

Overcoming toxicity: why boom-and-bust cycles are good for non-antagonistic microbes

Ming Yi Wang^a, Alexander Vladimirs^{b,1}, and Andrea Giometto^{c,1}

^aCenter for Applied Mathematics, Cornell University, Ithaca, NY 14853

^bDepartment of Mathematics and Center for Applied Mathematics, Cornell University, Ithaca, NY 14853

^cSchool of Civil and Environmental Engineering, Cornell University, Ithaca, NY 14853

Abstract

Antagonistic interactions are critical determinants of microbial community stability and composition, offering host benefits such as pathogen protection and providing avenues for antimicrobial control. While the ability to eliminate competitors confers an advantage to antagonistic microbes, it often incurs a fitness cost. Consequently, many microbes only produce toxins or engage in antagonistic behavior in response to specific cues like quorum sensing molecules or environmental stress. In laboratory settings, antagonistic microbes typically dominate over sensitive ones, raising the question of why both antagonistic and non-antagonistic microbes are found in natural environments and host microbiomes. Here, using both theoretical models and experiments with killer strains of *Saccharomyces cerevisiae*, we show that boom-and-bust dynamics caused by temporal environmental fluctuations can favor non-antagonistic microbes that do not incur the growth rate cost of toxin production. Additionally, using control theory, we derive bounds on the competitive performance and identify optimal regulatory toxin-production strategies in various boom-and-bust environments where population dilutions occur either deterministically or stochastically over time. Our findings offer a new perspective on how both antagonistic and non-antagonistic microbes can thrive under varying environmental conditions.

Keywords: Quorum sensing | Interference competition | Microbiome | Piecewise-deterministic processes | Stochastic optimal control

Introduction

Antagonistic interactions are found throughout the microbial tree of life [1–5], in almost any environment [6–10] and host-associated microbiomes [11–14]. Microbes have evolved a large variety of mechanisms to interact antagonistically with each other [1, 15], from contact dependent antagonism (e.g., via type IV, V, and VI secretion systems), to short-distance interaction mediated by diffusible antimicrobial metabolites (e.g., bacteriocins), to long-range interaction via secretion of volatile antimicrobials [16]. These antagonistic interactions are thought to be a strong determinant of microbial community structure [1, 17], to provide benefits to hosts, such as protection against pathogen invasion [13], and to be a promising avenue for antimicrobial control in both natural ecosystems and animal hosts [13, 18, 19].

¹A.V. (Author Two) and A.G. (Author Three) contributed equally to this work. To whom correspondence should be addressed.

E-mail: vladimirs@cornell.edu and giometto@cornell.edu.

Recent experimental investigations have explored the ecological [20–22] and evolutionary dynamics of microbial antagonism [23] in the laboratory, where toxin-producing strains are found to dominate over toxin-sensitive ones. In many microbiomes, however, one finds both antagonistic and non-antagonistic microbes [24], raising the question of whether the dominance of toxin-producing strains observed in the laboratory might be caused by idealized growth conditions in those settings. Here, we explore how the interplay of environmental fluctuations, costs associated with toxin production, and regulation of toxin production, play different roles in determining the outcome of the competition. Although spatial heterogeneity of populations is one of the popular explanations of why toxin-sensitive strains might be doing well in practice [4, 25–27], we show that environmental fluctuations (e.g., dilutions) can favor toxin-sensitive strains even in spatially homogeneous populations.

Ecological theory has long recognized the significant impact of environmental fluctuations on community composition [28–31]. Changes in temperature, nutrient levels, and other abiotic factors critically shape the structure and dynamics of these communities. Such disturbances not only disrupt resident populations, possibly allowing new colonizers to invade, but also affect both immediate and long-term ecological outcomes [32, 33]. Microbial communities in a diverse array of habitats are subject to “boom-and-bust” dynamics, where periods of rapid population growth are often followed by sharp declines. Such dynamics have been observed across various environments, including phytoplankton and particle-attached microbial communities in marine ecosystems [34–36], soil [37–40], host-associated microbiomes [41–44], and the built environment [45]. These boom-and-bust cycles are not only caused by abiotic factors such as nutrient availability and environmental disturbances, but also by biotic interactions including competition, predation, and parasitism. Particularly, the interactions with phages and predators can drastically alter microbial community structure, trigger population crashes, and thereby influence the overall dynamics of microbial ecosystems [46–48].

Despite the interest in the effect of microbial antagonism and environmental fluctuations on microbial community composition [33, 49, 50], the interplay between environmental fluctuations and antagonistic microbial interactions is surprisingly underexplored. Although data relating the relative abundance of toxin-producing strains to environmental fluctuations is very scarce, there is some evidence that environments with higher turnover rates harbor a reduced number of toxin-producing strains [51, 52]. Mathematical modeling and laboratory experiments in stationary environments suggest that the efficacy of toxin-mediated killing is dependent on high population

densities [21, 53] and consequently frequent population busts could favor sensitive strains that avoid the metabolic costs associated with toxin production. Certain microbial species have evolved regulatory mechanisms for toxin production that are triggered by quorum sensing signals [53, 54] or environmental indicators of stationary phase [55]. This regulation ensures that resources are not expended on toxin production at times when it would be least effective, possibly allowing microbes to optimize the cost-benefit ratio of toxin production. Activation of toxin production genes in response to quorum sensing signals produced by other strains, a phenomenon known as eavesdropping or cross-talk, has also been reported [56]. Regulation of toxin production in response to self and non-self abundances is thus theoretically possible and may be exploited to design synthetic genetic systems for pathogen eradication via targeted secretion of antimicrobials [57].

Our paper seeks to investigate how boom-and-bust cycles and toxin regulation affect the antagonistic competition dynamics between toxin-producing and non-producing microbes. This approach aims to elucidate the survival strategies of microbial populations and provide a deeper understanding of the ecological and evolutionary consequences of microbial interactions under variable environmental conditions.

Experimental antagonism with periodic dilutions.

To gain intuition for the impact of environmental disturbances on the dynamics of microbial antagonism, we conducted competition experiments between a sensitive (S) and a killer (K) strain of *Saccharomyces cerevisiae*, with the latter engineered to secrete the killer toxin K1 expressed from the galactose-inducible promoter *P_{GAL1}* [21]. The toxin K1 kills S cells by disrupting ion balance and membrane potential [58], but does not affect K cells that are immune to it [59].

In isolation and at low density, strain S grew at a faster growth rate ($r_S = 0.28 \pm 0.01 \text{ h}^{-1}$, mean \pm SD) than strain K ($r_S = 0.26 \pm 0.03 \text{ h}^{-1}$). Making these two strains compete against each other in well-mixed liquid cultures, we found that, for the lowest initial population sizes tested, the fraction of K cells in the population initially decreased and then increased (the lightest curves in Fig. 1). Increasing

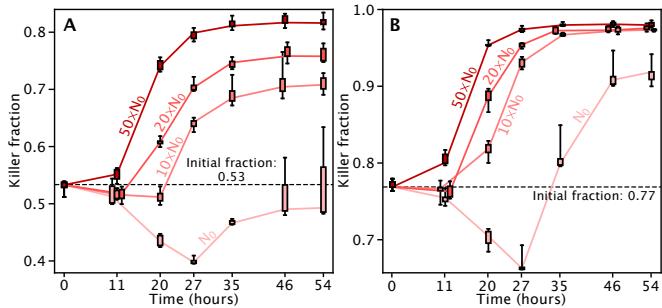


Figure 1. Uninterruption competitions between the killer and sensitive strains in experiments initialized with different population sizes (different shades, $N_0 = 5.8 \cdot 10^4 \text{ cells/mL}$) and killer fractions f_0 (53% in A and 77% in B). With small initial populations, the fraction of killers first decreases, due to the sensitive strain's higher growth rate, and later increases, due to the toxin's action. Lines connect subsequent median fractions. Some data points were spaced horizontally to ease visualization. Quartiles were computed via bootstrapping, whiskers report the minimum and maximum values. Could the toxin-sensitive microbes become dominant in the long term if recurrent environmental disturbances dilute the population when they are in the majority?

the initial population size or the initial fraction of K cells reduced the time window over which their fraction decreased. These dynamics are consistent with the intuition that at low cell densities, when the toxin is too dilute to significantly impact S cells, their relative abundance should temporarily increase due to their higher growth rate.

