

Robustness and Fragility in Immunosenescence

Sean P. Stromberg*, Jean Carlson*

Physics Department, University of California Santa Barbara, Santa Barbara, California, United States of America

We construct a model to study tradeoffs associated with aging in the adaptive immune system, focusing on cumulative effects of replacing naive cells with memory cells. Binding affinities are characterized by a stochastic shape space model. System loss arising from an individual infection is associated with disease severity, as measured by the total antigen population over the course of an infection. We monitor evolution of cell populations on the shape space over a string of infections, and find that the distribution of losses becomes increasingly heavy-tailed with time. Initially this lowers the average loss: the memory cell population becomes tuned to the history of past exposures, reducing the loss of the system when subjected to a second, similar infection. This is accompanied by a corresponding increase in vulnerability to novel infections, which ultimately causes the expected loss to increase due to overspecialization, leading to increasing fragility with age (i.e., immunosenescence). In our model, immunosenescence is not the result of a performance degradation of some specific lymphocyte, but rather a natural consequence of the built-in mechanisms for system adaptation. This “robust, yet fragile” behavior is a key signature of Highly Optimized Tolerance.

Citation: Stromberg SP, Carlson J (2006) Robustness and fragility in immunosenescence. PLoS Comput Biol 2(11): e160. doi:10.1371/journal.pcbi.0020160

Introduction

The adaptive immune system [1] of vertebrates has evolved in a manner that enables adaptation to the history of infections over the lifetime of each individual organism. It consists of a complex, heterogeneous collection of cells that is derived from stem cells in the bone marrow and proliferates in the lymph nodes. These cells are endowed with the remarkable ability to discriminate between self and nonself agents within the body and to remove the nonself elements [2–4]. B and T cells are the white blood cells (i.e., lymphocytes) that constitute the adaptive components of the immune system. They derive their ability to discriminate self from nonself with the binding specificity of their receptors: T cell receptors for T cells, and membrane-bound antibody for B cells. These receptors are assembled randomly from gene segments, producing a population of naive cells, in which each individual combination has a different binding specificity. The random combinations of genes give the immune system the ability to produce diverse cells capable of responding to many pathogens. During an infection, the cells whose receptors recognize the antigen proliferate and differentiate into antigen-removing effector cells and long-lived memory cells. The memory cells give rise to a more rapid and efficient response to a secondary exposure to the same antigen. However, due to homeostatic regulation of the lymphocyte population, the growth of memory cells reduces the naive cell population size. Over time, this has the effect of increasing sensitivity to novel infections.

We introduce a model that captures this tradeoff between resilience to repeated exposure and sensitivity to new pathogens. The model consists of coupled differential equations for immune-system cell populations, defined in terms of their primary immunological function and their binding characteristics. The relative population sizes evolve in time, stimulated by episodic infections. Antigens are drawn from a probability distribution of their characteristics, which enables estimation of the binding affinity of lymphocytes. We include a constraint on the total number of immune cells in the system, and define an immunological loss function that quantifies disease severity. The constraint on the number of

cells implies that memory, which is specific to infections the system has seen, comes at a price for unseen infections [5]. Our model illustrates how the immune system initially increases in effectiveness but eventually becomes overspecialized with age.

Results

Characteristics of the distinct populations in our model are summarized in Table 1. Note that in our simplified model, as in Segel and Perelson [4], lymphocytes (memory, naïve, and effector cells) are not specifically T or B cells, but a generalization having properties common to both types. We have also omitted helper T cells (that help to stimulate the immune response), as well as the complex germinal center reaction and somatic hypermutation (processes involving the proliferation and development of lymphocytes), assuming these features are not limiting factors in immunosenescence.

In our model, A , E , N , and M are all fields on a generalized shape space, introduced by Oster and Perelson [6] to represent the lock-and-key type specificity of antigen-receptor binding. The dimensions of the generalized Euclidean shape space correspond to quantities such as size, charge distribution, and hydrophobicity. This differs from other Hamming-type shape space models where each dimension pertains to a particular amino acid in the binding-region sequence [7]. The binding sites on the antigens and receptor proteins are described by the position

Editor: Michael Deem, Rice University, United States of America

Received: June 16, 2006; **Accepted:** October 10, 2006; **Published:** November 24, 2006

A previous version of this article appeared as an Early Online Release on October 10, 2006 (doi:10.1371/journal.pcbi.0020160.eor).

