R_0 \leq \frac{r\cos\theta^*}{b}$. It follows that a monotonically increasing (decreasing) function of s is maximized (minimized) over all traveling wave solutions at the limiting value $s = \frac{\cos\theta^*}{b}$. Thus for any values of D, b, θ^*, and N_h, there exists a maximum limit on v and minimum limits on σ and t_{MRCA} which are independent of m in the large-m limit. These bounds are plotted in the *Left* panels of Fig. 7.

1. S. Gupta et al., The maintenance of strain structure in populations of recombining infectious agents. *Nat. Med.* **2**, 437–442 (1996).
2. S. Gupta, N. Ferguson, R. Anderson, Chaos, persistence, and evolution of strain structure in antigenically diverse infectious agents. *Science* **280**, 912–915 (1998).
3. J. R. Gog, B. T. Grenfell, Dynamics and selection of many-strain pathogens. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 17209–17214 (2002).
4. M. Lipsitch, J. J. O'Hagan, Patterns of antigenic diversity and the mechanisms that maintain them. *J. R. Soc. Interface* **4**, 787–802 (2007).
5. B. T. Grenfell et al., Unifying the epidemiological and evolutionary dynamics of pathogens. *Science* **303**, 327–332 (2004).
6. D. Zinder, T. Bedford, S. Gupta, M. Pascual, The roles of competition and mutation in shaping antigenic and genetic diversity in influenza. *PLoS Pathog.* **9**, e1003104 (2013).
7. S. P. Yazdankhah et al., Distribution of serogroups and genotypes among disease-associated and carried isolates of *Neisseria meningitidis* from the Czech Republic, Greece, and Norway. *J. Clin. Microbiol.* **42**, 5146–5153 (2004).
8. J. Lucidarme et al., Genomic resolution of an aggressive, widespread, diverse and expanding meningococcal serogroup B, C and W lineage. *J. Infect.* **71**, 544–552 (2015).
9. N. Topaz et al., Phylogenetic relationships and regional spread of meningococcal strains in the meningitis belt, 2011–2016. *EBioMedicine* **41**, 488–496 (2019).
10. D. A. Caugant, O. B. Brynildsrud, *Neisseria meningitidis*: Using genomics to understand diversity, evolution and pathogenesis. *Nat. Rev. Microbiol.* **18**, 84–96 (2020).
11. C. D. Rao, P. Yergolkar, K. S. Shankarappa, Antigenic diversity of enteroviruses associated with nonpolio acute flaccid paralysis, India, 2007–2009. *Emerg. Infect. Dis.* **18**, 1833–1840 (2012).
12. N. M. Ferguson, A. P. Galvani, R. M. Bush, Ecological and immunological determinants of influenza evolution. *Nature* **422**, 428–433 (2003).
13. D. J. Smith et al., Mapping the antigenic and genetic evolution of influenza virus. *Science* **305**, 371–376 (2004).
14. R. T. Eguia et al., A human coronavirus evolves antigenically to escape antibody immunity. *PLoS Pathog.* **17**, e1009453 (2021).
15. K. A. Lythgoe et al., Lineage replacement and evolution captured by 3 years of the United Kingdom Coronavirus (COVID-19) Infection Survey. *Proc. R. Soc. B.* **290**, 20231284 (2023).
16. R. A. Neher, Contributions of adaptation and purifying selection to SARS-CoV-2 evolution. *Virus Evol.* **8**, veac113 (2022).
17. P. V. Markov et al., The evolution of SARS-CoV-2. *Nat. Rev. Microbiol.* **21**, 361–379 (2023).
18. K. E. Kistler, T. Bedford, An atlas of continuous adaptive evolution in endemic human viruses. *Cell Host Microbe* **31**, 1898–1909.e3 (2023).
19. K. Koelle, D. A. Rasmussen, The effects of a deleterious mutation load on patterns of influenza A/H3N2's antigenic evolution in humans. *eLife* **4**, e07361 (2015).
20. C. Castillo-Chavez, H. W. Hethcote, V. Andreasen, S. A. Levin, W. M. Liu, Epidemiological models with age structure, proportionate mixing, and cross-immunity. *J. Math. Biol.* **27**, 233–258 (1989).
21. V. Andreasen, J. Lin, S. A. Levin, The dynamics of cocirculating influenza strains conferring partial cross-immunity. *J. Math. Biol.* **35**, 825–842 (1997).