These observations led us to hypothesize that frequent and fortuitously timed disturbance events, such as dilutions, could favor the sensitive strain allowing it to become dominant over time, despite the presence of the toxin-producing killer strain, which would normally dominate without such dilutions. To test this hypothesis, we made the two strains compete against each other by periodically diluting them with different dilution intervals (Fig. 2). Unlike in the experiment of Fig. 1, here nutrients were replenished at each dilution, allowing the population to re-grow (Fig. S1). In agreement with our hypothesis, we found that both the initial fraction and the interval between successive dilutions influenced the competition outcome between S and K, favoring S when dilutions occurred more frequently (Fig. 2). While some heterogeneity of competition outcomes was observed with the 48 h cycle, the overall trend was clear: when the initial K fraction was low enough, S ultimately dominated, with the frequency of dilutions influencing only the rate at which S became dominant; when the initial K fraction was much higher, the same was true about the K domination; but for intermediate initial K fractions, the frequency of dilutions became a crucial predictor of the ultimate winner.

Overall, these experiments suggest that frequent dilutions can favor sensitive cells when their growth rates in isolation are larger than those of the killers. However, several factors may limit the general applicability of these results. For instance, many antagonistic mi-

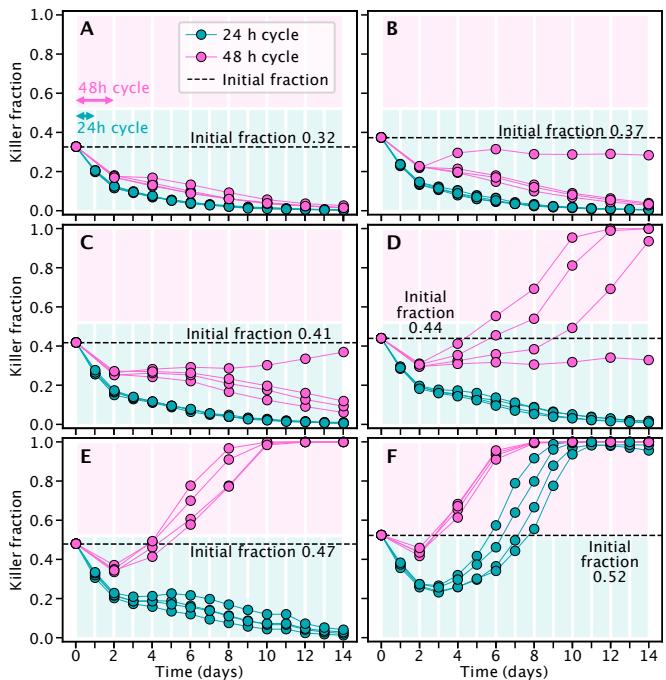


Figure 2. Experimental competitions between a toxin-producing (killer) and a sensitive strain of *Saccharomyces cerevisiae* in environments diluted periodically with periods of $T = 1 \text{ day}$ (teal) and $T = 2 \text{ days}$ (pink). Different panels show different initial fractions of the killer strain. Both the initial killer fraction and the period of the dilution cycles determine the outcome of the competition and the rate of extinction of the losing strain, with longer inter-dilution times favoring the killer strain. Shaded pink and teal rectangles display uninterrupted periods of growth in the two treatments. White vertical lines depict dilution events.

crobes regulate toxin production through quorum sensing or in response to environmental cues, rather than producing toxins constitutively. Additionally, nutrient turnover in more realistic environments could enable toxin-producers to completely eliminate toxin-sensitive microbes between population busts, unlike in our batch culture experiments where the relative proportions of K and S stabilize once nutrients are depleted (Fig. 1). To more comprehensively characterize how growth rates, toxin regulation, and environmental fluctuations jointly affect the dynamics of microbial antagonism, we turn to mathematical modeling.

Population dynamics.

We start by describing a basic competition model between a toxin-producing killer (K) and a toxin-sensitive strain (S). The goal is to keep the modeled mechanisms general, although our main ideas can be similarly applied to more complex models tied to specific microorganisms, experimental conditions, or environments.

We assume that the growth of both strains is logistic, with respective intrinsic growth rates (r_K, r_S) and a shared carrying capacity¹ C . We assume that the killer's *ability* to produce the toxin confers to it a growth rate deficit (i.e., $r_K < r_S$) independent of the actual toxin production. This deficit could arise from costs related to the regulation of toxin production, such as quorum sensing [60], but may also reflect other metabolic or genotypic differences that are unrelated to toxin production. Additionally, we assume an extra growth rate deficit that scales linearly with the normalized toxin production rate $a \in [0, 1]$, yielding the realized growth rate $r_K(1 - \varepsilon a)$, where $\varepsilon > 0$ is the cost of producing the toxin at the maximal rate. This deficit accounts for the metabolic cost of toxin production [53] and other expenses, such as toxin secretion, which may require cell lysis [3]. The rate of toxin-induced death of the sensitive strain is assumed to be proportional to the product of the strain densities (n_K and n_S , respectively) with the killing rate μ . After non-dimensionalizing, $n_K(t) \rightarrow n_K(t)/C$, $n_S(t) \rightarrow n_S(t)/C$, and $t \rightarrow r_{st}$, the resulting dynamics are

$$\begin{cases} \frac{dn_K}{dt}(t) = r_{KS}(1 - \varepsilon a)(1 - n_K - n_S)n_K, \\ \frac{dn_S}{dt}(t) = (1 - n_K - n_S)n_S - a\gamma n_K n_S, \end{cases} \quad (1)$$

where $r_{KS} = r_K/r_S$ is the intrinsic growth rates ratio and $\gamma = \mu C/r_S$ is the rescaled killing rate. In our numerical experiments, we set $\gamma = 1$ so that the strength of antagonism (the toxin-induced death term $-\gamma n_K n_S$ with the maximum rate of toxin-production $a = 1$) is comparable to the inter-strain competition for nutrients (the quadratic term $-n_K n_S$ in the logistic). For the sake of consistency, all rates in the rest of this paper are dimensionless (i.e., scaled by r_S).

Distinguishing between the relative values of r_{KS} and ε in published datasets is challenging, as both parameters influence the realized growth rate of the killer strain. A review of various studies [61–63] that compared the growth rates of killer and sensitive strains suggests that the ratio of realized growth rates $r_{KS}(1 - \varepsilon)$ typically falls between 0.68 and 0.98 for strains that produce the toxin at a constant, maximal rate (referred to as *constitutive killers*). Since proteinaceous toxins are often expressed from plasmids, one can estimate characteristic r_{KS} values by comparing the growth rates of *Escherichia coli* strains harboring such plasmids, with the toxin-producing genes deleted, to strains lacking these plasmids. From [64], we derive that

¹Unlike in the experiments reported in Fig. 1 (where nutrients were added at the start and never replenished) and in Fig. 2 (where nutrients were replenished at each dilution), the fixed carrying capacity used in Eq. (1) reflects a more realistic assumption of continuous nutrient supply.