Copyright: © 2006 Stromberg and Carlson. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abbreviations: HOT, Highly Optimized Tolerance

*To whom correspondence should be addressed. E-mail: stromberg@physics.ucsb.edu (SPS), and carlson@physics.ucsb.edu (JC)

Synopsis

The immune system can be viewed as a complex system, which adapts, over time, to reflect the history of infections experienced by the organism. This paper describes a model that captures this adaptation and corresponding robust, yet fragile behavior. To model immunological processes that rely on binding specificity, researchers typically utilize abstract shape space models. These models describe the binding characteristics of a receptor or antigen as points in a high dimensional vector space. Stromberg and Carlson have incorporated the concept of shape space into a dynamical model of immune response. They use this model to examine the development of the system over a series of infections and monitor the severity of disease for each infection. The diseases are drawn at random from a distribution having a few frequently reoccurring and many rare. The system is observed to adapt over a series of infections, becoming robust to the frequent diseases while developing fragility to the rare diseases. This age-correlated weakness arises from the underlying dynamics of system adaptation rather than from an accumulation of defects. This robust, yet fragile behavior is a signature of Highly Optimized Tolerance, a mechanism for complexity based on robustness tradeoffs.

vectors in the shape space, \vec{x} and \vec{y} , respectively. The binding of antigen to dendritic cells as described in Table 1, is not shape space-dependent. All antigens in this model bind to dendritic cells with the same affinity. Recent calculations indicate that the shape space is best described with between five and eight dimensions [7]. We use two dimensions here for visualization. Using higher dimensions in the model changes the distribution of affinities, but does not dramatically affect the results of the paper. Extensions to higher dimensions, as well as to the more complex interactions listed above, will be considered in future work.

Vector values of the antigen \vec{x} and immune cells \vec{y} describe complementary binding characteristics, so that the binding affinity for $A(\vec{x})$ and a receptor at \vec{y} , given by $\gamma(\vec{x}, \vec{y})$, is maximal for $\vec{x} = \vec{y}$. For $\vec{x} \neq \vec{y}$, the binding affinity is a decaying function of the distance from \vec{x} to \vec{y} in the shape space $\gamma = \gamma(|\vec{x} - \vec{y}|)$. Following Segel and Perelson [4], we take the affinity function to be a Gaussian:

$$\gamma(\vec{x}, \vec{y}) = \gamma_{max} e^{-|\vec{x} - \vec{y}|^2 / (2b^2)}, \quad (1)$$

where γ_{max} sets the overall scale for the strength of the

immune response, and b sets the mismatch tolerance between antigens and receptors. Replacing γ with different decaying functions of distance (e.g., exponential) does not significantly alter the results of this work.

Using the species listed in Table 1, we next define our model of immune system response and adaptation. The process can be broken down into three stages: (i) antigen proliferation and immune response in an individual infection, (ii) recovery and stasis between infections, (iii) long-term adaptation of lymphocyte populations over the lifetime of the individual. We assume that rates of infection are small enough that the immune system completely eliminates one pathogen, and relaxes to the uninfected state long before the next infection occurs. This allows us to introduce our model in three stages, corresponding to the increasing time scales (i)–(iii) above, beginning with an individual infection.

Periods of infection are associated with introduction of antigen. Different diseases are associated with different shape space coordinates \vec{x} , and have different rates of infection. Upon infection, a pathogen proliferates at an exponential rate so that the antigen population grows at a rate $\beta A(\vec{x})$ (Equation 2). In our model, $t = 0$ marks the time when the pathogen is mixed into the lymph and begins to stimulate an immune response. We assume a finite value of $A(\vec{x}, 0)$ at this onset to account for the delay in the start of the immune response. This represents how once a small amount of the pathogen bypasses the physical barriers of the innate immune system it will proliferate until it finds its way into the blood and then lymph nodes, at which point the immune response is triggered.

