Accelerated crossing of fitness valleys through division of labor and cheating

Natalia L. Komarova, Erin Urwin, and Dominik Wodarz

Supplementary Information

Contents

1 Selection-mutation balance in the absence of cheaters

Some important aspects of the system can be described by means of ordinary differential equations. In this section we present these equations and show under which circumstances the emergence of m-hit mutants can be observed. We start with a simple, 2-component example, and then expand it to the case relevant for this paper.

1.1 A two-component system

Let us consider a very simple selection-mutation network of the form $a \to A$ where the type a has the division rate r_0, the type A had the division rate r_1, both types have the death rate of d, and the mutation rate is given by u. We are assuming that both are viable types, such that

$$r_0(1-u) > d, \quad r_1 > d. \tag{1}$$

We further assume that the second type is disadvantageous,

$$r_1 < r_0.$$

Imposing logistic growth of the species with a carrying capacity, K, we can describe the dynamics of such a system by a set of ODEs,

$$\dot{x}_0 = r_0 x_0 \left(1 - \frac{x_0 + x_1}{K}\right)(1-u) - d x_0, \tag{2}$$

$$\dot{x}_1 = (r_0 x_0 u + r_1 x_1)\left(1 - \frac{x_0 + x_1}{K}\right) - d x_1, \tag{3}$$

where x_0 is the number of cells of the first type (unmutated), and x_1 is the number of cells of the second type (mutated). Note that the symbol K stands for carrying capacity in the ODEs, and the same symbol is used for the replication/cooperation radius in the main paper. Insights obtained from the theory developed here are applicable to spatial systems when the replication and cooperation radii are the same and equal to K.

Equations (2-3) have three equilibria:

- The trivial equilibrium, $x_0^t = x_1^t = 0$. This equilibrium is unstable as long as any of the conditions (1) holds.

- The selection-mutation balance equilibrium,

$$x_0^{(0)} = \frac{K(r_0(1-u)-d)\left(1-\frac{r_1}{r_0}-u\right)}{(r_0-r_1)(1-u)}, \quad x_1^{(0)} = \frac{K(r_0(1-u)-d)u}{(r_0-r_1)(1-u)}.$$

  This equilibrium is stable if

$$u < 1 - \frac{r_1}{r_0}, \tag{4}$$

  that is, if the mutation rate is smaller than the relative difference in the two types' growth rates. At this equilibrium, the weaker phenotype is maintained by a constant production through mutations from the stronger phenotype.

- The dominance of the second phenotype equilibrium,

$$x_0^{(1)} = 0, \quad x_1^{(1)} = K\left(1-\frac{d}{r_1}\right),$$

  where the second phenotype reaches the effective carrying capacity. This equilibrium is stable when condition (4) is reversed. If the mutation rate is high enough to offset the difference in division rates, the second phenotype becomes effectively stronger and dominates the system.

This picture changes somewhat in the presence of back-mutations. Now, the equations are

$$\dot{x}_0 = (r_1x_1u_1 + r_0x_0(1-u))\left(1-\frac{x_0+x_1}{K}\right) - dx_0, \tag{5}$$

$$\dot{x}_1 = (r_0x_0u + r_1x_1(1-u_1))\left(1-\frac{x_0+x_1}{K}\right) - dx_1, \tag{6}$$

where we denoted by u_1 the rate of back-mutations. For a non-zero rate u_1, there is only one stable solution, as illustrated in figure 1(a). When the value of u is small, the stronger phenotype, x_0, dominates the system, and the weaker phenotype is maintained at a selection-mutation balance. As u increases, the frequency of x_1 increases. When the value of u approaches $1 - r_1/r$, the relative fitness difference, the type x_1 increases in proportion, and for sufficiently large u it starts dominating the system. Exactly when

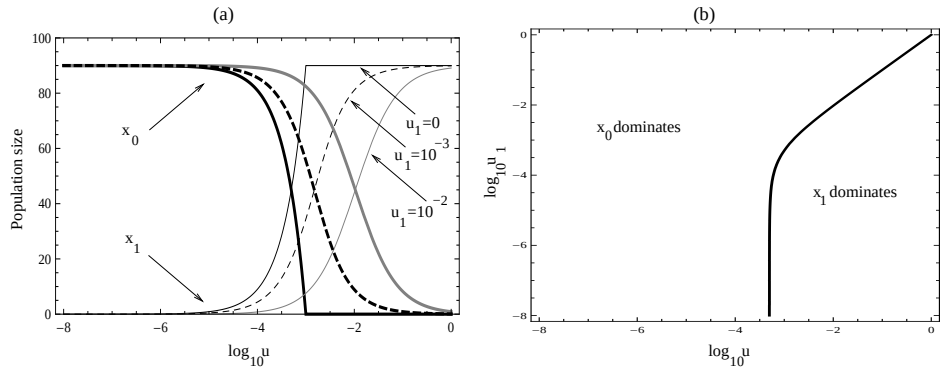

Figure 1: Dynamics of the two-component system in the presence of back-mutations, equations (5-6). (a) Stable equilibrium values of for the variables x_0 (thick lines) and x_1 (thin lines). The horizontal axis is $\log_{10} u$, and the solutions for three different values of u_1 are presented. (b) A contour plot showing the regions in the $u - u_1$ space where types x_0 and x_1 dominate (the contour corresponds to the equation $x_1 = K(1 - d/r)/2$). Other parameters are: $K = 100$, $r_0 = 1$, $r_1 = 0.999$, $d = 0.1$.

that happens depends on the value of u_1. The larger u_1, the harder it is for x_1 to dominate, see figure 1(b). For small values of u_1, there is a more or less constant threshold value of u after which the weaker type can dominate. For larger values of u_1, the value of u has to be larger than u_1 to sustain a large frequency of x_1.

In the important case where $u_1 = u$, we can see that making the mutation rates large does not lead to a dominance of the weaker phenotype. Instead, x_1 at equilibrium is maintained at a frequency similar to that of x_0 (figure 1(b)).

1.2 A multiple-component system

Next, we consider a more general system relevant for this work. Assuming m sites that can be mutated, we have the total of 2^m possible types, connected by mutation processes. If we assume that the division rates of different types are only defined by the number of mutated sites, then all the types can be split into $m + 1$ symmetric classes according to the number of mutated sites. The equations for all the types within one class are identical, and the solutions are also identical (if we assume that the initial conditions are the same within classes). Let us denote by r_i the division rate of cells with i sites mutated.

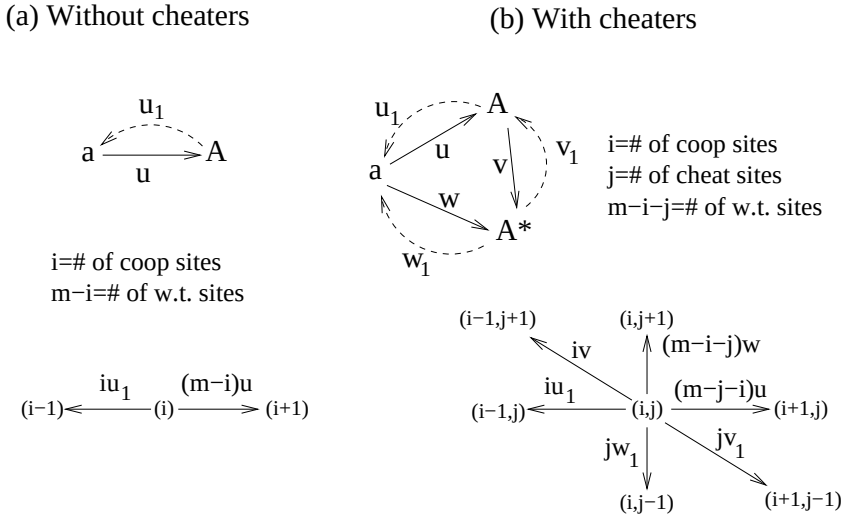

Figure 2: A schematic illustrating the derivation of equations (8) (without cheaters, (a)) and (9) (with cheaters, (b)). Index i refers to the number of cooperating sites, and index j to the number of cheating sites.

We further denote by y_i the number of cells of any one type in class i (i sites mutated). The number of types within class i is given by $m!/i!/(m-i)!$, and therefore the total number of cells in class i is $y_i m!/i!/(m-i)!$. The equations that these variables satisfy (and which are a direct generalization of equations (2-3)) are as follows:

$$\dot{y}_i = [ir_{i-1}y_{i-1}u + r_i y_i(1-(m-i)u)]W - dy_i, \quad 0 \le i \le m, \tag{7}$$

where the logistic growth restriction factor is

$$W = 1 - \frac{\sum_{j=0}^{m} \frac{y_j m!}{j!(m-j)!}}{K}.$$

To illustrate the derivation of equation (7), we use a specific example. Suppose that $m = 5$, and $i = 3$. Then y_i represents any type with three mutations, for example type $xxx00$ (where x stands for a mutated site and 0 for a wild-type site). Such cells can be generated by one mutation in only three ways: from type $xx000$, from type $x0x00$ and from type $0xx00$. There are $i = 3$ such types, and the abundance of each of them is y_{i-1}. Therefore, the corresponding production term in equation (7) is $ir_{i-1}y_{i-1}$.