22. J. Lin, V. Andreasen, S. A. Levin, Dynamics of influenza a drift: The linear three-strain model. *Math. Biosci.* **162**, 33–51 (1999).
23. J. H. P. Dawes, J. R. Gog, The onset of oscillatory dynamics in models of multiple disease strains. *J. Math. Biol.* **45**, 471–510 (2002).
24. J. Lin, V. Andreasen, R. Casagrandi, S. A. Levin, Traveling waves in a model of influenza A drift. *J. Theor. Biol.* **222**, 437–445 (2003).
25. L. J. Abu-Raddad, N. M. Ferguson, The impact of cross-immunity, mutation and stochastic extinction on pathogen diversity. *Proc. Biol. Sci.* **271**, 2431–2438 (2004).
26. L. J. Abu-Raddad, N. M. Ferguson, Characterizing the symmetric equilibrium of multi-strain host-pathogen systems in the presence of cross immunity. *J. Math. Biol.* **50**, 531–558 (2005).
27. F. Tria, M. Lässig, L. Peliti, S. Franz, A minimal stochastic model for influenza evolution. *J. Stat. Mech.* **2005**, P07008 (2005).
28. M. Recker, O. G. Pybus, S. Nee, S. Gupta, The generation of influenza outbreaks by a network of host immune responses against a limited set of antigenic types. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7711–7716 (2007).
29. S. Cobey, M. Pascual, Consequences of host heterogeneity, epitope immunodominance, and immune breadth for strain competition. *J. Theor. Biol.* **270**, 80–87 (2011).
30. A. Kucharski, J. R. Gog, Influenza emergence in the face of evolutionary constraints. *Proc. Biol. Sci.* **279**, 645–652 (2012).
31. T. Bedford, A. Rambaut, M. Pascual, Canalization of the evolutionary trajectory of the human influenza virus. *BMC Biol.* **10**, 38 (2012).
32. K. Koelle, S. Cobey, B. Grenfell, M. Pascual, Epochal evolution shapes the phylodynamics of interpandemic influenza A (H3N2) in humans. *Science* **314**, 1898–1903 (2006).
33. V. Andreasen, S. Levin, J. Lin, A model of influenza A drift evolution. *ZAMM Z. Angew. Math. Mech.* **76**, 421–424 (1996).
34. L. Yan, R. A. Neher, B. I. Shraiman, Phylodynamic theory of persistence, extinction and speciation of rapidly adapting pathogens. *eLife* **8**, e44205 (2019).
35. H. Schenk, H. Schulenburg, A. Traulsen, How long do red queen dynamics survive under genetic drift? A comparative analysis of evolutionary and eco-evolutionary models *BMC Evol. Biol.* **20**, 8 (2020).
36. J. Marchi, M. Lässig, A. M. Walczak, T. Mora, Antigenic waves of virus-immune coevolution. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2103398118 (2021).
37. T. Bedford et al., Global circulation patterns of seasonal influenza viruses vary with antigenic drift. *Nature* **523**, 217–220 (2015).
38. F. M. Davenport, A. V. Hennessy, T. Francis Jr., Epidemiologic and immunologic significance of age distribution of antibody to antigenic variants of influenza virus. *J. Exp. Med.* **98**, 641–656 (1953).
39. F. M. Davenport, A. V. Hennessy, C. H. Stuart-Harris, T. Francis Jr., Epidemiology of influenza; Comparative serological observations in England and the United States. *Lancet* **269**, 469–474 (1955).
40. A. V. Hennessy, F. M. Davenport, T. Francis Jr., Studies of antibodies to strains of influenza virus in persons of different ages in sera collected in a postepidemic period. *J. Immunol.* **75**, 401–409 (1955).
41. F. M. Davenport, A. V. Hennessy, A serologic recapitulation of past experiences with influenza A; antibody response to monovalent vaccine. *J. Exp. Med.* **104**, 85–97 (1956).
42. F. M. Davenport, A. V. Hennessy, Predetermination by infection and by vaccination of antibody response to influenza virus vaccines. *J. Exp. Med.* **106**, 835–850 (1957).
43. T. Francis, On the doctrine of original antigenic sin. *Proc. Am. Philos. Soc.* **104**, 572–578 (1960).
44. F. M. Davenport, A. V. Hennessy, J. Drescher, J. Mulder, T. Francis Jr., Further observations on the relevance of serologic recapitulations of human infection with influenza viruses. *J. Exp. Med.* **120**, 1087–1097 (1964).