Next the unoccupied dendritic cells, F , begin to trap antigen and become activated to F^* at a rate $\rho F A(\vec{x})$ (Equations 2 and 3). The activated dendritic cells F^* now present antigen to the naive and memory cells, stimulating them to divide (Equations 3–5). Overall, the stimulation occurs at a rate $\alpha \gamma(\vec{x}, \vec{y}) F^* (N + M)(\vec{y})$. Here the factor of $\gamma(\vec{x}, \vec{y})$ gives the highest-affinity lymphocytes the most rapid stimulation. The daughters of the cellular division of either $N(\vec{y})$ or $M(\vec{y})$, are $E(\vec{y})$ cells with fraction f , or $M(\vec{y})$ cells with fraction $(1 - f)$ (Equations 4–6). Several generations of memory cells may therefore be produced through an immune response. In stimulating naive and memory cells, F^* reverts back to F (Equation 3), keeping the total $F + F^* = H$ constant. This rate is the integral of the rate of N and M stimulation

Table 1. Immune System Model Ingredients

Species	Symbol	Primary Function
Antigen	$A(\vec{x}, t)$	Chemical that stimulates an immune response, which acts to remove it from the body.
Effector	$E(\vec{y}, t)$	Short-lived cells that remove antigen from the body.
Naive	$N(\vec{y}, t)$	Short-lived detector cells generated in the bone marrow with randomly assembled receptors. When stimulated, these divide into memory and effector cells.
Memory	$M(\vec{y}, t)$	Long-lived detector cells having exactly the same receptor as the parent. Like naive cells, when stimulated these also divide into more memory and effector cells.
Dendritic	$F(t)$ and $F^*(t)$	Antigen-presenting cells that act as catalysts. F traps antigen, converting to F^* . F^* facilitates the stimulation of naive and memory cells, which convert it back to F .

Population sizes are functions of time t , and evolve as defined below. Aside from the dendritic cells, all populations have binding characteristics represented as vectors in an abstract shape space. Increasing proximity between an antigen and an effector corresponds to increasing efficacy of the immune response.
doi:10.1371/journal.pcbi.0020160.t001

over the entire shape space. Effector cells eliminate antigen from the system with rate $A(\vec{x})\gamma(\vec{x}, \vec{y})E(\vec{y})$. The total rate of antigen removal is the integral of this rate over the shape space of effector cells (Equation 2). Effector cells are short-lived and die with rate $\delta E(\vec{y})$ (Equation 6).

These short time scale reactions are described by the following system of equations (we drop the explicit t dependence in all populations to simplify notation):

$$\frac{\partial A(\vec{x})}{\partial t} = \beta A(\vec{x}) - A(\vec{x}) \int \gamma(\vec{x}, \vec{y}) E(\vec{y}) d\vec{y} - \rho F A(\vec{x}) \quad (2)$$

$$\frac{\partial F^*}{\partial t} = \rho F A(\vec{x}) - \alpha \int \gamma(\vec{x}, \vec{y}) F^* (M + N)(\vec{y}) d\vec{y} \quad (3)$$

$$\frac{\partial N(\vec{y})}{\partial t} = -\alpha \gamma(\vec{x}, \vec{y}) F^* N(\vec{y}) \quad (4)$$

$$\frac{\partial M(\vec{y})}{\partial t} = (2 - 2f)\alpha \gamma(\vec{x}, \vec{y}) F^* N(\vec{y}) + (1 - 2f)\alpha \gamma(\vec{x}, \vec{y}) F^* M(\vec{y}) \quad (5)$$

$$\frac{\partial E(\vec{y})}{\partial t} = 2f\alpha \gamma(\vec{x}, \vec{y}) F^* (M + N)(\vec{y}) - \delta E(\vec{y}) \quad (6)$$

Here α is an affinity-independent factor that accounts for the difference in $\gamma(\vec{x}, \vec{y})$ -dependent rates of lymphocyte stimulation and removal of antigen. Note that our model does not include any spatial variables for position of antigen and lymphocytes in the body, which corresponds to assuming a well-mixed system. This system of equations exhibits many features we expect from an immune system model, such as rapid secondary response and affinity selection.

During the immune response, the naive and memory cells are indistinguishable. In our model their difference becomes apparent on intermediate time scales. Therefore, we consider their combined effect using a single variable $D(\vec{y}) = N(\vec{y}) + M(\vec{y})$. Figure 1 shows a typical response to a repeated inoculation with antigen \vec{x} . Although other lymphocytes also bind less effectively to the antigen, for illustrative purposes we plot only populations $E(\vec{y})$ and $D(\vec{y})$ for $\vec{x} = \vec{y}$, as well as F^* (for which binding is independent of shape space characteristics).