$r_{KS} \approx 0.85$ is a plausible value, based on comparisons of growth rates between cells with and without ColE1-type plasmids. In our experiments with killer *S. cerevisiae*, the cost of constitutive toxin production (ε) is minor [21], but the deletion of the hexokinase isoenzyme 2 gene that allows expression of the toxin from the galactose-inducible promoter P_{GAL1} in K carries a growth-rate cost that reduces r_K compared to r_S , resulting in the ratio of realized growth rates $r_K(1 - \varepsilon)/r_S = 0.92 \pm 0.03$ (mean \pm SE). Unless otherwise noted, we will adopt $r_{KS} = 0.85$ and $\varepsilon = 0.2$ in our computations, aligning with the lower bound $r_{KS}(1 - \varepsilon) = 0.68$ of the range reported in the literature.

Box 1: Population growth model

Switching from the normalized strain abundances (n_K, n_S) to the normalized population size ($N(t) = n_K(t) + n_S(t) \in [0, 1]$) and the fraction of killers ($f(t) = n_K(t)/N(t) \in [0, 1]$), we re-write Eq. (1) as

$$\begin{cases} \frac{df}{dt} = \overbrace{f(1-f)}^{\frac{F(f, N, a)}{F(f, N, a)}} \overbrace{\left((1-N)[r_{KS}(1 - \varepsilon a) - 1] + a\gamma f N \right)}^{\frac{G(f, N, a)}{G(f, N, a)}}, \\ \frac{dN}{dt} = \overbrace{N(1-N)}^{\frac{G(f, N, a)}{G(f, N, a)}} \overbrace{\left(1 + [r_{KS}(1 - \varepsilon a) - 1]f \right) - a\gamma N^2 f(1-f)}^{\frac{F(f, N, a)}{F(f, N, a)}}. \end{cases} \quad (2)$$

Under this transformation, the entire horizontal line $N = 0$ maps to the origin $(n_K, n_S) = (0, 0)$ in the original Eq. (1) coordinates. Similarly, the horizontal line $N = 1$ corresponds to $n_K + n_S = 1$, indicating that the system is at carrying capacity.

Definitions and Parameters:

- $f(0) = f_0, N(0) = N_0$: initial killer fraction and population size;
- $r_{KS} := r_K/r_S$: ratio between intrinsic growth rates;
- $a(t) \in [0, 1]$: toxin-production rate;
- $\varepsilon > 0$: cost of producing the toxin;
- $\gamma = \mu C/r_S$: rescaled killing rate.

It is also more convenient to restate the dynamics in terms of the normalized total population size $N(t) = n_K(t) + n_S(t)$ and the fraction of killers $f(t) = n_K(t)/N(t)$. This change of coordinates yields the ODE model Eq. (2) on a unit square, as summarized in Box 1. Fig 3A shows the phase portrait for constitutive killers, with all trajectories approaching $f = 1$ and $N = 1$ (i.e., $(n_K, n_S) = (1, 0)$ – the competitive exclusion of sensitives by killers), which is the only attracting fixed point of Eq. (2) for any fixed $a > 0$. Consistently with experimental data in Fig. 1, when starting from a small initial population N , the trajectories in Fig 3A bend left (i.e., decreasing the killer fraction) for a significant amount of time. This reduction is due to the killers' growth-rate disadvantage ($r_{KS}(1 - \varepsilon) < 1$), which does not prevent their eventual domination.

Do regular dilutions protect the sensitive?

We begin by assuming that dilutions occur regularly every T time units and that at each dilution the relative strain abundances are preserved, but only a fixed fraction ρ of the total population survives². We observe that with relatively frequent dilutions ($T = 1$, i.e., the inter-dilution time equal to the inverse growth rate of sensitive cells), the killers may either progressively increase their relative abundance (Fig. 3B) or decrease it (Fig. 3C), depending on the initial condition.

²We focus here on such "strictly proportional" dilutions for the sake of simplicity and computational efficiency. The results in SI Appendix §S9 show that for most initial conditions the conclusions remain largely the same even with a probabilistic dilution model, where each cell has probability ρ of surviving each dilution.

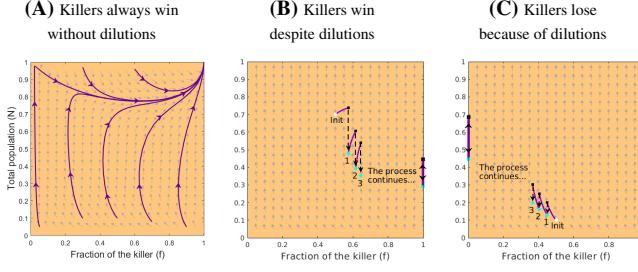


Figure 3. Trajectories of competitions between constitutive killers and sensitives in undisturbed (A) and periodically-diluted (B,C) populations. (A) Killers always win without dilutions. (B,C) With periodic dilutions, their fate depends on the initial condition. (B) With a high enough initial population size, e.g., $(f_0, N_0) = (0.5, 0.7)$, a sequence of dilutions carries killers to an eventual victory (i.e., $f \rightarrow 1$). (C) Starting at a lower population size, e.g., $(f_0, N_0) = (0.5, 0.1)$, the dilutions lead to their demise (i.e., $f \rightarrow 0$). In both cases, the period of dilutions is $T = 1$, and the pre- and post-dilution states are shown with black squares and cyan dots respectively. Once either strain dominates, the population oscillates between the terminal cyan dot and black square at $f = 0$ or $f = 1$. Temporal trajectories associated with panel B are shown in Fig. 4B. Gray arrows denote the vector field directions corresponding to Eq. (2) with $\alpha = 1$. Parameter values: $\varepsilon = 0.2$, $r_{KS} = 0.85$, $\gamma = 1$, and $\rho = 0.65$ used here and throughout the paper unless stated otherwise.

This suggests that while dilution events can disrupt the killers' dominance, a further investigation is needed to determine under which circumstances these interventions help the sensitives if T is small.

Numerical methods make it easy to analyze the performance of constitutive killers for all possible initial states over one cycle; i.e., we use linear partial differential equations (PDEs) to compute the pre-dilution $f(T^-) = f(T^+)$ and $N(T^-) = N(T^+)/\rho$ corresponding to all initial (f_0, N_0) ; see Fig. 4A and SI Appendix §S7.1. Once this mapping is computed, we iterate it to determine the asymptotic outcomes as the number of dilutions approaches infinity. Fig. 4A shows a general trend: larger initial killer fractions f_0 and population sizes N_0 allow constitutive killers to moderately increase their relative abundance by the end of the first cycle. The region below the black-dashed curve in Fig. 4A indicates initial conditions for which killers decrease their fraction in the first cycle; i.e., $f(T^-) < f_0$. Thus, one may expect that populations that start below the black-dashed curve are exactly the ones that become dominated by the sensitive strain with successive dilutions. However, our calculation of the asymptotic limit $\lim_{n \rightarrow \infty} f((nT)^-)$, shows that this is not the case. Fig. 4C-D show that, in the limit of infinite dilutions, the state space is divided into two regions corresponding to the competitive exclusion of the killer by the sensitive strain (blue) and vice versa (red). The actual shades of blue and red in Figs. 4C-D represent the time it takes from the initial (f_0, N_0) to come within the machine accuracy of the asymptotic limit for killers' pre-dilution fraction (which in real systems would be correlated with the time until the competitive exclusion). We highlight five features generic in these computations, which mirror the observations from the experiments of Fig. 2:

1. For every initial population size, there exists a critical killer fraction threshold that determines which strain will eventually dominate.
2. The closer the initial killer fraction is to that threshold, the longer it takes to approach competitive exclusion.
3. Changes in relative abundance in the first cycle are not predictive of the asymptotic limit. Accordingly, in our experiments of

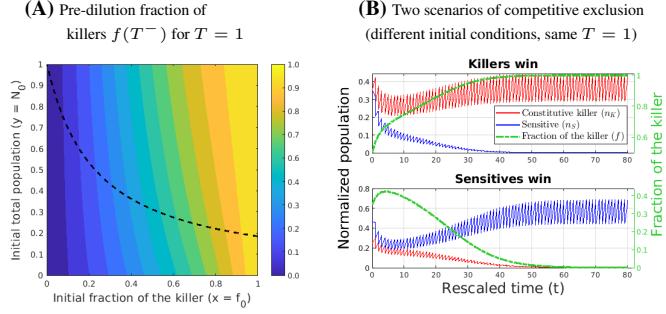


Figure 4. Constitutive killers with regular dilutions: pre-dilution fraction and limiting behavior. (A) Pre-dilution fraction of the killer at the end of the first cycle, $f(T^-)$, for $T = 1$ and any initial condition. Initial conditions below the black, dashed curve lead to $f(T^-) < f_0$. (B) Temporal trajectories of killer (red) and sensitive (blue) population sizes (left axes), and killer fraction (green, right axes), with two different initial conditions (corresponding to the \blacklozenge and \blacktriangle markers in panel C) and dilution period $T = 1$. (C) Time until competitive exclusion (red/blue shades) and limiting killer fraction (red and blue indicate $f = 1$ and $f = 0$ limits, respectively) vary with the initial condition. Within the red and blue regions, the absolute killer and sensitive population sizes reach $n_K \approx 0.45$ and $n_K \approx 0.69$, respectively (see also panel B). Black dashed line as in panel A. (D) Doubling the dilution period extends the range of initial conditions leading to domination by the killer and reduces/increases the timescale over which the killer/sensitive reach domination, respectively. Both initial conditions marked by \blacklozenge and \blacktriangle now lead to killer domination.

Fig. 2, the killer's fraction always decreased in the first cycle due to the low initial population size.

4. When the dilution period grows, this might change the asymptotic limit and the time necessary to approach it. For example, the \blacklozenge -marked initial condition leads to the victory of sensitives when $T = 1$ and of killers when $T = 2$ (Fig. 4C-D). While the \blacktriangle -marked initial condition was already leading to killers' victory even with $T = 1$, with $T = 2$ this exclusion happens much faster. This is consistent with the comparison of 24 h and 48 h dilution trajectories in the experiments.
5. With $T = 1$, neither strain can reach the carrying capacity within one cycle even after the other strain is excluded in the limit. The range of such oscillations can be found analytically (SI Appendix §S3.1), is observed in Fig. 4B, and also consistent with experimental evidence in Fig. S1.

Do toxin-producers benefit from population-sensing?

The results for constitutive killers are revealing and align well with our experiments. However, antagonistic strains often use quorum sensing or environmental signals [53–55] to regulate toxin production, and thus may not engage in antagonistic behavior at all times,

decreasing their growth-rate penalty when the toxin is not produced. Here, rather than modeling specific types of quorum sensing mechanisms, we use a phenomenological approach and explore different notions of optimality for toxin production policies of “omniscient” killers. The results will serve as an upper bound on how well more realistic killers could do given the limits to their sensing abilities.

Supposing that the killers can sense the current size and composition of the population, they could use this to regulate their rate of toxin production, also taking into account the remaining time until the next dilution. More precisely, we will consider a theoretical possibility of their evolving the optimal toxin production rate *in feedback form*: $a_* = a_*(f, N, t)$ to optimize the resulting pre-dilution fraction $f(T^-)$. We will refer to such killers as “myopically-optimal” or simply “myopic” since they optimize the results over a single cycle only, without any regard to the sequence of future dilutions. We use the methods of control theory [65] to find this optimal toxin-production policy in the framework of dynamic programming,³ see the summary in Box 2 and computational details in SI Appendix §S7. The structure of this control problem guarantees that the optimal policy is *bang-bang*: i.e., for generic (f, N, t) , it will be optimal to either not produce the toxin at all ($a_* = 0$) or produce it at the maximum rate ($a_* = 1$).

Box 2: Myopic optimal toxin production policy for population-sensing killers under regular dilutions

The *value function* $u(x, y, t) = \sup_{a(\cdot)} f(T^-)$ is the best pre-dilution killer fraction achievable starting from $f(t) = x$, $N(t) = y$. This value function satisfies a Hamilton-Jacobi-Bellman (HJB) PDE

$$-\frac{\partial u}{\partial t}(x, y, t) = \max_{a \in [0, 1]} \left\{ \nabla u(x, y, t) \cdot \begin{bmatrix} F(x, y, a) \\ G(x, y, a) \end{bmatrix} \right\} \quad (3)$$

with terminal conditions $u(x, y, T) = x$ if $y > 0$ and $u(x, 0, T) = 0$.

F and G are defined in Eq. (2). The myopically optimal toxin-production policy $a_*(x, y, t)$ is an argmax in Eq. (3), see SI Appendix §S7.1.

The corresponding pre-dilution population size $\phi(x, y, t) = N(T^-)$ starting from $f(t) = x$, $N(t) = y$, and using policy $a_*(\cdot)$ satisfies

$$-\frac{\partial \phi}{\partial t}(x, y, t) = \nabla \phi(x, y, t) \cdot \begin{bmatrix} F(x, y, a_*(x, y, t)) \\ G(x, y, a_*(x, y, t)) \end{bmatrix} \quad (4)$$

with the terminal condition $\phi(x, y, T) = y$.

Iterating (u, ϕ) one can obtain the limiting pre-dilution fraction of killers $\hat{u}^\infty(x, y)$ under an infinite sequence of dilutions (SI Appendix §S7.2).

As Fig. 5C shows, these myopic killers try to maximize the early exponential growth by opting not to produce the toxin at first if the initial population and/or their initial fraction are low. Fig. 5A demonstrates that they do better than the constitutive killers over the first cycle, but at least for these parameter values, this sensing-based advantage is minor and mostly pronounced when the initial populations are low. Compared to constitutive killers, the region starting from which the myopic killers eventually dominate expands only slightly, for small N_0 and relatively large f_0 (Fig. 5B). Consequently, the sensitive strain is still protected by periodic dilutions, which allow it to achieve competitive exclusion starting from a broad range of initial configurations.

³Our use of dynamic programming ensures that the computed policy is *globally* optimal and is obtained “in feedback form”. These are important advantages over the *Pontryagin Maximum Principle* approach, which is more commonly used in biological applications [66].