45. F. de St Groth, R. G. Webster, Disquisitions of original antigenic sin. I. evidence in man. *J. Exp. Med.* **124**, 331–345 (1966).
46. F. de St Groth, R. G. Webster, Disquisitions on original antigenic sin. II. proof in lower creatures. *J. Exp. Med.* **124**, 347–361 (1966).
47. S. L. Linderman, S. E. Hensley, Antibodies with 'original antigenic sin' properties are valuable components of secondary immune responses to influenza viruses. *PLoS Pathog.* **12**, e1005806 (2016).
48. A. Schiepers et al., Molecular fate-mapping of serum antibody responses to repeat immunization. *Nature* **615**, 482–489 (2023).
49. T. Weber et al., Enhanced SARS-CoV-2 humoral immunity following breakthrough infection builds upon the preexisting memory B cell pool. *Sci. Immunol.* **8**, eadk5845 (2023).
50. J. S. Turner et al., Human germinal centres engage memory and naive B cells after influenza vaccination. *Nature* **586**, 127–132 (2020).
51. S. Cobey et al., Poor immunogenicity, not vaccine strain egg adaptation, may explain the low H3N2 influenza vaccine effectiveness in 2012–2013. *Clin. Infect. Dis.* **67**, 327–333 (2018).
52. D. J. Smith, S. Forrest, D. H. Ackley, A. S. Perelson, Variable efficacy of repeated annual influenza vaccination. *Proc. Natl. Acad. Sci. U.S.A.* **96**, 14001–14006 (1999).
53. V. I. Zarnitsyna, J. Lavine, A. Ellebedy, R. Ahmed, R. Antia, Multi-epitope models explain how pre-existing antibodies affect the generation of broadly protective responses to influenza. *PLOS Pathog.* **12**, e1005692 (2016).
54. A. M. Davey, S. K. Pierce, Intrinsic differences in the initiation of B cell receptor signaling favor responses of human IgG(+) memory B cells over IgM(+) naive B cells. *J. Immunol.* **188**, 3332–3341 (2012).
55. S. G. Tangye, D. T. Avery, E. K. Deenick, P. D. Hodgkin, Intrinsic differences in the proliferation of naive and memory B cells as a mechanism for enhanced secondary immune responses. *J. Immunol.* **170**, 686–694 (2003).
56. K. L. Good, D. T. Avery, S. G. Tangye, Resting human memory B cells are intrinsically programmed for enhanced survival and responsiveness to diverse stimuli compared to naive B cells. *J. Immunol.* **182**, 890–901 (2009).
57. S. Cobey, S. E. Hensley, Immune history and influenza virus susceptibility. *Curr. Opin. Virol.* **22**, 105–111 (2017).
58. G. E. Rydell, K. Prakash, H. Norder, M. Lindh, Hepatitis B surface antigen on subviral particles reduces the neutralizing effect of anti-HBs antibodies on hepatitis B viral particles in vitro. *Virology* **509**, 67–70 (2017).
59. J. Jin, C. Park, S. H. Cho, J. Chung, The level of decoy epitope in PCV2 vaccine affects the neutralizing activity of sera in the immunized animals. *Biochem. Biophys. Res. Commun.* **496**, 846–851 (2018).
60. S. Gouma et al., Middle-aged individuals may be in a perpetual state of H3N2 influenza virus susceptibility. *Nat. Commun.* **11**, 4566 (2020).
61. C. P. Arevalo et al., Original antigenic sin priming of influenza virus hemagglutinin stalk antibodies. *Proc. Natl. Acade. Sci. U.S.A.* **117**, 17221–17227 (2020).
62. J. W. Yewdell, J. J. Santos, Original antigenic sin: How original? How sinful? *Cold Spring Harb. Perspect. Med.* **11**, a038786 (2020).
63. J. H. Kim, I. Skountzou, R. Compans, J. Jacob, Original antigenic sin responses to influenza viruses. *J. Immunol.* **183**, 3294–3301 (2009).
64. A. J. Kucharski, J. R. Gog, The role of social contacts and original antigenic sin in shaping the age pattern of immunity to seasonal influenza. *PLoS Comput. Biol.* **8**, e1002741 (2012).
65. A. J. Kucharski, J. R. Gog, Age profile of immunity to influenza: Effect of original antigenic sin. *Theor. Popul. Biol.* **81**, 102–112 (2012).