Initially there are 15 memory cells with $\vec{x} = \vec{y}$, $D(\vec{y}, 0) = 15\delta(\vec{x} - \vec{y})$, $F = H$, $E = 0$, and an antigen inoculation $A(\vec{x}, 0) = 110$. After the first immune response is complete, there is a second identical inoculation at a later time. In each exposure, the population size of the antigen increases, until a sufficient number of effector cells are created from the memory and naive cell populations to eliminate the infection. The total number of lymphocytes $N + M + E$, increases during an immune response, corresponding to swelling of the lymph nodes. As the effectors die and the memory and naive cells are no longer stimulated, the swelling subsides. Additionally, this model predicts that symptoms associated with elevated E levels peak just as the pathogen is cleared. The more rapid secondary response is due to the elevated number of memory cells. (The initial steep decline in the secondary response is due to the trapping of antigen by the dendritic cells.) All other model parameters remain the same from the first exposure to the second. Between infections all short-lived effector cells die, and the F^* cells all revert to F .

To quantify the severity of an individual infection, we




Figure 1. Immune Responses to Two Sequential Inoculations by the Same Antigen

Results are shown for maximum binding affinity pairs $\vec{x} = \vec{y}$. The rapid response to the secondary inoculation (represented by the smaller size of the second peak) is due to the elevated number of memory cells. The immunological loss, Equation 15, is defined to be the area under the antigen population peak. For the primary and secondary peaks, the values are 8,860 and 1,525, respectively. The model parameters used in this and all simulations in this paper are as follows: $\alpha = 1.5$, $\beta = 0.083$, $f = 0.38$, $\delta = 0.01$, $\rho = 1$, $\varphi = 10^{-4}$, $H = 10$, $\gamma_{max} = 0.005$, $b = 2$. These parameter values give typical behavior for the model. doi:10.1371/journal.pcbi.0020160.g001

define a loss function $L(\vec{x})$ as the integral of the antigen population size with respect to time:

$$L(\vec{x}) = \int A(\vec{x}, t) dt. \quad (7)$$

While physiologically, severity of disease depends on many factors, we believe that this is a simple natural choice, as it is a rough measure of the amount of the body's resources a pathogen may consume and the amount of toxin the pathogen may secrete. This immunological loss function serves as a tool for quantifying statistics of infection size, and provides a meaningful target for sensitivity analysis. In the context of this investigation and immunosenescence, it allows us to quantify fitness and to monitor how it changes over the development of the immune system. Additionally, loss can be used to compare the effects of additional immune system components and reactions in more detailed immune system models, and to quantify the efficacy of drugs and therapies based on their effect on loss.

We can obtain analytic estimates for loss as well as memory-cell population growth as functions of pre-infection memory and naive cell population sizes based on several simplifying approximations to Equations 2–6. For $\rho A \gg \alpha \gamma$, F^* is approximately equal to the total number of dendritic cells H , $F^* \approx H$, and $F \approx 0$. Since A levels will be high when an immune response is initiated, this approximation is reasonable. Equations 2–6 with M and N replaced with $D = M + N$, and the approximation $F^* \approx H$, reduce to:

$$\frac{\partial A(\vec{x})}{\partial t} = \beta A(\vec{x}) - A(\vec{x}) \int \gamma(\vec{x}, \vec{y}) E(\vec{y}) d\vec{y}, \quad (8)$$

$$\frac{\partial E(\vec{y})}{\partial t} = 2f\alpha \gamma(\vec{x}, \vec{y}) H D(\vec{y}) - \delta E(\vec{y}), \quad (9)$$

$$\frac{\partial D(\vec{y})}{\partial t} = (1 - 2f)\alpha\gamma(\vec{x}, \vec{y})HD(\vec{y}). \quad (10)$$

These equations can be easily integrated, yielding solutions that approximate the antigen population size during an infection. The complete expression for $A(\vec{x}, t)$ is tractable, but cumbersome, and takes the form

$$A(\vec{x}, t) = A(\vec{x}, 0)e^{\beta t - \int d\vec{y} S(\vec{x}, \vec{y}, t)}. \quad (11)$$

A simple expansion of the function $S(\vec{x}, \vec{y}, t)$ to second order in t yields a Gaussian approximation for the $A(\vec{x}, t)$ peaks (e.g., in Figure 1):

$$A(\vec{x}, t) = A(\vec{x}, 0)e^{\beta t - B(\vec{x})t^2} \quad (12)$$

where,

$$B(\vec{x}) = f\alpha H \int d\vec{y} \gamma^2(\vec{x}, \vec{y})D(\vec{y}, 0). \quad (13)$$

This approximation describes the $A(\vec{x}, t)$ pulse as a function of the initial value of D .