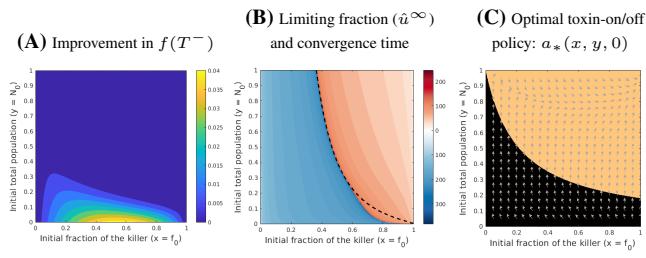


Figure 5. Myopic killers with regular dilutions: pre-dilution and limiting behaviors with regular dilutions and $T = 1$. (A) Population sensing and myopic planning provide a minor improvement to the killer fraction $f(T^-)$ at the end of the first cycle, mostly for initial conditions with low population size. Shown here is the maximized $f(T^-)$ for myopic killers, minus the corresponding $f(T^-)$ for constitutive killers. (B) In the infinite-dilution limit, myopic population sensing expands the set of initial conditions leading to the killer dominance, compared to constitutive killers (black dashed curve reports the blue/red boundary of Fig. 4C). (C) The optimal toxin production policy $a_*(f, N, t)$ (shown here for the initial time $t = 0$ only) is *bang-bang*, equal to 1 in the orange region, and to 0 in the black region. Gray arrows denote the vector field directions corresponding to Eq. (2) with $a = a_*(f, N, 0)$. When the killers switch from no toxin production to constitutive production (i.e., as the competition trajectory crosses from the black to orange region in (f, N) space), the vector field exhibits a jump discontinuity. However, the angular difference between the vectors on either side is consistently small (always $< 1.27^\circ$), making the discontinuity not visually obvious. The policy is time-dependent and the toxin production region changes for $t > 0$.

Who benefits from randomness in dilution times?

The periodic setting investigated up till now provides useful insights but may oversimplify the complex dynamics of populations in natural environments. For example, boom-and-bust dynamics are often driven by fluctuations that occur randomly in time, rather than periodically. Here, we extend our analysis to randomly distributed perturbations.

We model dilutions as random events governed by a Poisson process; so, the duration of inter-dilution time intervals \mathcal{T} are exponentially distributed with rate $\lambda > 0$. As before, we will consider proportional dilutions, preserving f but instantaneously switching from N to ρN . Mathematically, this continuous evolution of $f(t)$ and $N(t)$ punctuated by randomly timed jumps in N can be described as a Piecewise-Deterministic Markov Process (PDMP) [67] and we take advantage of a well-developed theory for optimal control of PDMPs throughout the rest of this paper. We first note that any toxin-production policies will now be independent of time since the last dilution; we will use $\alpha = \alpha(f, N)$ to denote such feedback policies, to distinguish them from $a(f, N, t)$ used in the periodic case above.

Second, we note that our criteria for evaluating initial conditions and the quality of policies become more subtle. Starting from the same initial (f_0, N_0) and using any reasonable fixed policy α , a sequence of randomly timed dilutions might lead to a competitive exclusion of either strain. Fig. 6 illustrates this for the simplest case of constitutive killers; from now on, we will use $\alpha_0 = 1$ to denote their policy. Since a trajectory reaches neither $f = 0$ nor $f = 1$ in finite time, we will select small threshold values γ_d and $(1 - \gamma_v)$, declaring killers’ victory⁴ as soon as $f(t) > \gamma_v$ or sensitives’ victory

⁴It is similarly possible to declare the victory/defeat criteria in terms of strain populations rather than fractions. For most initial conditions, this does not affect the qualitative policy and victory probabilities; see SI Appendix S8.

as soon as $f(t) < \gamma_d$. We then define a new probabilistic metric for policy performance: $\hat{w}^\alpha(f_0, N_0)$ is the probability of the killers attaining their victory (before the toxin-sensitives do) after an arbitrary number of dilutions starting from (f_0, N_0) and using policy α . The rigorous definition and the PDE that \hat{w}^α must satisfy are covered in Box 3. We will use this metric to compare the performance of all toxin-production policies in stochastic environments.

Box 3: Probabilistic performance metric for policies under randomly-timed dilutions

Definitions and Parameters:

- $\Delta_v = \{(x, y) \in [0, 1]^2 \mid x > \gamma_v\}$, victory zone (γ_v —victory threshold);
- $\Delta_d = \{(x, y) \in [0, 1]^2 \mid x < \gamma_d\}$, defeat zone (γ_d —defeat threshold);
- $\Delta = \Delta_v \cup \Delta_d$, terminal set.

(Random) victory time for the killer:

$$T_v(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t) \in \Delta_v; f(0) = x, N(0) = y, \text{ with many dilutions} \right\}. \quad (5)$$

(Random) defeat time for the killer:

$$T_d(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t) \in \Delta_d; f(0) = x, N(0) = y, \text{ with many dilutions} \right\}. \quad (6)$$

(Random) termination time:

$$\hat{T} := \hat{T}(x, y, \alpha(\cdot)) = \min \{T_v(x, y, \alpha(\cdot)), T_d(x, y, \alpha(\cdot))\}. \quad (7)$$

Terminal cost: $g = \begin{cases} 1, & \text{if } (x, y) \in \Delta_v, \\ 0, & \text{if } (x, y) \in \Delta_d. \end{cases}$

Performance metric (probability of killers' victory with α):

$$\hat{w}^\alpha(x, y) = \mathbb{P}(T_v(x, y, \alpha) < T_d(x, y, \alpha)) \quad (8)$$

can be found by numerically solving a first-order linear equation:

$$\lambda [\hat{w}^\alpha(x, \rho y) - \hat{w}^\alpha(x, y)] + \left(\nabla \hat{w}^\alpha(x, y) \cdot \begin{bmatrix} F(x, y, \alpha(x, y)) \\ G(x, y, \alpha(x, y)) \end{bmatrix} \right) = 0, \quad (9)$$

with the boundary condition $\hat{w}^\alpha = g$ on Δ .

See Remark VII in SI Appendix §S7.3 for the numerics.

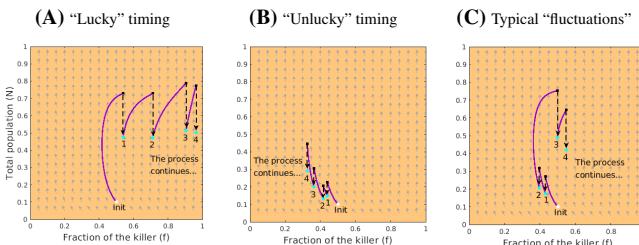


Figure 6. Constitutive killers with random dilution times. Killers can either progressively increase or decrease their fraction in a “lucky” or “unlucky” scenario depicted in panels (A, B), respectively. However, in most cases, their fraction will fluctuate instead (panel C). In all cases, the initial configuration $(f_0, N_0) = (0.5, 0.1)$ is plotted with a white dot and followed by 4 dilution events. The pre-dilution / post-dilution states are again shown with black squares and cyan dots, respectively. Gray arrows denote the vector field directions corresponding to Eq. (2) with $a = 1$.

Since the inter-dilution intervals are random, this also affects the notion of optimal policy. Extending the same myopic approach to stochastic setting, we will call the killers *tactically-optimal* if they follow a policy α_1 chosen to maximize the *expected* killers' fraction just before the next dilution; i.e., $\mathbb{E}[f(\mathcal{T}^-)]$. The subscript in α_1 indicates the temporal horizon for optimizing this policy (one cycle). This bang-bang policy can be found by solving the so-called “randomly-

terminated” problem [68]; mathematical details are in Box 4, part 2.