In the presence of back-mutations, the system corresponding to equations (5-6) is as follows:

$$\begin{aligned} \dot{y}_i &= [ir_{i-1}y_{i-1}u + (m-1)r_{i+1}y_{i+1}u_1 + r_i y_i(1-(m-i)u - iu_1)]W \\ &- dy_i, \quad 0 \le i \le m. \end{aligned} \tag{8}$$

Figure 2(a) illustrates the derivation of these equations. Each cell type characterized by i mutations can give rise to $(m-i)$ other types by a forward mutation, u, and to i other types by a back-mutation, u_1. This is because type i has $m-i$ unmutated sites, which can give rise to a type $i+1$ cell. It has i mutated sites, which can give rise to a type $i-1$ cell. Thus the total mutation rate of type i is $(m-i)u + iu_1$; this is reflected in the term $r_i y_i(1-(m-i)u - iu_i)$ in equation (8). On the other hand, the same types that can originate from type i, can give rise to type i: types $i-1$ result in type i by means of a forward mutation (this corresponds to the term $ir_{i-1}y_{i-1}u$ in equation (8)), and types $i+1$ result in type i by means of a back-mutation (term $(m-1)r_{i+1}y_{i+1}u_1$ in equation (8)).

The fitness parameters are as follows. The various classes of mutants have division rates $R < r_i = R^+ - if < R^+$. The wild-type cells have the division rate $r_0 = R$ if they do not participate in the division of labor, and $r_0 = R^+$ otherwise. Note that in this description we are assuming that the partial mutants always cooperate with each other, and therefore their fitness never falls to the value $R^- < R$. The m-hit mutants are the least fit of all the mutants, because they share the largest number of gene products. The question is, can the m-hit mutants get established at significant frequencies in this setting?

Analysis of equation (7) shows that m-hit mutants can dominate the system in the absence of back-mutations as long as the fitness cost of cooperation (given by $mf/(aR)$) is less than the mutation rate, u, see figure 3. When we add back-mutations, however, the picture changes. If we assume that $u_1 = u$, then, although the population of the m-hit mutants can reach the abundance similar to that of the other types (as in figure 1(b)), the combinatorial factors make the overall relative frequency of the m-hit mutants low. We conclude that even for high mutation rates, in the presence of back-mutations, m-hit mutants cannot dominate without cheaters in the system.

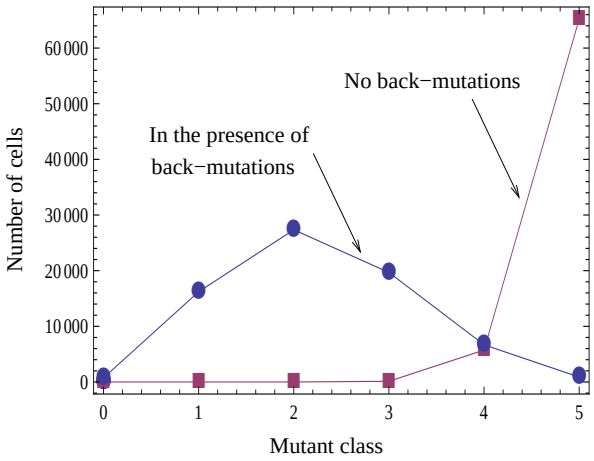

Figure 3: The m-component system without cheaters in the absence and in the presence of back-mutations, equations (8). The number of cells of each type at equilibrium is shown. The parameters are $K = 90,000$, $R = 0.4$, $R^+ = 0.5$, $d = 0.1$, $f = 1/200$, $u = 10^{-2}$. In the presence of back-mutations, $u_1 = u$.

2 Cheaters

We can include cheaters in the equations, as illustrated in figure 2(b). Let us denote by i the number of cooperating sites and by j the number of cheating sites. All the types can be split into classes, (i, j). We will denote by y_{ij} the abundance of cells of a particular type from class (i, j). There are $m!/(i!j!(m-i-j)!$ different types in such a class. In the most general case, there are three possible forward mutation rates, u, v and w, as shown in figure 2(b). Each of them has a corresponding back-mutation rate, u_1, v_1 and w_1. A type (i, j) can give rise to six different types (and be created from 6 different types). All the connections result in the following equations:

$$\begin{aligned}
\dot{y}_{i,j} &= [r_{i,j}y_{i,j}(1-(m-i-j)u - iu_1 - (m-i-j)w - jw_1 - iv - jv_1) \\
&+ r_{i-1,j}y_{i-1,j}iu + r_{i+1,j}y_{i+1,j}(m-i-j)u_1 + r_{i,j+1}y_{i,j+1}(m-i-j)w_1 \\
&+ r_{i,j-1}y_{i,j-1}jw + r_{i-1,j+1}y_{i-1,j+1}iv_1 + r_{i+1,j-1}y_{i+1,j-1}jv]W \\
&- dy_{i,j}, \quad 0 \leq i+j \leq m. \qquad (9)
\end{aligned}$$

The division rates of the various types are defined by the number of cooperating sites. In the case where the wild type does not participate in the

division of labor, we have

$$r_{0,0} = R, \quad R < r_{i,j} = R^+ - if \leq R^+, \quad ij > 0. \tag{10}$$

In the case where the wild type participates in the division of labor, we have

$$r_{0,0} = R^+,$$

with the rest of the fitness values the same as in (10).

It is convenient to split the population into the following 4 groups:

- The wild type, $Y_{w.t.} = y_{0,0}$;
- The full cheaters: m-hit mutants that cheat at all the sites, $Y_{full\ cheat} = y_{0,m}$;
- The cooperators: cells that cooperate at one or more sites, $Y_{coop} = \sum_{i=1}^{m} \sum_{j=0}^{m-i} \frac{y_{i,j} m!}{i!j!(m-i-j)!}$;
- The partial cheaters: cells that do not cooperate, and have between 1 and $m-1$ cheating sites, $Y_{part.\ cheat} = \sum_{j=1}^{m-1} \frac{y_{0,j} m!}{j!(m-j)!}$.

Note that the population of group Y_{coop} includes m-hit mutants (types $y_{i,m-i}$ for $1 \leq i \leq m$). However, these types have a slower replication rate than the type $Y_{full\ cheat}$, because they are penalized for cooperation. Therefore, the m-hit mutant which has a potential to dominate the population is the $Y_{full\ cheat}$ type.

In the presence of cooperators, the replication rate of the class $Y_{part.\ cheat}$ is identical to that of $Y_{full\ cheat}$. In this case the groups $Y_{part.\ cheat}$ and $Y_{full\ cheat}$ have the highest fitness. If the class Y_{coop} is not present, the cooperation disappears and the class $Y_{part.\ cheat}$ operates at low fitness, $R^- < R$, and class $Y_{full\ cheat}$ wins. If only some of the components are missing from the cooperating pool, the group $Y_{part.\ cheat}$ has a lowered average fitness compared to that of $Y_{full\ cheat}$.

Let us suppose that cooperation is fully present in the system. Then the following is observed. If all the mutation rates are equal to each other, then because of combinatorial effects, the relative proportion of the middle types (with $j \sim m/2$) is the largest among the fittest types (the types $(0, j)$, with $1 \leq j \leq m$). If we set $w_1 = 0$, then the group $Y_{full\ cheat}$ acquires an advantage because it has fewer mutations radiating from it than the group $Y_{part.\ cheat}$.

If we set $w_1 = v_1 = 0$, then group $Y_{full\ cheat}$ is the absolute winner, no mutations come out of it, and this makes it the fittest type. This demonstrates that in the absence of some back-mutations, m-hit mutants can arise quite rapidly in this setting; this can happen on a much faster time-scale than in the absence of cooperators or cheaters (sequential evolution). However, this requires the assumption about the unidirectionality of mutations. Below we will concentrate on the more subtle scenario where all back-mutations are present.

Let us suppose that all the mutation rates are the same. Then the only way for group $Y_{full\ cheat}$ to dominate is to get rid of cooperation, such that the fitness of group $Y_{part.\ cheat}$ plunges. Cooperators (the group Y_{coop}) have a smaller division rate than $Y_{full\ cheat}$ (given that cooperation is present), and thus they are maintained at a selection-mutation balance. The population of Y_{coop} is dominated by types $y_{1,j}$, $0 \leq j \leq m-1$, and their abundance is proportional to the mutation rate. Decreasing the mutation rate will change the equilibrium number of the cooperators. If it is sufficiently low, cooperation can be absent from the system in part for a lot of the time, thus giving type $Y_{full\ cheat}$ a selective advantage. This is a possible stochastic mechanism by which an m-hit mutant can arise in this system. Figure 4 shows a typical steady-state solution for different classes in the case where the wild types do not participate in the division of labor. The picture is similar, except for a higher abundance of the wild types, if they do participate in division of labor. In the next section, we find the steady state of system (9) analytically.