Using this approximate solution for $A(\vec{x}, t)$, we estimate the increase in memory cell–population values after the infection is cleared. We take the value of $M(\vec{y}, t)$ to be constant after time t_e when A has been reduced to half its initial value, in the tail of the pulse, and we round it to integer value.

$$M(\vec{y}, t_e) = D(\vec{y}, 0)e^{(1-2f)\alpha\gamma H t_e} - N(\vec{y}, 0)e^{-\gamma\alpha H t_e} \quad (14)$$

This analytical result gives close agreement with memory cell–growth levels given by our original model.

We estimate loss by integrating our analytical solution for the antigen population peak from $-\infty$ to ∞ (rather than starting at $t = 0$) to obtain:

$$L(\vec{x}) \approx A(\vec{x}, 0) \sqrt{\frac{\pi}{B(\vec{x})}} \exp\left[\frac{\beta^2}{4B(\vec{x})}\right]. \quad (15)$$

Note that extending the range of integration to $-\infty$ makes a relatively small difference in the result and simplifies this expression. Furthermore, it may in a certain sense be more accurate, as it accounts for the proliferation of antigen before the antigen enters the lymph nodes.

On intermediate time scales the system relaxes, homeostasis adjusts naive cell number, and the naive cell population is recycled. These processes are considered fast enough to reach a steady state during the time between infections, but not so fast as to be a factor during an immune response. In the absence of antigen, the populations of effector and activated dendritic cells (which are both responsible for removing antigen from the body) relax back to zero ($E(\vec{y}) = 0$, $F^* = 0$, and $F = H$), as illustrated in Figure 1. Though during an immune response N and M cells play an identical role (represented as D in Figure 1), during the homeostatic period their differences become important. The memory cells are long-lived, and in the absence of antigen their population is static. Naive cells have a shorter lifetime than memory cells and die by apoptosis. As the naive cells die, homeostatic mechanisms stimulate the cells of the bone marrow to randomly repopulate the system with new naive cells. The repopulation is constrained by the total number of D cells, R :

$$R = \int [M(\vec{y}) + N(\vec{y})] d\vec{y} = M_{tot} + N_{tot} \quad (16)$$

This constraint is violated during an immune response as the lymphocytes rapidly proliferate. Once the antigen is cleared, the total relaxes back to R . Thus, as memory cell populations rise, homeostasis effectively depletes the naive cell population. The replacement of naive cell populations with memory cells with increasing age is described by Linton and Dorshkind [8].

Next, using Equations 14–16 and simulated naive cell recycling, we study the long-term adaptation of lymphocyte populations over the lifetime of the individual. Using these approximations, our model reduces to a cellular automaton describing the population changes of lymphocytes on the shape space after each infection under our homeostatic constraint, Equation 16. Initially the system is composed of R naive cells. The naive cells randomly populate the shape space with uniform probability. The system is then inoculated with antigen at position \vec{x} with probability $P(\vec{x})$. The corresponding loss is computed in Equation 15, as well as the change in the memory cell population in Equation 14. The naive cells are then redistributed with their number adjusted to satisfy Equation 16. A subsequent inoculation of the same antigen will make use of these memory cells for a more rapid response, but an inoculation at another point in shape space will have a reduced number of naive cells with which to respond and the loss will be higher.

We monitor the evolution of loss on long-time scales by considering a 70×70 lattice with $n = 36$ possible infections, at sites evenly distributed, indexed i , occurring with probability p_i . The infections are far enough apart that cross-reactivity is not a factor. The distribution of infection probabilities is taken to have a few chronic infections that are very likely to recur, and many rare infections. The probabilities are given by an exponential distribution:

$$p_i = \zeta e^{-i/\xi}, \quad (17)$$

where we set $\xi = 20/3$, and $\zeta = \sum_{j=1}^n e^{-j/\xi}$ normalizes the distribution over the discrete set of n infections. We choose a distribution of this form to have a mix of frequent and rare infections. We have used other distributions as well (e.g., a power law distribution), and obtained similar results. Different distributions alter the memory cell population–growth rate, which affects the time scale for the onset of immunosenescence. Realistically, the distribution of diseases and their respective infection probabilities is itself a dynamic coevolving system with new diseases constantly arising. In such a dynamic disease distribution, when the naive cell population is depleted there will be fragility similar to the observations reported here. With the kind of static distribution we consider here, in order for fragility to develop, the rare diseases must have low enough probability that one of them is likely to be experienced for the first time once the naive cell population is depleted. Changing the numerical values of parameters in the model will in general change the rate at which the naive cell population is depleted. We ran the simulation for 400 infections drawn at random from the above distribution. Figure 2A (top) shows the loss $L(\vec{x})$ for each event in a representative sequence.