Box 4: Different types of toxin production policies under randomly-timed dilutions

1. Constitutive policy α_0 :

2. Tactically-optimal policy α_1 :

Maximizes the *expected* killers' fraction just before the next dilution. Given the value function

$$v(x, y) = \sup_{\alpha(\cdot)} \mathbb{E}[f(\mathcal{T}^-; \alpha(\cdot)) \mid f(0) = x, N(0) = y], \quad (10)$$

the tactically-optimal policy α_1 can be found by solving a first-order HJB PDE satisfied by v :

$$\lambda[x - v(x, y)] + \max_{a \in \{0, 1\}} \left\{ \nabla v(x, y) \cdot \begin{bmatrix} F(x, y, a) \\ G(x, y, a) \end{bmatrix} \right\} = 0, \quad (11)$$

with the boundary condition

$$v(x, y) = \begin{cases} 1, & \text{if } x = 1 \text{ and } y \neq 0 \\ 0, & \text{if } y = 0 \text{ or } x = 0. \end{cases} \quad (12)$$

See SI Appendix §S6.1 for the derivation.

3. Strategically-optimal policy α_∞ :

Maximizes the probability of killers winning before the sensitives do (without any limit on the number of dilutions). Given the value function

$$w(x, y) = \sup_{\alpha(\cdot)} \mathbb{P}(T_v(x, y, \alpha(\cdot)) < T_d(x, y, \alpha(\cdot))), \quad (13)$$

the strategically-optimal policy α_∞ can be found by solving a first-order non-local HJB equation satisfied by w :

$$0 = \lambda[w(x, \rho y) - w(x, y)] + \max_{a \in \{0, 1\}} \left\{ \nabla w(x, y) \cdot \begin{bmatrix} F(x, y, a) \\ G(x, y, a) \end{bmatrix} \right\}, \quad (14)$$

with the boundary condition

$$w(x, y) = \begin{cases} 1, & \text{if } x > \gamma_v \text{ and } y \neq 0, \\ 0, & \text{if } x < \gamma_d \text{ or } y = 0. \end{cases} \quad (15)$$

See SI Appendix §S6.2 for the derivation and §S7.3 for the numerics.

We first focus on the case $\lambda = 1$, to ensure that $\mathbb{E}[\mathcal{T}] = 1/\lambda = 1$ matches the period of regular dilutions $T = 1$ considered in the previous section. Fig. 7 compares the performance of constitutive and tactically-optimal killers. Unlike in the periodic case (Fig. 5A), here the advantage of population-sensing (tactically-optimal) killers is significant: they have noticeably better chances of winning than constitutives starting from most initial conditions. Another simple comparison is to focus on the previous boundaries between the blue (deterministic defeat) and red (deterministic winning) in Figs. 4C and 5B. Plotting these boundaries as black dashed lines in Figs. 7A and 7B respectively, we provide a different quantitative measure of tactically-optimal killers' advantage: their average chances for success starting near this deterministic “no microbe’s land” are $\approx 72\%$, compared to only $\approx 46\%$ for the constitutive killers. Interestingly, the tactically-optimal policy $\alpha_1(f, N)$ shown in Fig. 7C is quite close to the zeroth time-slice of the deterministic myopic optimal policy $a_*(f, N, 0)$ from Fig. 5C. The difference in performance comes primarily from the fact that $\alpha_1(f, N)$ is stationary and that occasional long intervals ($\mathcal{T} > 1/\lambda$) really help the tactically-optimal killers. Nevertheless, the toxin-sensitives still have a significant probability of winning on a large set of initial conditions, particularly when the toxin-producers are not starting in the majority.

For each toxin production policy, it is also interesting to ask whether it performs better in deterministic or stochastic environments (assuming that the average dilution frequency and strength are the same in both cases). In the SI Appendix §S4, we perform this comparison by fixing a specific initial population size $N_0 = 0.5$ and averaging over possible initial killer fractions f_0 . The observed perfor-

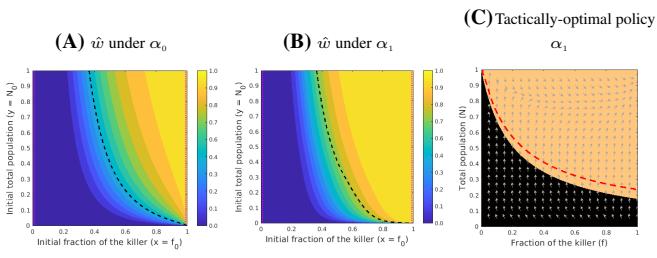


Figure 7. Probabilistic performance of (A) constitutive and (B) tactically-optimal killers with random dilution times ($\lambda = 1$). The probability of attaining competitive exclusion is noticeably higher for the tactically-optimal toxin-producers starting from most initial configurations. Dashed black lines show the boundary of the set from which they could (deterministically) win under periodic dilutions with $T = 1$. In the current random dilutions setting, starting near that boundary gives the tactically-optimal killers an average $\approx 72\%$ chance of winning, compared to only $\approx 46\%$ for the constitutive killers. This is mainly because the tactically-optimal killers do not produce the toxin when their fraction or the overall population size is low (black region in C). In both (A) & (B), the victory and defeat barriers (γ_v and γ_d , respectively) are indicated by vertical magenta dotted lines. In (C), gray arrows represent the vector field directions corresponding to Eq. (2) with $a = \alpha_1(f, N)$, and the red-dashed line denotes the lower boundary of the toxin-on region for a different *strategically-optimal* policy α_∞ discussed below. All parameter values are the same as in Fig. 4C.

mance differences depend non-trivially on ρ and λ ; see Fig. S6. But in general, constitutive killers perform better in stochastic environments when λ is large and ρ is small, and in deterministic environments when λ is small and ρ is large. In contrast, tactically-optimal population-sensing killers always fare better under stochastic dilutions.

Can toxin-producers do better if they are non-myopic?

The optimality of toxin-production policy α_i is tactical (or myopic) because it is selected with only one upcoming dilution in mind, ignoring the ultimate goal of killers to win after arbitrarily many dilutions. It is natural to ask whether they would gain a substantial advantage by selecting a policy that maximizes the probability of attaining their victory before the sensitives, \hat{w}^α . For fixed values of dilution frequency λ and survival factor ρ , such *strategically-optimal* policy $\alpha_\infty(f, N)$ can be computed as described in Box 4 part 3. The subscript in α_∞ indicates that the policy maximizes the probability of eventual victory without any regard to how many dilutions will be required. The strategically-optimal α_∞ prescribes producing toxin slightly more conservatively than the tactical α_i (see the red-dashed line in Fig. 7C). But in the end, α_∞ 's performance is only marginally better for our chosen parameter values; see SI Appendix §S5 and also Fig. 8C.

Is the outcome affected by dilution strength and frequency?

The competition dynamics described so far appear to be robust, holding true for a variety of stochastic environments. In Fig. 8, we focus on a single initial condition $(f_0, N_0) = (0.5, 0.1)$ and compare the performance of constitutive, tactically-optimal and strategically-optimal toxin-production policies for a range of (ρ, λ) values. Predictably, all three of them yield higher chances of winning against the toxin-sensitive strain when the dilutions are rare and weak (small λ , large ρ) – in these regimes, the population gets closer to the carrying capacity in between dilutions, the growth of both strains slows

down, and the toxin's effect becomes more noticeable. As expected, the constitutives are far less effective on most of this map. The biggest surprise is how well the tactically-optimal killers do – their chances of winning are in the worst case only 2.5% below those of the strategically-optimal killers. This is impressive since the tactically-optimal policy α_1 is formulated without any reference to ρ , but works well across a fairly broad range of dilution strengths. But even against the strategically optimized α_∞ , the toxin-sensitives still have a chance of winning above 50% on at least half of this (ρ, λ) map.