3 The analytic solution for the steady state

3.1 Wild type cells do not participate in the division of labor

We can find the steady state for system (9) analytically, in the limit of small mutation rates. Let us for simplicity assume that all the mutation rates are equal,

$$u_1 = w = w_1 = v = v_1 = u,$$

and that $u \ll 1$. We will use the following expansion:

$$y_{0,0} = 0 + u y_{0,0}^{(1)} + O(u^2), \tag{11}$$

$$y_{0,j} = y_{0,j}^{(0)} + u y_{0,j}^{(1)} + O(u^2), \quad 1 \leq j \leq m, \tag{12}$$

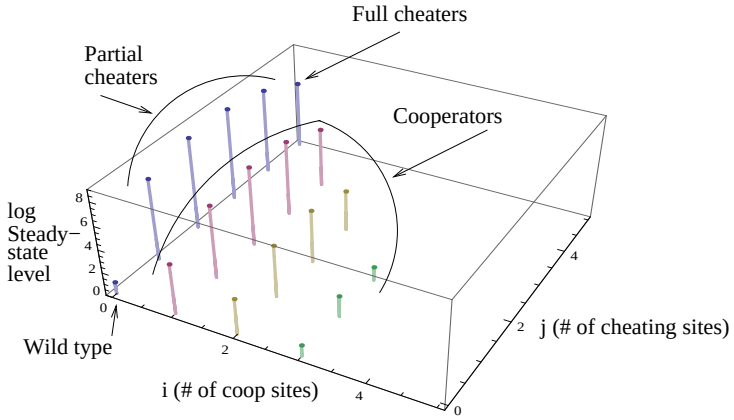

Figure 4: The logarithms of the steady-state values for different classes of variables are shown (the values $m!/[i!j!(m-i-j)!]y_{i,j}$). Parameters are $m = 5$, $K = 120 \times 120$, all mutation rates are $10^{-2.5}/2$, $aR = 0.5$, $R = 0.15$, $f = 1/70$.

$$y_{1,j} = 0 + uy_{1,j}^{(1)} + O(u^2), \quad 0 \le j \le m-1, \tag{13}$$

$$y_{i,j} = 0 + 0u + O(u^2), \quad 1 \le i \le m, \quad 1 \le i+j \le m. \tag{14}$$

This expansion is motivated by the fact that in the absence of mutations, only the types with the highest fitness will be present, which are the cheaters ($y_{0,j}$ with $j > 0$). Further, if a small mutation rate is added, the quasispecies will contain a cloud of other types, each proportional to the number of mutational steps separating it from the winning phenotype. The classes most closely related to the winning phenotype are the wild type ($y_{0,0}$) and the cooperators with one cooperating site ($y_{1,j}$ with $0 \le j \le m-1$). Their abundances scale with u. All the other types will be significantly less frequent.

For the saturation term, we have $W = W^{(0)} + uW^{(1)} + O(u^2)$, where

$$W^{(0)} = 1 - \frac{1}{K}\sum_{j=1}^{m} \frac{y_{0,j}^{(0)} m!}{j!(m-j)!}, \tag{15}$$

$$W^{(1)} = -\frac{1}{K}\left(\sum_{j=0}^{m} \frac{y_{0,j}^{(1)} m!}{j!(m-j)!} + \sum_{j=0}^{m-1} \frac{y_{1,j}^{(1)} m!}{j!(m-1-j)!}\right). \tag{16}$$

The zeroth order term of equations (9) yields,

$$r_{0,j} y_{0,j}^{(0)} W^{(0)} - d y_{0,j}^{(0)} = 0, \quad 1 \le j \le m. \tag{17}$$

Noting the $r_{0,j} = R^+$ for all $j > 0$, we obtain the nontrivial solution

$$W^{(0)} = \frac{d}{R^+}, \tag{18}$$

or equivalently,

$$\sum_{j=1}^{m} \frac{y_{0,j}^{(0)} m!}{j!(m-j)!} = K(1 - d/(R^+)). \tag{19}$$

The left hand side of this equation is the total number of cheaters. At the zeroth order, it is equal to the effective carrying capacity of the system, $K(1 - d/(R^+))$.

At the first order, after some manipulations, we obtain the following system,

$$y_{0,j+1}^{(0)}(m-j) + y_{0,j-1}^{(0)} j = \lambda y_{0,j}^{(0)}, \quad 1 \le j \le m, \tag{20}$$

where we denoted

$$\lambda = 2m - \frac{R^+}{d} W^{(1)}. \tag{21}$$

This is an eigenvalue problem whose eigenvectors do not depend on any parameters of the system. In particular, below we list the eigenvalues for several small values of m, in the form $(\lambda_1, \ldots, \lambda_m)$, such that $\lambda_1 < \ldots < \lambda_m$:

$$\begin{aligned}
&m = 2, && (-\sqrt{2}, \sqrt{2}),\\
&m = 3, && (-\sqrt{7}, 0, \sqrt{7}),\\
&m = 4, && (-\sqrt{2(4+\sqrt{10})}, -\sqrt{2(4-\sqrt{10})}, \sqrt{2(4-\sqrt{10})}, \sqrt{2(4+\sqrt{10})}),\\
&m = 5, && (-\sqrt{15+2\sqrt{19}}, -\sqrt{15-2\sqrt{19}}, 0, \sqrt{15-2\sqrt{19}}, \sqrt{15+2\sqrt{19}})
\end{aligned}$$

The largest of these eigenvalues, λ_m, is the "ground state" of the system corresponding to the only eigenvector whose components are of the same sign. This is the relevant solution in our case. For several values of m, the eigenvector is given below, in the form $(v_1, \ldots, v_m)$:

$$\begin{aligned}
&m = 2, && \left(\frac{1}{\sqrt{2}}, 1\right),\\
&m = 3, && \left(\frac{2}{3}, \frac{\sqrt{7}}{3}, 1\right),\\
&m = 4, && \left(\frac{1}{2}\sqrt{\frac{1}{6}(8+\sqrt{10})}, \frac{1}{6}(2+\sqrt{10}), \frac{1}{2}\sqrt{2+\sqrt{\frac{5}{2}}}, 1\right),
\end{aligned}$$

$$m=5, \qquad \left(\frac{2}{15}(1+\sqrt{19}), \sqrt{\frac{94}{225}+\frac{7\sqrt{19}}{90}}, \frac{1}{10}(5+\sqrt{19}), \frac{1}{5}\sqrt{15+2\sqrt{19}}, 1\right).$$

We have

$$y_{0,j}^{(0)} = Av_j, \tag{22}$$

where A is a scaling factor, which can be determined from the normalization condition (19). The above eigenvectors together with condition (19) uniquely define the components $y_{0,j}^{(0)}$ for $1 \le j \le m$. In particular, we have

$$A = \frac{K(1-d/(R^+))}{\sum_{j=1}^{m} \frac{m! v_j}{j!(m-j)!}}.$$

The number of wild-type cells is obtained from the equation for $i=0, j=0$ in the first order of u,

$$y_{0,0}^{(1)} = \frac{y_{0,1}^{(0)} m R^+}{R^+ - R}. \tag{23}$$

Also, at the first order of system (9), we have the following equations,

$$y_{1,0}^{(1)} = \frac{R^+}{f} y_{0,1}^{(0)}, \quad y_{1,j}^{(1)} = \frac{R^+}{f}(y_{0,j}^{(0)} + y_{0,j+1}^{(0)}), \quad 1 \le j \le m-1, \tag{24}$$

which define the components of the cooperating types in terms of the cheating types. In particular, the total number of cooperators in the system is given by

$$\sum_{j=0}^{m-1} \frac{y_{1,j}^{(1)} m!}{j!(m-1-j)!} = \frac{uR^+}{f}\left(\sum_{j=1}^{m-1} \frac{y_{0,j}^{(0)} m!}{j!(m-1-j)!} + \sum_{j=0}^{m-1} \frac{y_{0,j+1}^{(0)} m!}{j!(m-1-j)!}\right) + O(u^2).$$

It it easy to show that

$$\frac{\sum_{j=1}^{m-1} \frac{v_j m!}{j!(m-1-j)!} + \sum_{j=0}^{m-1} \frac{v_{j+1} m!}{j!(m-1-j)!}}{\sum_{j=1}^{m} \frac{v_j m!}{j!(m-j)!}} = m.$$

Therefore, we have for the total number of cooperators,

$$\sum_{j=0}^{m-1} \frac{y_{1,j} m!}{j!(m-1-j)!} \approx \frac{uKmR^+}{f}\left(1 - \frac{d}{R^+}\right).$$

An important observable is the total number of cooperators that cooperate on one particular site, N_c,

$$\sum_{i=1}^{m}\sum_{j=0}^{m-i}\frac{y_{i,j}(m-1)!}{(i-1)!j!(m-i-j)!}.$$

In our approximation, this is given by $1/m$ times the total number of cooperators, which is approximated by

$$N_c=\sum_{j=0}^{m-1}\frac{y_{1,j}(m-1)!}{j!(m-1-j)!}\approx\frac{R^+}{f}uK\left(1-\frac{d}{R^+}\right). \tag{25}$$

Note that $\lambda_m/m<1$ for all m, and $\lim_{m\to\infty}\lambda_m/m=1$.