Figure 2B (top) illustrates the corresponding lymphocyte populations ($D(\vec{y}) = M(\vec{y}) + N(\vec{y})$) on the shape space at three stages in the adaptive development: initial, after 250 infections, and after 400 infections. The corresponding loss fields are illustrated below (these illustrate the loss that would




Figure 2. One Realization of System Development for 400 Infections on a 70×70 Lattice

There are 36 possible infections evenly distributed throughout shape space with $R = 33,000$. The distribution of probabilities for the infections approximates an exponential, Equation 17.

(A) The top curve illustrates the actual losses based on the history of infections. For most of the simulation, expected loss J , Equation 18, tends to decrease, yet rare events are increasing in size resulting in increased variance σ , Equation 19. Eventually this results in catastrophic failure.

(B) Illustrates the shape space representation of populated receptor sites in the immune system initially, after 250 infections, and after 400 infections (top), and the corresponding distributions of losses for subsequent infections (bottom). The initial configuration (left) is randomly populated by naive cells. After each infection there are elevated populations of memory cells in the vicinities of the infection site. The bottom figures illustrate how the immune system becomes skewed in favor of rapid response to repeated exposures at the expense of novel infections, by illustrating the immunological losses, Equation 15, that would be incurred by inoculations at each lattice point before and after building up memory-cell populations.

doi:10.1371/journal.pcbi.0020160.g002

be incurred for a subsequent infection as a function of the antigen characteristics \vec{x}). In Figure 2B, top left, D is strictly composed of naive cells. In Figure 2B, top middle, D includes a mix of memory cells that form in the neighborhood of the inoculations and the recycled naive cells. The rightmost figure is almost entirely depleted of naive cells. The bottom images show $L(\vec{x})$ for the naive state and after inoculations. Figure 2B, bottom, illustrates what the loss would be, given the D values in Figure 2B, top, for an inoculation at each point on the lattice (though we only consider the 36 points to be possible infections). Initially there are few vulnerabilities, associated with potentially large losses (red), in the system. Instead, the system is uniformly protected. However, after 250 infections, the system develops structure and has areas of high potential loss around the rare antigens. The points

around the most common infections are well-protected after 250 infections, indicating low values of loss (dark blue) in subsequent infections. However, because of the overall constraint on the number of cells, many outlying areas are left more vulnerable than they were initially.

Based on the probability distribution of infections, we calculate the expected loss J at each stage of the system's adaptation. Here J corresponds to the average value of $L(\vec{x})$ computed over the full spectrum of possible infections and weighted according to the probability of each infection:

$$J = \int P(\vec{x})L(\vec{x})d\vec{x}. \quad (18)$$

The standard deviation σ gives a measure of the corresponding variation in the possible loss values:

$$\sigma = \sqrt{\int P(\vec{x})L(\vec{x})^2d\vec{x} - J^2} \quad (19)$$

The expected loss J (Figure 2A, middle) initially decreases from the starting value, associated with the random population of shape space. As the system adapts, J takes its minimum value at roughly 250 infections, which we refer to as the "optimal" state. In later stages J begins to rise, due to overspecialization. It is this increase that we associate with immunosenescence. Throughout the simulation, adaptation is accompanied by a steady increase in the variability σ (Figure 2A, bottom), associated with increasing breadth in the distribution of losses as the system becomes increasingly specialized. At the latest stages of the simulation the increase in σ sharpens, which is indicative of extreme vulnerability to rare events.

Figure 3 illustrates the cumulative statistical distribution of loss sizes obtained by combining data from 600 simulations of the form illustrated in Figure 2. The initial state is characterized by a narrow (note the logarithmic axes) and flat distribution, which reflects the uniform coverage of shape space by the random population of naive cells. The green curve corresponds to the optimal state, where the expected loss J takes its minimum value. Compared with the initial state, here the distribution of losses is both broader and more variable, indicative of adaptation that optimizes the inherent tradeoff between reducing loss sizes for frequent events, at the cost of larger losses for less frequent infections, which arises because of the overall resource constraint (Equation 16). The red curve shows the result at the end of our simulation, when the system has overspecialized, and exhibits immunosenescence. In this case, the distribution of losses is extremely heavy-tailed, corresponding to the increase in J . Any distribution containing very rare events leads to heavy-tailed loss statistics as the naive cell population becomes depleted. This heavy-tailed distribution of loss shows immunosenescence in the increased fragility of an aged immune system to as yet unseen diseases.