66. B. Yang *et al.*, Long term intrinsic cycling in human life course antibody responses to influenza A(H3N2): An observational and modeling study. *eLife* **11**, e81457 (2022).

67. A. Yisimayi *et al.*, Repeated omicron exposures override ancestral SARS-CoV-2 immune imprinting. *Nature* **625**, 148–156 (2023).

68. J. H. Kim, W. G. Davis, S. Sambhara, J. Jacob, Strategies to alleviate original antigenic sin responses to influenza viruses. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 13751–13756 (2012).

69. S. Kryazhimskiy, U. Dieckmann, S. A. Levin, J. Dushoff, On state-space reduction in multi-strain pathogen models, with an application to antigenic drift in influenza A. *PLoS Comput. Biol.* **3**, e159 (2007).

70. K. M. Gostic, M. Ambrose, M. Worobey, J. O. Lloyd-Smith, Potent protection against H5N1 and H7N9 influenza via childhood hemagglutinin imprinting. *Science* **354**, 722–726 (2016).

71. S. L. Linderman *et al.*, Potential antigenic explanation for atypical H1N1 infections among middle-aged adults during the 2013–2014 influenza season. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 15798–15803 (2014).

72. T. Zhang, W. Wang, Existence of traveling wave solutions for influenza model with treatment. *J. Math. Anal. Appl.* **419**, 469–495 (2014).

73. D. Deng, J. Li, D. Zhang, Existence of traveling waves for a nonlocal dispersal SIR epidemic model with treatment. *J. Math. Anal. Appl.* **499**, 125009 (2021).

74. J. Wei, J. Zhou, Z. Zhen, L. Tian, Time periodic traveling waves in a three-component non-autonomous and reaction-diffusion epidemic model. *Int. J. Math.* **32**, 2150003 (2021).

75. D. A. Kessler, H. Levine, Multiple possible patterns can emerge from virus-immune coevolution. arXiv [Preprint] (2023). https://doi.org/10.48550/arXiv.2311.00417 (Accessed 16 April 2024).

76. J. M. Fonville *et al.*, Antibody landscapes after influenza virus infection or vaccination. *Science* **346**, 996–1000 (2014).

77. K. Kim *et al.*, Measures of population immunity can predict the dominant clade of influenza a (H3N2) and reveal age-associated differences in susceptibility and specificity. medRxiv [Preprint] (2023). https://doi.org/10.1101/2023.10.26.23297569 (Accessed 1 January 2024).

78. J. M. Lee *et al.*, Mapping person-to-person variation in viral mutations that escape polyclonal serum targeting influenza hemagglutinin. *eLife* **8**, e49324 (2019).

79. F. C. Welsh *et al.*, Age-dependent heterogeneity in the antigenic effects of mutations to influenza hemagglutinin. bioRxiv [Preprint] (2023). https://doi.org/10.1101/2023.12.12.571235 (Accessed 1 January 2024).

80. A. J. Greaney *et al.*, Comprehensive mapping of mutations in the SARS-CoV-2 receptor-binding domain that affect recognition by polyclonal human plasma antibodies. *Cell Host Microbe* **29**, 463–476.e6 (2021).

81. A. J. Greaney *et al.*, Complete mapping of mutations to the SARS-CoV-2 spike receptor-binding domain that escape antibody recognition. *Cell Host Microbe* **29**, 44–57.e9 (2021).

82. L. McGough, OAS-strain-replacement. Github. https://github.com/cobeylab/OAS-strain-replacement/tree/main. Deposited 24 April 2024.