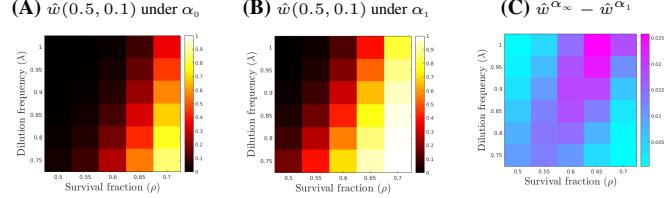


Figure 8. Comparison of probabilistic performance for different toxin production policies starting from $(f_0, N_0) = (0.5, 0.1)$ for a range of survival fractions and dilution frequencies, with random dilution times. Policy α_1 is recomputed for each λ , while policy α_∞ is recomputed for each (ρ, λ) combination. In general, a larger survival rate ρ combined with a smaller dilution frequency λ increases the chances of toxin-producers to win for all three policies. It is clear that the tactically-optimal (panel B) killers significantly outperform the constitutives (panel A). The differences in $\hat{w}(0.5, 0.1)$ between α_∞ and α_1 are small, with the discrepancy increasing toward the upper right corner (panel C).

In conclusion, whether dilutions are randomly timed or periodic and despite the benefits brought by toxin-production regulation to killers, the sensitive strain can prevail by taking advantage of disruptions caused by dilutions.

Discussion

The impact of environmental variability on population size and diversity is well-established in the literature. For example, it is well known that *environmental switches* can make co-existence possible even if all the “un-switched” versions of the environment would result in a competitive exclusion with the same winner [69]. Similarly, the variability in the resulting population is influenced not only by the frequency but also by the timing (periodic vs random) of switching events [70]. While most studies employ environmental switches to model abrupt changes in population dynamics (e.g., instantaneous reductions in carrying capacity or resource processing efficiency), there is also a growing interest in understanding instantaneous exogenous events that impact all subpopulations similarly without affecting the dynamics in between. Dilutions provide a prime example of this, resulting in boom and bust cycles, where the bust phase is nearly instantaneous [71, 72]. Within this emerging focus, our paper is the first demonstration that for antagonistic interactions such population busts can fundamentally alter outcomes, determining which species achieves competitive exclusion.

We have shown both theoretically and experimentally that dilution events can benefit toxin-sensitive strains, when their rate of growth in the exponential phase is larger than that of toxin-producing ones. Using experiments with strains of *S. cerevisiae*, one engineered to constitutively produce the killer toxin K1, and the other sensitive to it, we found that the outcome and dynamics of competition between the two varied with the frequency of periodic dilution events. Because toxin production is often regulated in response to quorum sensing

[54, 55], we used tools of optimal control theory to explore how different toxin-production policies can benefit the producers. Following this approach, we developed two efficient algorithms that (i) calculate the deterministic limit of the relative abundances of these strains as the number of periodic dilutions approaches infinity; and (ii) address the non-local Hamilton-Jacobi-type equation that includes the jumps introduced by dilution events in a stochastic environment. Our numerical experiments consistently support the conclusion that, regardless of toxin production regulation, dilution events often disrupt the dominance of the killer, thereby protecting the sensitive strain from extinction.

Rather than focusing on specific mechanistic models of toxin production regulated by quorum sensing, which would vary across species and would require a large number of parameters and modeling assumptions, we adopted a phenomenological approach to identify theoretical performance bounds for “omniscient killers” capable of measuring population density and fractions. The optimal policies derived here should thus be regarded as upper bounds on the ability of toxin-producing strains to out-compete sensitive ones in fluctuating environments. Future work will explore how more realistic, mechanistic models of toxin production regulated by quorum sensing compare to the optimal policies investigated here. For example, it will be useful to consider models in which the concentration of toxin and/or the concentration of quorum sensing molecules are additional state variables.

In a seminal paper on allelopathy in spatially distributed populations [25], Durrett and Levin showed that the competition of toxin-producing and toxin-sensitive strains displays bi-stability in well-mixed, undisturbed competitions, with either the killer or sensitive strain dominating in the long-term limit depending on the initial condition. In their model, such bi-stability depends critically on the magnitude of a (toxin unrelated) death rate term in their governing equations. In the SI Appendix §S2, we show that using typical laboratory-measured values for growth and death rates, the range of initial conditions for which the sensitive strain competitively excludes the killer is extremely small even under the Durrett-Levin model. This theoretical bi-stability alone is thus unlikely to explain why non-antagonistic strains are found in many microbial communities, and our analysis of disruption effects highlights another contributing factor, which is arguably at least as important. It is of interest to ask how the perturbations investigated here would affect the spatial competition between killer and sensitive cells explored theoretically in [25], where the two strain types were shown to coexist by forming dynamic, single-strain clusters. Similarly, it would be interesting to explore the impact of dilutions on the nucleation criteria that control the killers’ invasion success of a spatially-distributed, resident population [21].

As in any model, simplifying assumptions were made to focus on the main observed phenomena and to ensure the computational efficiency of our approach. But in the future, it might be desirable to relax some of these assumptions to reflect additional features present in realistic environments. One such extension would be to consider more general dilutions, for example, modeling survival rates ρ using a Binomial distribution, rather than assuming fixed relative abundances. Although Monte Carlo simulations (SI Appendix §S9) suggest that results remain qualitatively similar, random dilution outcomes can sometimes produce variability in which strain dominates, as seen in our experiments (Fig. 2C-D). An even more realistic model could also incorporate random variability in parameters such as growth rates (r_K, r_S), killing rates γ , and dilution factors ρ and λ , which are influenced by environmental and biological fluctuations [73, 74]. These fluctuations may better capture population dynamics but would increase computational complexity, requiring a hybrid

model with discrete and continuous random perturbations based on jump-diffusion processes [75, 76].

Incorporating evolutionary adaptation, such as mutational dynamics, would also enhance our understanding of the competition between toxin-producing and toxin-sensitive strains in natural environments over longer timescales. Experiments indicate that killer strains can lose or alter their toxin-producing ability, and sensitive strains may develop resistance to the toxin [21, 23, 77]. The ultimate success of the killer strain in our model is negatively impacted by the costs associated with toxin production, suggesting that evolutionary adaptation may aim to minimize these costs [77]. Additionally, evolutionary adaptation may enable antagonistic strains to regulate toxin production in response to their environment, population abundance, or the presence of competitors [55], possibly approaching the performance of the optimal policies described here. Toxin resistance can arise through various mechanisms, such as alterations in toxin receptors or translocation pathways, which may have antagonistic pleiotropic effects where resistance incurs a cost in terms of growth rate [78]. This growth rate penalty will, in turn, influence the competitive dynamics with the killer strain. Finally, in environments experiencing disturbances, evolutionary adaptation may promote increased retention (ρ) [79] or even alter the rate of disturbances (λ). For example, production of surface-attachment molecules, pili or fimbriae by microbes such as *Pseudomonas aeruginosa*, *Vibrio cholerae*, *Clostridium difficile* and *Streptococcus salivarius* can help them adhere to surfaces in their environment and prevent them from being washed away in fluid environments such as the gastrointestinal tract, the oral cavity, or natural water bodies [11, 80, 81]. *C. difficile* and *V. cholerae*, in addition to adhering to surfaces such as the intestinal mucosa, can also cause diarrhea [82] and thus potentially control the environment dilution rate and intensity, at least transiently.