Discussion

Our model is representative of the HOT mechanism [9,10], in which robustness tradeoffs provide a mechanism for complexity and power laws through either deliberate design or biological evolution, both of which favor configurations that minimize loss (Equation 18) subject to resource

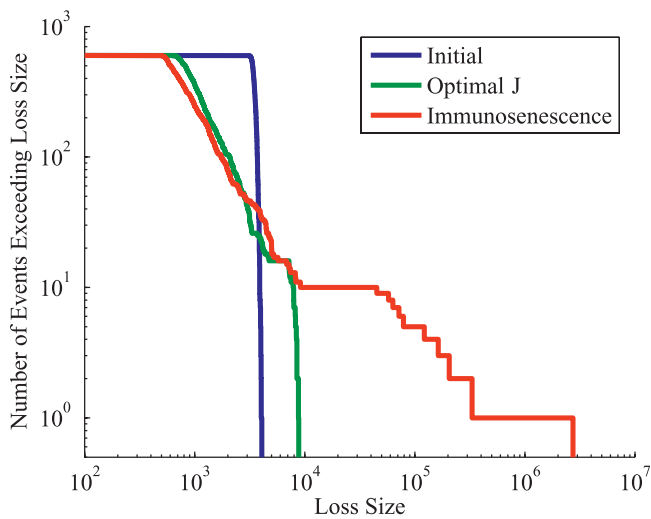


Figure 3. Distribution of Losses over 600 Realizations of the System
Results are shown after the first infection, after 250 infections (when J is at its lowest value, corresponding to the optimal state), and after 400 infections (corresponding to immunosenescence), in blue, green, and red respectively. The infection probabilities have the same distribution as in Figure 2.
doi:10.1371/journal.pcbi.0020160.g003

constraints (Equation 16). The simplest examples are referred to as “Probability Loss Resource” (PLR) HOT models [11–13], which incorporate physically motivated relationships between resource allocations and loss sizes of individual events to define a constrained resource optimization problem involving a set of events with prescribed probabilities. In the cases that have been studied to date, resources have acted as barriers to propagation of cascading events, such as wildfires [14] or power outages [15]. In our case, the analogy is more akin to a sprinkler system, populated by lymphocytes, in the shape space of possible pathogens. Over time, adaptation leads to specialized states, through replacement of naive cells with memory cells, which are tuned to the history of past exposures. This results in a system that is increasingly robust to common disturbances, yet increasingly fragile to rare events—a key signature of HOT. In our model, this age-correlated effect is a result of overspecialization rather than of an accumulation of defects. Other possible factors, such as deterioration, may contribute to immunosenescence as well, though it has been experimentally observed that some symptoms are due to system dynamics [16].

Consequences of overspecialization were studied previously in a HOT model of evolution, based on Darwinian mechanisms [17,18], leading to extreme vulnerability, similar to our observations here. In that case, offspring of lattice organisms evolved through random mutation relative to their parent lattice, and fitness was based on disturbances over the lifetime of individual lattices. Competition resulted in development of generalists and specialists. While specialists

flourished during common circumstances, they experienced episodic extinction during rare events, which parallels the extreme fragility in our model associated with immunosenescence. In that case, the mutation rate itself was subject to mutation, and high mutation rates played an important role in rapid diversification and evolution following an extinction of the specialists. In the immune system, rapid mutation is associated with somatic hypermutation, which gives the daughter cells of the lymphocyte stimulation a receptor that is a mutation of the parent’s corresponding receptor. This gives rise to higher affinity and more efficient responses [19], and will be considered in future work.