The approximation used here works well as long as the mutation rate is significantly smaller than the fitness differences,

$$u\ll\frac{f}{R^+}.$$

In this case, the selection-mutation balance is in place. When this inequality is reversed, expansion (11-14) breaks down. Figure 5 demonstrates how well the approximation (formula (25)) works by plotting the number of cooperators as a function of the parameter f. We can see that for small values of the quantity u/f, the relative error of approximation is small (the inset).

To get the second correction to the solution, we extend expansion (11-14) to the second order in u, and use equation (9) to write down the $O(u^2)$ terms. We have $m+1$ equations with $i=0$, m equations with $i=1$, and the additional equation (21), where $W^{(1)}$ is given by equation (16), and $\lambda=\lambda_m$. The total number of equations is $2m+2$. These are nonhomogeneous, nondegenerate equations.

The unknown variables to be determined at this order are:

$$\begin{aligned}&y_{00}^{(2)},\\&y_{01}^{(1)},\dots,y_{0m}^{(1)}\quad(m\text{ variables}),\\&y_{10}^{(2)},\dots,y_{1,m-1}^{(2)}\quad(m\text{ variables}),\\&W^{(2)},\end{aligned}$$

totaling $2m+2$ unknown variables. Note that variables $y_{01}^{(2)},\dots y_{0,m}^{(2)}$ do not enter the equations. There are also $m-1$ equations with $i=2$ which we did not count above. These equations contain $m-1$ additional unknowns, $y_{20}^{(2)},\dots,y_{2,m-2}^{(2)}$, which only appear in these equations and can be determined once the rest of the variables in this order are known.

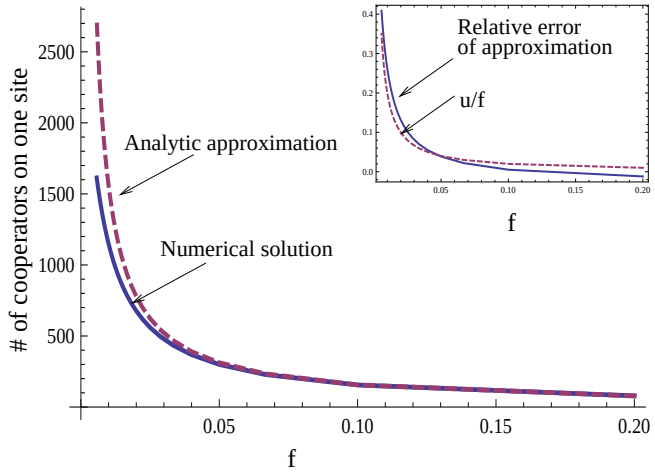

Figure 5: The number of cooperators (which cooperate on a particular site) as a function of parameter f: the numerical solution of the steady state of the ODEs (solid line) and the analytical estimate of formula (25) (dashed line). The inset: the relative error of this approximation and the value u/f as a function of f. Other parameters are: $m = 5$, $u = 2 \cdot 10^{-3}$, $K = 10^4$, $R^+ = 0.9$, $d = 0.1$.

3.2 Wild type cells participate in the division of labor

In this case, the wild type has fitness R^+ in the presence of cooperation, and instead of expansion (11) we have

$$y_{0,0} = y_{0,0}^{(0)} + u y_{0,0}^{(1)} + O(u^2),$$

that is, the wild type solution has a nonzero component in the absence of mutations. This makes the analysis simpler compared to the one presented in the previous section. The wild type now contributes to the zeroth order in the saturation term, such that instead of equation (15) we have

$$W^{(0)} = 1 - \frac{1}{K} \sum_{j=0}^{m} \frac{y_{0,j}^{(0)} m!}{j!(m-j)!},$$

and in equation (17) the indices now start from zero, $0 \leq j \leq m$. Equation (18) holds with the new definition of $W^{(0)}$, and the eigenvalue problem, equation (20), has indices $0 \leq j \leq m$. This fact simplifies the task of finding the eigenvalues and eigenvectors. Note that now, the matrix on the left of (20)

is stochastic (that is, its entries in each row add up to m). This matrix has the eigenvalue m corresponding to the only nonnegative eigenvector whose entries are all equal to one. Therefore, we have from (22)

$$y_{0,j}^{(0)} = A, \quad 0 \leq j \leq m, \tag{26}$$

and the normalization factor can be found explicitly from normalization (19) with the summation starting from zero. We have

$$A = \frac{K(1 - d/R^+)}{2^m}.$$

Equation (23) now does not apply as $y_{0,0}^{(0)}$ is now determined by (26). Expressions (24) now simplify to

$$y_{1,j}^{(1)} = \frac{R^+}{f}(y_{0,j}^{(0)} + y_{0,j+1}^{(0)}), \quad 0 \leq j \leq m-1,$$

and the total number of cooperators that cooperate on a particular site is again given by equation (25).

4 Numerical simulations

4.1 Patterns of cooperation and cheating

It was observed that cheaters can speed up the emergence of m-hit mutants compared to the sequential evolution in the absence of cooperation. Below we will describe all the patterns of cooperation and cheating that are observed in the simulations and can be explained by our theory. We assume that all mutation rates are equal to each other (or, more generally, the back-mutations are characterized by rates comparable to forward-mutations). We will refer to the case where the wild types do not participate in division of labor as the WT^- model, and the case where the wild types do participate in division of labor as the WT^+ model. The results depend on the size of the local steady state solution for the cooperators, N_c, approximated by equations (25). We start the simulation with wild-type cells only. These are the observed scenarios, depending on the mutation rates.

(i) For a relatively large mutation rate, the cooperators gradually appear and are maintained at a relatively high level, totaling approximately mN_c.

The partial cheaters typically have a much higher abundance. The wild type cells quickly disappear from the system in case WT^-, and they remain in the system in case WT^+. The full cheaters are produced, but disappear relatively quickly. This situation can persist for quite a long time. However if we look at the population of cooperators more closely, we observe interesting dynamics. If we track separately all the cooperators that cooperate on a given site, this population sometimes takes a "dip", such that no more product from this site is available for the cheaters to utilize. This leads to a temporary dip in the cheaters population. The lower the steady-state population of cooperators, the more often this happens. If it so happens that during such a "dip" full cheaters are generated, it is possible that they rise to a very high level, because their fitness in unaffected by the presence of cooperators. Once they rise, they can be maintained at a level higher than the level of partial cheaters. Interestingly, this is not because their steady-state level is higher (it is in fact lower than that of partial cheaters). They are winners of the evolutionary competition because the partial cheaters constantly take "hits" from the disappearance of one or another public goods products from the system. The partial cheaters simply cannot rise high enough (without crushing temporarily) to displace the full cheaters.

(ii) Next, we consider the scenario where the mutation rate is lower, and the steady-state level of cooperators, mN_c, is lower than in the previous case. Now, the public-goods products disappear from the system quite regularly, which leads to a significant decrease in the fitness of the partial cheaters (and the wild types in case WT^+). This happens because more than one, or even all m products can simultaneously disappear from the system. At such moments, the partial cheaters greatly decrease in levels. In the WT^- case, the partial cheaters give way to the wild-type, which temporarily invades the system. This is because in the absence of cooperation the fitness of wild-types exceeds that of partial cheaters. After a while, the cooperators, and thus the cheater populations build up again. Now the dynamics consist of periods where partial cheaters are abundant, and periods where the wild-type dominates. Depending on the mutation rate, the periods of wild-type domination can be longer (for lower mutation rates) or shorter. This proceeds until an m-hit mutant is generated which then replaces the partial cheaters (and the wild types) next time the cooperation breaks down, leading to its subsequent domination. For the WT^+ case, the dynamics remain similar as those described in the previous paragraph because the wild-type and the partial cheater populations are neutral towards each other (both of their fit-

nesses are influenced by cooperators, and neither contain costly cooperating sites). Therefore, the regular switching in their dominance is not observed.

It is clear that in the scenarios described in (i) and (ii), the m-hit mutant will rise faster than in the scenario with sequential evolution. The reason for this is the periods during which partial cheaters exist at elevated abundances, which increase the probability of generating an m-hit cheater.

(iii) When the mutation is even lower ($N_c < 1$), no cooperating population can get established and the dynamics are identical to those of sequential evolution.

(iv) When the mutation rate is very high, the cooperation never disappears from the system, and the m-hit mutants never come to dominate the system. In this case, the sequential dynamics are faster than those governed by cooperation and cheating.

4.2 The number of reproductive events

In this paper we report that over a wide range of parameters, cooperation and cheating accelerate evolution. This is for example demonstrated in figure 2 of the main text, where we present histograms of times it takes for an m-hit mutant to spread to 90% of the total population. It is clear that the amount of time it takes is significantly less in the cooperation/cheating scenario compared with the sequential evolution scenario.

One of the reasons why cooperation/cheating interactions accelerate evolution is the relatively higher number of reproductive events among partial mutants, which makes the rate of the m-hit mutant production faster. This is what is quantified in figure 4(a) of the main text. However, the difference in the numbers of reproductive events in the two scenarios is not the whole story.

The increase in the reproductive events of mutants is naturally coupled with the increase in the total reproductive events in the cooperation/cheating system. This comes from the fact that in the cooperation/cheating system, there is a large abundance of cooperating and cheating phenotypes, all of which have an elevated fitness (at least, most of the time), thus increasing the overall turnover in the system. Hypothetically it could happen that the acceleration of the evolutionary process observed in the cooperation/cheating scenario is simply due to the larger total number of reproductive events that occurs per unit time. Is this the only mechanism of acceleration that is at work?

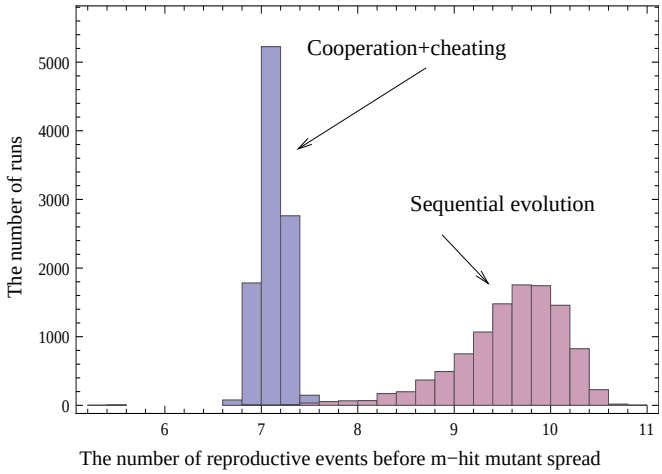

Figure 6: Distribution of times until the complex phenotype reaches 90% of the total population, based on repeated runs of the computer simulation. The simulations are the same as in figure 2(a) of the main text, with the same parameter values, except the time is measured in terms of the number of reproductive events.

To answer this question, we ran simulations in the sequential and cooperation/cheating scenarios, where instead of measuring time in terms of uniform units, we measured it in terms of reproductive events, see figure 6. This (generally, non-uniform) time-scale eliminates any advantage gained by a system with a larger abundance of fitter individuals (in our case, the cooperation/cheating system). We created histograms of m-hit mutant generation for the two scenarios, and we observed that despite the control over the number of reproductive events, the cooperation/cheating dynamics still give rise to an m-hit mutant at a significantly shorter time-scale. This highlights additional, less trivial aspects of the cooperation/cheating dynamics (apart from simply accelerating divisions). Partially-mutated individuals (cooperators and cheaters) exist at much higher levels compared to those of sequential evolution, thus making them much more likely to produce an m-hit mutant, and once produced, the m-hit mutant has an advantage over partially-mutated cheaters because it can resist the plunge of cooperators and survive the lean periods, when one or more gene products stop being supplied by cooperators.

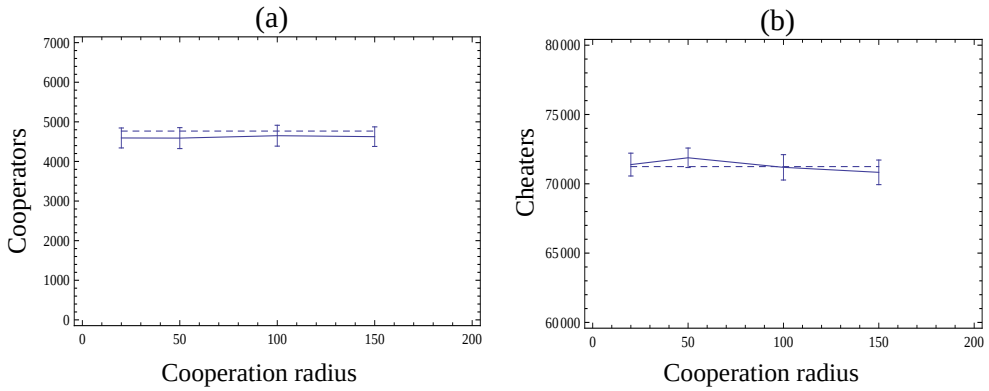

Figure 7: The mean and the standard deviation of the number of cooperators (a) and cheaters (b) in the cellular automaton simulations, as a function of the cooperation radius. The dashed lines correspond to the steady-state solution of the ODEs, equation (9). The parameters are $R = 0.33$, $R^+ = 0.8$, $R^- = 0.72$, $d = 0.1$, $f = 0.05$, $u = 7.92 \cdot 10^{-4}$, total grid size 300×300. The wild type participates in the division of labor. The fully-mutated cheaters are suppressed.

4.3 Parameter dependence of the results

Theory in developed in sections 2 and 3, although derived under the mass-action assumption, presents a surprisingly useful tool to study parameter dependence of the division of labor dynamics. Figure 7 shows that the steady-state levels of various species are independent of the cooperation radius. We ran time-evolution simulations, where we changed the cooperation radius. For example, if the cooperation radius is 50, this means that the cells cooperate in a 101×101 square. We suppressed the rise of m-hit cheaters by setting their reproduction rate to zero, and observed the steady-state levels of other species. The mean values and the standard deviations were then calculated to produce the plots in figure 7. The steady-state solutions of equations (9) are plotted with dashed lines for comparison. We can see that results from the mass-action theory can be used to study the behavior of the spatial systems.

Although the overall total steady-state levels of various species remain the same across different cooperation radii, the dynamics of division of labor systems change. In particular, the time it takes for the m-hit mutant to arise and dominate the system is a function of all the parameters, including the neighborhood size, K. For example, for a fixed number m, the cooperation

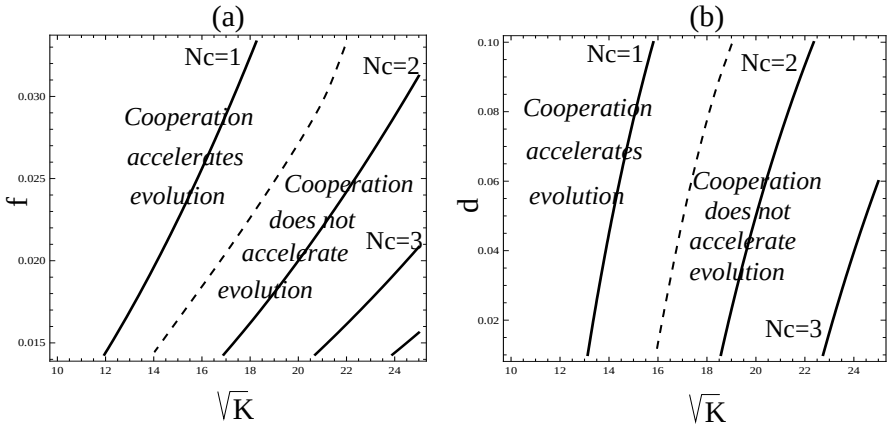

Figure 8: The contour-plots of the number of cooperators per site per neighborhood of size K, as given by equation (25), plotted as a function of different parameters. (a) Parameters f and K are varied, and $d = 0.1$. (b) Parameters d and K are varied, and $f = 1/70$. The other parameters are $R = 0.15$, $R^- = 0.135$, $R^+ = 0.3$, $u = 5 \cdot 10^{-4}$.

radius may be too small for the neighborhood K to simultaneously contain all m cooperating types. In this case, cooperation is not likely to arise, and the mutation accumulation will proceed by sequential evolution.

Other aspects of the parameter dependence can be seen from formula (25), where K has the meaning of the local neighborhood size. The quantity N_c tells us the steady-state number of cooperators per site, per neighborhood, before the m-hit mutants have arisen. If the value N_c is too large, then cooperators will not experience frequent local extinction, and the m-hit mutant will not have the necessary advantage to rise above the level of the partial cheaters.

In the last figure of the main text, we showed how the dependence of N_c given by formula (25) can serve as a predictor of the stochastic dynamics of the agent-based system. There, we explored the dependence of the results on parameters K and R^+. Figure 8 presents additional contour-plots of the quantity N_c as a function of parameters f, d, and K. We expect the cooperation/cheating dynamics to accelerate evolution to the left of the dashed lines on the graphs representing portions of the parameter space.