While we have focused on immunosenescence, there are numerous additional robustness tradeoffs associated with the immune system. For example, the immune system has the ability to attack and remove nonself elements from the body with no prior knowledge of nonself features. Normally this is done with little harm to the body itself. However, the immune system can make mistakes in recognition, leading to autoimmune disease, a fragility which would not be present if an organism had no immune system to begin with. In addition, on adaptive time scales, the ability to retain memory of past exposures enables development of effective vaccines, and reduces the severity of outbreaks of communicable diseases within populations. However, in some instances vaccinations may also lead to increased susceptibility to similar diseases [19,20]. This “robust yet fragile” behavior is a key feature of HOT, a statistical theory for complexity in designed, evolved, or adaptive systems. The immune system can be viewed as a complex system in which robustness tradeoffs play a central role in evolution of the basic operating mechanisms as well as adaptation of cell populations within an individual. We emphasize the importance of tradeoffs associated with a spectrum of possible events. Evolution and adaptation favor increased robustness to common disturbances, but this is inevitably paired with increased fragility, both to rare events and to new opportunities for diseases and disturbances to hijack the system, which would not be available were the system not in place.

Materials and Methods

All computations and stochastic and numerical simulations were done using MATLAB (The Mathworks, <http://www.mathworks.com>) on a personal computer.

Acknowledgments

Author contributions. JC and SPS conceived, designed, and performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, and wrote the paper.

Funding. This work was supported by the David and Lucile Packard Foundation, the James S. McDonnell Foundation, and the Institute for Collaborative Biotechnologies through grant DAAD19-03-D-0004 from the US Army Research Office.

Competing interests. The authors have declared that no competing interests exist.

References

1. Goldsby RA, Kindt TJ, Osborne BA, Kuby J (2003) Immunology. Fifth edition. New York: W. H. Freeman. 551 p.
2. Perelson AS, Weisbuch G (1997) Immunology for physicists. *Rev Mod Phys* 69: 1219–1267.
3. Segel LA, Perelson AS (1988) Theoretical immunology. Part one. SFI studies in the sciences of complexity. Redwood City (California): Addison-Wesley. 404 p.
4. Segel LA, Perelson AS (1988) Theoretical immunology. Part two. SFI Studies in the sciences of complexity. Redwood City (California): Addison-Wesley. 407 p.
5. Haynes L, Swain SL, Cambier J, Fuldner R (2005) Aging and immune function: Summary of a workshop held at Trudeau Institute, Saranac Lake, New York. *Mech Ageing Dev* 126: 822–825.
6. Perelson AS, Oster GF (1979) Theoretical studies of clonal selection:

- Minimal antibody repertoire size and reliability of self–non-self discrimination. *J Theor Biol* 81: 645–670.
7. Smith DJ, Forrest S, Hightower RR, Perelson AS (1997) Deriving shape space parameters from immunological data. *J Theor Biol* 189: 141–150.
 8. Linton PJ, Dorshkind K (2004) Age-related changes in lymphocyte development and function. *Nat Immunol* 5: 133–139.
 9. Carlson JM, Doyle J (1999) Highly optimized tolerance: A mechanism for power laws in designed systems. *Phys Rev E* 60: 1412–1427.
 10. Carlson JM, Doyle J (2000) Highly optimized tolerance: Robustness and design in complex systems. *Phys Rev Lett* 84: 2529–2532.
 11. Doyle J, Carlson J (2000) Power laws, HOT, and generalized source coding. *Phys Rev Lett* 84: 5656–5659.
 12. Carlson J, Doyle J (2002) Complexity and robustness. *Proc Nat Acad Sci U S A* 99: 2538–2545.
 13. Manning M, Carlson JM, Doyle J (2005) Highly optimized tolerance and power laws in dense and sparse resource regimes. *Phys Rev E* 72: 016108.
 14. Moritz MA, Morias M, Carlson JM, Doyle J (2005) Wildfires: An example of highly optimized tolerance? *Proc Nat Acad Sci U S A* 102: 17912–17917.
 15. Stubna MD, Fowler J (2003) An application of the highly optimized tolerance model to electrical blackouts. *Int J Bifurcation Chaos* 13: 237–242.
 16. Geiger H, Zant GV (2002) The aging of lympho-hematopoietic stem cells. *Nat Immunol* 3: 329–333.
 17. Zhou T, Carlson JM, Doyle J (2002) Mutation, specialization, and hypersensitivity in highly optimized tolerance. *Proc Nat Acad Sci U S A* 99: 2049–2054.
 18. Zhou T, Carlson JM, Doyle J (2005) Evolutionary dynamics and highly optimized tolerance. *J Theor Biol* 236: 438–447.
 19. Deem MW, Lee HY (2003) Sequence space localization in the immune system response to vaccination and disease. *Phys Rev Lett* 91: 837–840.
 20. de St Groth SF, Webster RG (1966) Disquisitions on original antigenic sin. *J Exp Med* 12: 331.