We also explored parameter dependencies that are beyond the applicability of formula (25). Figure 9 presents results of stochastic simulations. In Figure 9(a), we vary the total grid size of the simulation as well as the param-

eter R^-. We can see that the cooperator/cheater pathway can accelerate the emergence of the m-hit mutant if the total population size is sufficiently large (indicated by green). If the population size is too low, the rise of cheaters drives all populations extinct (indicated by red). The higher the value of R^-, the lower this threshold population size. Figure 9(b) assumes a higher mutation rate and corresponds to the regime where u/f is not small, such that formula (25) breaks down. In this case, the cooperator-cheater pathways can accelerate the emergence of the m-hit mutant if the population size lies below a threshold (indicated by green). If the population size is larger, the higher mutation rate prevents the extinction of individual public goods, and thus prevents the drop in fitness that is required for the m-hit mutant to emerge (indicated by purple). A more detailed analysis of the parameter dependence of the division of labor dynamics is subject of future work.

4.4 Alternative assumptions

We have explored the model in the context of a certain set of assumptions, some of which can be formulated differently. This section shows that alterations in these non-essential assumptions does not change the basic behavior of the model.

For simplicity, it was assumed that the cooperation and reproduction radii are identical. This need not be the case, and figure 10(a) shows that relaxing this assumption still leads to the accelerated evolution of the complex phenotype through the cooperator/cheater pathway. Figure 10(a) presents the result of stochastic simulation where two parameters were varied: the cooperation radius and the mutation rate. The replication radius was kept constant. The mutation rate was varied because it influences the level of cooperators within the cooperation radius, which is a crucial determinant of extinction events and thus of the outcome as explained above. Different colors in the graph indicate whether or not the cooperator/cheater pathway on average resulted in a faster emergence of the m-hit mutant than sequential evolution. As the cooperation radius is increased, the cooperator population within the radius rises, and this makes cheater-induced extinction of individual cooperator types less likely. The chances of cooperator extinction, however, are increased as the mutation rate is lowered. Therefore, the graph shows that for larger cooperation radii, an accelerated emergence of the m-hit mutant requires lower mutation rates. Note that for relatively small cooperation radii and small mutation rates, the cooperator/cheater pathway does

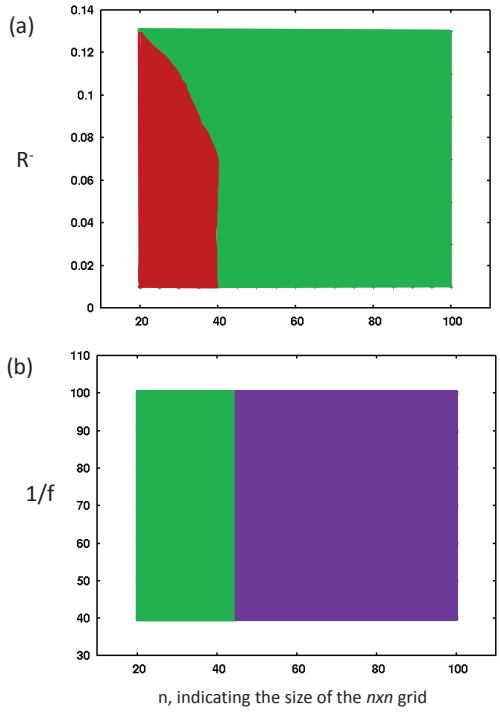

Figure 9: Parameter dependence of the cooperation dynamics vs sequential dynamics (numerical results from stochastic simulations). (a) The dependence on parameters n (the linear size of the simulation domain) and R^-. Red indicates the regime where the total population size is too small, and the population growth is not sustainable. Green indicates the regime where the division of labor dynamics leads to a faster production of an m-hit mutant compared to the sequential evolution scenario. Other parameters are $u = 3.17 \cdot 10^{-3}$, $f = 1/70$. (b) The dependence on parameters n and f (the cost of cooperation). Green is as in part (a), and blue indicates the regime where the division of labor dynamics do not led to a faster production of an m-hit mutant compared with the sequential evolution scenario. Other parameters are $R^- = 0.135$, $u = 7.96 \cdot 10^{-3}$. For both parts (a) and (b) we have $m = 5$, $K = 10 \times 10$, $R = 0.15$, $R^+ = 0.5$, $d = 0.1$.

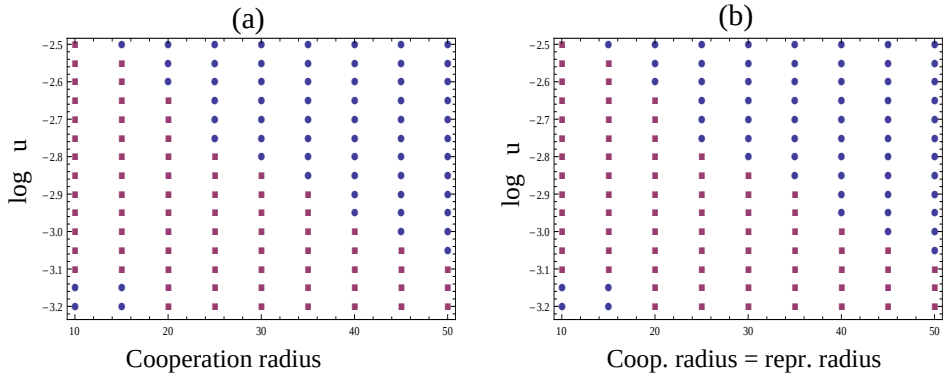

Figure 10: Altering the cooperation and reproduction radii in the model where wild-types do not benefit from shared goods. (a) Varied are the cooperation radius and the mutation rate, u. The reproduction radius is kept constant. Red squares correspond to cases where the m-hit mutant is created significantly faster in the division-of-labor scenario compared to sequential evolution. Blue circles correspond to the cases where division of labor does not accelerate the emergence of m-hit mutants compared to sequential evolution. For each point, 10 simulations were run both under the division of labor and sequential evolution scenarios and the results were averaged. (b) Same type of simulation. Instead of varying only the cooperation radius, both the reproduction and the cooperation radius were varied, keeping the size of the two radii equal. The same picture emerges, showing that varying the reproduction radius does not influence the qualitative outcome of the evolutionary dynamics. Parameters were chosen as follows: the grid size is 100×100; $m = 5$; $R = 0.15$; $R^+ = 0.5$; $R^- = 0.135$; $D = 0.1$; $f = 1/70$. In the grid size under consideration, a radius of 50 corresponds to mass-action.

not accelerate the emergence of the m-hit mutant. The reason is that in this parameter regime, not enough cooperating types are generated by mutations within the cooperation radius, such that the division of labor dynamics cannot emerge.

Further, simulations indicate that keeping the cooperation radius constant while varying the replication radius does not change the outcome of the evolutionary dynamics on a qualitative level, because it does not alter the level of cooperators in local neighborhoods. For comparison, we also varied both radii together from relatively small values to mass action and plotted the outcome in dependence of the mutation rate, figure 10(b). The outcome is the same as varying only the cooperation radius. This figure also shows that the validity of the evolutionary dynamics are not restricted to the

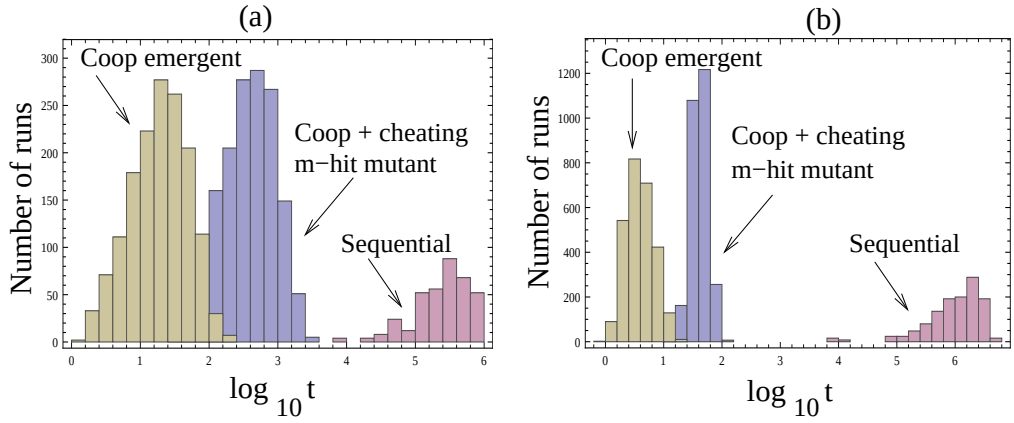

Figure 11: Altering assumptions about the fitness of individuals. As in figure 2(a) of the main text, simulations of three types were run for each of the pannels (a,b): (1) sequential evolution scenario, stopping at the m-hit mutant invasion (marked "sequentisl"), (2) cooperation scenario in the absence of cheaters, stopping at the invasion of a heterogeneous couud of cooperators (marked "coop emergent"), and (3) cooperation scenario in the presence of cheaters, stopping at the m-hit mutant invasion (marked "coop+cheating m-hit mutant"). (a) This simulation assumes that cooperators are characterized by a fitness cost f, cheaters by a smaller fitness cost f_2, and wild-type by the absence of a fitness cost. Compared to figure 2(a) of the main text, a lower mutation rate has been chosen to make sure that the level of cooperators remains low enough such that individual cooperators go extinct frequently. (b) This simulation, also based on figure 2(a) in the main text, assumes that each intermediate mutant carries an independent fitness cost of 10%. Hence, the more intermediate mutations are accumulated, the higher the fitness cost of the mutant. Parameters were chosen as follows: grid size is 100×100; $m = 5$; $R = 0.15$; $R^+ = 0.5$; $R^- = 0.135$; $D = 0.1$; $f = 1/70$; $f_2 = 1/140$; $u = 3.17 \times 10^{-3}$. Cooperation and replication radii are both 10.

relatively small radii observed in the paper, but hold for all radii.

Another assumption we made was that cooperators carry a fitness cost while other populations do not. The rational behind this assumption is that cooperators produce larger amounts of a product, sufficient to be shared with others. However, one can alternatively assume that cheaters also carry a fitness cost, although a smaller one, because they also produce the product, but in lesser quantities. Hence, we assumed that cheaters carry half the fitness cost of cooperators, while wild-types do not carry a fitness cost. Because the fitness of the cheaters is lowered under this assumption, the level of cooperators becomes higher, making the extinction of individual cooperator types,

and thus the emergence of the m-hit mutant, less likely. To increase the chance that individual cooperator types go extinct, leading to the rise of the m-hit mutant, the mutation rate needs to be lowered. Thus, we re-ran the simulation shown in figure 3(a) of the main text with a lower mutation rate, and the outcome is qualitatively the same, as shown in figure 11(a).

Finally, we explored alterations in the assumption about the fitness cost of intermediate mutants. In our simulations, all intermediate mutants carried the same fitness cost, regardless of how many mutations an agent had accumulated. Instead, we now assume that each new mutation adds a fitness cost, until an agent has accumulated m mutations and attains an advantage compared to the wild-type. We re-ran the simulation of Figure 3(a) of the main text to demonstrate that the evolutionary dynamics remain qualitatively identical, see figure 11(b).

5 Application to the in vivo evolution of HIV

In this section, we apply our model to the in vivo evolution of human immunodeficiency virus (HIV) infection. In particular, we consider the emergence of mutants that escape cytotoxic T lymphocyte (CTL) responses.

Following infection, HIV replicates to high levels in the acute phase, and virus load is eventually suppressed to lower levels, which marks the beginning of the asymptomatic, or chronic, phase of the infection. The reduction in virus load is at least in part due to anti-viral CTL responses that can act both lytically by killing infected cells and non-lytically by secreting factors that inhibit viral replication. Over time, the immune system loses control of the infection, leading to the development of AIDS. Viral evolution in vivo, and escape from CTL in particular (Ganusov *et al.*, 2011) is thought to contribute to the progression of the disease. In the presence of a narrow CTL response (i.e. where just one immunodominant response largely fights the virus), escape can readily occur, given that it can be achieved by simple point mutations. In the presence of broader CTL responses, however, escape is more difficult to achieve (Ganusov *et al.*, 2011). It has been argued that a sufficiently broad CTL response could prevent viral escape (Ganusov *et al.*, 2011). Yet, viral escape from multiple CTL responses has been observed in patients that suppress virus relatively well (Crawford *et al.*, 2009). Here, we show that while escape from a relatively broad CTL response can indeed be

an unlikely event if the virus accumulates mutations sequentially, it can occur within a biologically realistic time frame in the context of the evolutionary dynamics described in this paper.

A central component in these evolutionary dynamics is the presence of cooperators and cheaters. Escape typically occurs if a mutation leads to loss of binding to MHC class I molecules, or if a mutation destroys the ability of the T cell receptor to bind its epitope (McMichael & Phillips, 1997). However, mutations can also alter the binding properties between the peptide and the T cell receptor such that besides failed lysis of the mutant-infected cell, activity against the wild-type is also abolished. This is called T cell antagonism and has been demonstrated to occur in HIV infection (Klenerman *et al.*, 1994; Meier *et al.*, 1995). Antagonistic mutants are cooperative in nature because they alleviate the wild-type virus from CTL-mediated activity (Davenport, 1995). In contrast, escape mutants that merely prevent killing of mutant-infected cells, but do not benefit other types can be considered cheaters in this context. We will refer to the cooperators as antagonistic mutants, and the cheaters as escape mutants.

We have adapted our model to describe the escape of HIV from multiple CTL responses directed against different epitopes of the virus, concentrating on scenarios in which virus load is relatively low and the susceptible target cells are abundant. This corresponds to relatively strong immunological control of the virus, a situation in which escape from a broad CTL response is likely to be most challenging. Because the number of uninfected, susceptible target cells is assumed to be not limiting, this population of cells is not explicitly taken into account in our simulations. The model tracks the number of infected cells. Consider wild-type viruses first. With a probability R, the infected cell gives rise to a newly infected cell, a process that includes viral replication in the source cell as well as the transmission of offspring virus to a target cell. With probability D, an infected cell dies. This probability includes both virus-induced cell death, and CTL-induced cell death. The CTL population is not modeled explicitly, but expressed as a constant which adds to the death rate of infected cells. Next, consider escape mutants. They are known to carry a fitness cost relative to the wild-type virus (Crawford *et al.*, 2009; Ganusov & De Boer, 2006). Although compensatory mutations may occur that lessen this cost, the escape mutant is still thought to be less fit than the wild-type. Hence, we assume that escape

mutants are characterized by a replication probability R^-, where $R^- < R$. If a virus has escaped all CTL responses under consideration, the infected cell has a reduced death probability, D^-. In addition, it is also assumed to have a raised replication probability because of reduced non-lytic CTL-mediated activity, which is denoted by R^{-*}. Finally, consider antagonistic mutants. They are assumed to carry the same replicative cost as escape mutants, and their replication probability is thus given by R^-. If a virus antagonizes all CTL responses under consideration, the appropriate infected cell is characterized by the lower death probability D^- and by the replication probability R^{-*}. The same probabilities also apply to any infected cell if an antagonist of each CTL response is present within its cooperation radius. In contrast to escape mutants, antagonists are assumed to carry an additional cost, f, for cooperating, which is subtracted from their replication probability. As with the model presented in the main text, we also assume a replication radius, defining the area in which a target cell has to be located to be infected by a source cell. Mutations occur with a probability u, and follow the scheme outlined in Figure 1 in the main text.

Many parameters are known for HIV infection, and these parameter values were used in simulations. We take the time between updates in the simulation to be 0.2 days. The mutation rate is assumed to be 3×10^{-5} per base pair per generation (Mansky & Temin, 1995). It is unclear how many nucleotide substitutions can lead to viable escape within an epitope. It is likely to be a small number that varies from epitope to epitope, and we have assumed it to be ten in the model, although results do not depend on this assumption. Back mutations occur through one specific mutation. The death rate of (productively) infected cells is around $0.5 day^{-1}$, which translates into $D = 0.1$ (Perelson *et al.*, 1996). The basic reproductive ratio of HIV has been estimated to be around $R_0 = 8$ (Little *et al.*, 1999; Ribeiro *et al.*, 2010). For this to be true in our model, the replication probability must be $R = 0.8$. These parameters are assumed to apply to the wild-type. For mutants, the following additional considerations apply. While the exact cost of escape has not been quantified (Ganusov & De Boer, 2006), compensatory mutations have been documented which, however, still leave the mutant with a disadvantage compared to the wild-type. We assume a 10% fitness cost of the escape mutant such that R^-=0.72. The contribution of CTL killing to the infected cell death rate has been examined (Asquith *et al.*, 2006; Wick *et al.*, 2005). While exact estimates vary, it was consistently found to be relatively

low. We assumed that 20% of infected cell death is attributable to CTL (including all clones). Hence, $D^- = 0.08$. This death rate applies to all infected cells that are either infected with a full escape mutant, or escape all CTL clones through cooperation. Such infected cells are also characterized by an increased replication rate due to a reduction in CTL-mediated non-lytic activity. The extent of CTL-mediated non-lytic reduction of viral replication is unclear, so a modest 10% is assumed, such that $R^{-*} = 0.8$. It is currently not known whether cooperation through antagonism carries a cooperation cost. Since cooperative interactions often do involve costs, we assumed a relatively small cost of $f = 0.02$. The total number of infected cells varies greatly among patients. In some patients, the total count of infected cells has been estimated to be around 10^7-10^8 (Chun *et al.*, 1997). Because we are interested in escape in the context of relatively efficient virus control, we assume a total infected cell count of the order of about 10^5. Regarding the cooperation and replication radii in the model, the following was assumed. Perfect mixing or mass action was implemented for viral replication. Although lymph nodes have some spatial structure, cells do mix relatively well and there is so far no evidence for spatial effects in HIV dynamics. The cooperation radius was taken to be 10. That is, a cooperator can only benefit a subset, and not all infected cells in the body, which is a reasonable assumption.

Figure 12 shows the results of the simulations, assuming that a virus needs to escape 4 different CTL clones. While a mutant that escapes all CTL responses comes up in about 100 years in the sequential evolution scenario, it arises within about 5 years through the cooperator/cheater pathway. The use of measured parameters therefore demonstrates that the cooperator/cheater pathway can significantly speed up the emergence of escape from broad CTL responses in HIV infection. While some parameters are known precisely, others could only be estimated roughly. Our model allows us to specify ranges of these parameters for which the cooperator-cheater dynamics are expected to speed up the emergence of a full escape mutant.

An interesting observation is that while escape mutants are quite readily observed in acute HIV infection, they are much less likely to emerge during the chronic phase (Ganusov *et al.*, 2011). In acute infection, virus load can rise to very high levels, enhancing mutant generation, and initial immune responses can be more narrow. Lower virus loads during chronic infection and possibly the presence of broader responses makes it more difficult for

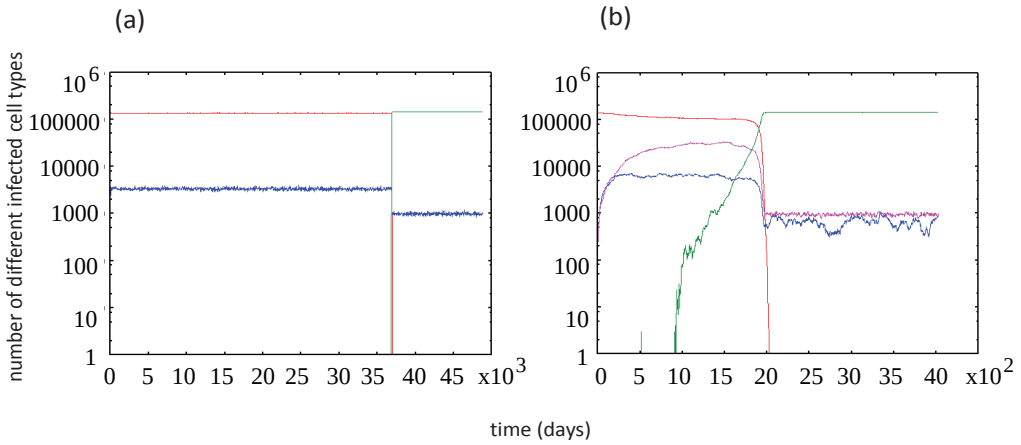

Figure 12: Application to CTL escape in HIV infection, based on an adapted version of the model that has been parameterized using estimates available in the literature. Cooperator mutants are T cell antagonists, while cheater mutants are escape strains that only benefit themselves. We consider the evolution of mutants that simultaneously escape 4 CTL clones. (a) In the sequential evolution scenario, it takes approximately 100 years for the complete escape mutant to emerge. Wild-type-infected cells are shown in red, cells infected with partial mutants in blue, and cells infected with complete escape mutants in green. (b) With the cooperator/cheater pathway, the complete escape mutant arises within 5 years. Wild-type infected cells are shown in red, cells infected with partial mutant viruses that cooperate at least in one site are shown in blue, cells infected with partial mutant viruses that cheat in all sites are shown in pink, and cells infected with complete escape mutant viruses are shown in green. The plots are based on individual realizations of the simulation, and are representative examples of the behavior of the stochastic model under the two assumptions.

escape mutants to arise. It is intriguing to hypothesize that the rise of escape mutants during chronic infection requires special evolutionary pathways, as described in this paper.

While it is interesting to adapt the model to HIV infection and to show that the evolutionary dynamics described in this paper could be applicable in the context of measured parameters, it is important to note that this is a simplified model that ignores several aspects of HIV infection. In the evolutionary context, the process of recombination is not included in the model. HIV has a diploid genome, and if a cell is simultaneously infected with different strains, two different genomes can be packaged in the same viral particle. Upon infection of a new target cell, recombination can occur during the process of reverse transcription (Jung *et al.*, 2002; Levy *et al.*, 2004), allowing different mutations to be brought together in a single viral genome. Mathematical models have shown that recombination can slow down or accelerate multi-resistant mutants in HIV infection, or have no effect, depending on the particular scenarios and assumptions considered (REF). In the context of the dynamics examined here, simulations have shown that recombination does not accelerate the emergence of the m-hit mutant in the sequential evolution scenario. This would require that different partial mutants infect the same cell. The partial escape mutants, however, are characterized by reduced fitness relative to the wild-type and are therefore only maintained at low levels at a selection-mutation balance. This makes it extremely unlikely that a cell is infected with different partial escape mutants, and recombination is not likely to influence the evolutionary dynamics in this context.

Also note that the literature suggests different mechanisms of CTL antagonism. One study (Sewell *et al.*, 1997) has argued that antagonism requires the wild-type and the mutant epitopes to be presented on the same infected cell, a scenario that is different from the one assumed in our model. Another study (Meier *et al.*, 1995), however, clearly shows that wild-type virus can benefit from the antagonistic mutant even if it is presented on a different target cell, confirming the assumption made in our model.

References

Asquith, B., Edwards, C., Lipsitch, M. & McLean, A. 2006. Inefficient cytotoxic t lymphocyte–mediated killing of hiv-1–infected cells in vivo. PLoS

biology, 4 (4), e90.

Chun, T., Carruth, L., Finzi, D., Shen, X., DiGiuseppe, J., Taylor, H., Hermankova, M., Chadwick, K., Margolick, J., Quinn, T. *et al.* 1997. Quantification of latent tissue reservoirs and total body viral load in hiv-1 infection. Nature, 387 (6629), 183–188.

Crawford, H., Lumm, W., Leslie, A., Schaefer, M., Boeras, D., Prado, J., Tang, J., Farmer, P., Ndung'u, T., Lakhi, S. *et al.* 2009. Evolution of hla-b* 5703 hiv-1 escape mutations in hla-b* 5703–positive individuals and their transmission recipients. The Journal of experimental medicine, 206 (4), 909.

Davenport, M. 1995. Antagonists or altruists: do viral mutants modulate t-cell responses? Immunology today, 16 (9), 432–436.

Ganusov, V. & De Boer, R. 2006. Estimating costs and benefits of ctl escape mutations in siv/hiv infection. PLoS Computational Biology, 2 (3), e24.

Ganusov, V., Goonetilleke, N., Liu, M., Ferrari, G., Shaw, G., McMichael, A., Borrow, P., Korber, B. & Perelson, A. 2011. Fitness costs and diversity of the cytotoxic t lymphocyte (ctl) response determine the rate of ctl escape during acute and chronic phases of hiv infection. Journal of virology, 85 (20), 10518–10528.

Jung, A., Maier, R., Vartanian, J., Bocharov, G., Jung, V., Fischer, U., Meese, E., Wain-Hobson, S. & Meyerhans, A. 2002. Recombination: multiply infected spleen cells in hiv patients. Nature, 418 (6894), 144–144.

Klenerman, P., Rowland-Jones, S., McAdam, S., Edwards, J., Daenke, S., Lalloo, D., Köppe, B., Rosenberg, W., Boyd, D., Edwards, A. *et al.* 1994. Cytotoxic t-cell activity antagonized by naturally occurring hiv-1 gag variants. Nature, 369, 403–407.

Levy, D., Aldrovandi, G., Kutsch, O. & Shaw, G. 2004. Dynamics of hiv-1 recombination in its natural target cells. Proceedings of the National Academy of Sciences of the United States of America, 101 (12), 4204.

Little, S., McLean, A., Spina, C., Richman, D. & Havlir, D. 1999. Viral dynamics of acute hiv-1 infection. The Journal of experimental medicine, 190 (6), 841–850.

Mansky, L. & Temin, H. 1995. Lower in vivo mutation rate of human immunodeficiency virus type 1 than that predicted from the fidelity of purified reverse transcriptase. Journal of Virology, 69 (8), 5087–5094.

McMichael, A. & Phillips, R. 1997. Escape of human immunodeficiency virus from immune control. Annual review of immunology, 15 (1), 271–296.

Meier, U., Klenerman, P., Griffin, P., James, W., Köppe, B., Larder, B., McMichael, A. & Phillips, R. 1995. Cytotoxic t lymphocyte lysis inhibited by viable hiv mutants. Science, 270 (5240), 1360.

Perelson, A., Neumann, A., Markowitz, M., Leonard, J. & Ho, D. 1996. Hiv-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. Science, 271 (5255), 1582–1586.

Ribeiro, R., Qin, L., Chavez, L., Li, D., Self, S. & Perelson, A. 2010. Estimation of the initial viral growth rate and basic reproductive number during acute hiv-1 infection. Journal of Virology, 84 (12), 6096–6102.

Sewell, A., Harcourt, G., Goulder, P., Price, D. & Phillips, R. 1997. Antagonism of cytotoxic t lymphocyte-mediated lysis by natural hiv-1 altered peptide ligands requires simultaneous presentation of agonist and antagonist peptides. European journal of immunology, 27 (9), 2323–2329.

Wick, W., Yang, O., Corey, L. & Self, S. 2005. How many human immunodeficiency virus type 1-infected target cells can a cytotoxic t-lymphocyte kill? Journal of virology, 79 (21), 13579–13586.