In conclusion, we posit that the fitness cost incurred by toxin-producing strains may be particularly detrimental in boom-and-bust environments in which populations undergo regular or stochastic dilutions. We propose this as a possible explanation for why both antagonistic and non-antagonistic microbes are found in nature, and why environments with higher turnover rates may favor the latter [51, 52].

Materials and Methods

Competition experiments.

The killer strain yAG171b expressed the toxin gene K1 from the chromosome from the galactose-inducible *P_{GAL1}* promoter. To enable strain differentiation at the flow cytometer, strain yAG171b expressed the fluorescent protein CyOFP1opt [83], whereas strain yAG177 expressed the fluorescent protein ymCitrine (Table 3). Cells of the two strains could be clearly identified via flow cytometry and separated in the FL-1A (533±30 nm) vs FL-2A (585±40 nm) emission channels with excitation at 488 nm. Dead yAG177 cells (confirmed via Propidium-Iodide staining) formed a separate cluster in the FL-1A vs FL-2A scatter plot, and were excluded when calculating killer fractions. No dead cells were observed when growing killer cells in isolation. To prevent metabolization of galactose and to make the expression of K1 titratable, both *GAL1* and *GAL10* were deleted, and *GAL3* was placed under the constitutive promoter *P_{ACT1}* [84], in both strains. In addition, to prevent catabolite repression that would have prevented expression of the K1 toxin from *P_{GAL1}* in the presence of glucose, we deleted *HXK2* in yAG171b. This deletion confers a growth-rate deficit to yAG171b, compared to yAG177, in addition to the cost of toxin production. Strain construction is described in the SI Appendix §S1.2. Experiments were performed with strains

yAG171b (killer) and yAG177 (sensitive) in filter-sterilized Complete Synthetic Medium (CSM) buffered at pH 4.5 with succinic acid. The medium composition is described in the SI Appendix §S1.2. The medium was supplemented with 500 μ M galactose for induction of the killer toxin gene K1. Experiments were performed in 96 well plates with flat bottom and a transparent lid, incubated in a plate reader at 23°C with orbital shaking at 425 rpm with an amplitude of 3 mm. For the experiment of Fig. 1, we initialized six technical replicates for each initial fraction and initial population size. Initial fractions of killers were measured via flow cytometry. The experiment was initialized with overnights grown in CSM at 30°C, which were diluted in CSM with 500 μ M galactose by a factor 3125 (N_0 treatment), 312.5 (10 $\cdot N_0$ treatment), 156.25 (20 $\cdot N_0$ treatment) and 62.5 (50 $\cdot N_0$ treatment). For the experiments of Fig. 2, we initialized four technical replicates for six different 24 h and 48 h dilution cycle treatments, targeting initial killer fractions of 32%, 36%, 40%, 44%, 48% and 52%. Actual initial fractions were measured via flow cytometry and were within $\pm 1\%$ of the targeted ones; see Fig. 2. The experiment was initialized with overnights grown in CSM at 30°C, which were diluted by a factor 3125 in CSM with 500 μ M galactose. Successive daily dilutions were performed with a dilution factor of 225. In both experiments, the culture volume was 150 μ L per replicate.

Acknowledgments

We thank Andrew W. Murray, Stephen P. Ellner, and Tobias Dörr for insightful comments on the manuscript. AG thanks Marie F. Gorwa-Grauslund for gifting a plasmid with *CyOFFlopt*. AG acknowledges support by the National Institute of General Medical Sciences of the National Institutes of Health, United States of America under award number 1R35GM147493. MW and AV acknowledge support by the National Science Foundation (awards DMS-1645643 and DMS-2111522) as well as the Air Force Office of Scientific Research (award FA9550-22-1-0528).

Author Declaration

The authors declare no conflict of interest.

References

- [1] S Brook Peterson, Savannah K Bertolli, and Joseph D Mougous. The central role of interbacterial antagonism in bacterial life. *Current Biology*, 30(19):R1203–R1214, 2020.
- [2] Nina S Atanasova, Maija K Pietilä, and Hanna M Oksanen. Diverse antimicrobial interactions of halophilic archaea and bacteria extend over geographical distances and cross the domain barrier. *MicrobiologyOpen*, 2(5):811–825, 2013.
- [3] Margaret A Riley and John E Wertz. Bacteriocins: evolution, ecology, and application. *Annual Reviews in Microbiology*, 56(1):117–137, 2002.
- [4] Elisa T Granato, Thomas A Meiller-Legrand, and Kevin R Foster. The evolution and ecology of bacterial warfare. *Current biology*, 29(11):R521–R537, 2019.
- [5] Thomas J Travers-Cook, Jukka Jokela, and Claudia C Buser. The evolutionary ecology of fungal killer phenotypes. *Proceedings of the Royal Society B*, 290(2005):20231108, 2023.
- [6] Richard A Long and Farooq Azam. Antagonistic interactions among marine pelagic bacteria. *Applied and environmental microbiology*, 67(11):4975–4983, 2001.
- [7] Sijmen E Schoustra, Jonathan Dench, Rola Dali, Shawn D Aaron, and Rees Kassen. Antagonistic interactions peak at intermediate genetic distance in clinical and laboratory strains of *pseudomonas aeruginosa*. *Bmc Microbiology*, 12:1–9, 2012.
- [8] Rocío-Anaís Pérez-Gutiérrez, Varinia López-Ramírez, Africa Islas, Luis David Alcaraz, Ismael Hernández-González, Beatriz Carely Luna Olivera, Moisés Santillán, Luis E Eguiarte, Valeria Souza, Michael Travisano, et al. Antagonism influences assembly of a bacillus guild in a local community and is depicted as a food-chain network. *The ISME Journal*, 7(3):487–497, 2013.
- [9] Jakob Russel, Henriette L Røder, Jonas S Madsen, Mette Burmølle, and Søren J Sørensen. Antagonism correlates with metabolic similarity in diverse bacteria. *Proceedings of the National Academy of Sciences*, 114(40):10684–10688, 2017.
- [10] Sara Arbulu and Morten Kjos. Revisiting the multifaceted roles of bacteriocins. *Microbial Ecology*, 87(1):1–14, 2024.
- [11] Maria Santagati, Marina Scillato, Francesco Patane, Caterina Aiello, and Stefania Stefanì. Bacteriocin-producing oral streptococci and inhibition of respiratory pathogens. *FEMS Immunology & Medical Microbiology*, 65(1):23–31, 2012.
- [12] Jinshui Zheng, Michael G Gänzle, Xiaoxi B Lin, Lifang Ruan, and Ming Sun. Diversity and dynamics of bacteriocins from human microbiome. *Environmental microbiology*, 17(6):2133–2143, 2015.
- [13] Leonor García-Bayona and Laurie E Comstock. Bacterial antagonism in host-associated microbial communities. *Science*, 361(6408):eaat2456, 2018.
- [14] Simon Heilbronner, Bernhard Krismer, Heike Brötz-Oesterhelt, and Andreas Peschel. The microbiome-shaping roles of bacteriocins. *Nature Reviews Microbiology*, 19(11):726–739, 2021.
- [15] Sean C Booth, William PJ Smith, and Kevin R Foster. The evolution of short-and long-range weapons for bacterial competition. *Nature Ecology & Evolution*, 7(12):2080–2091, 2023.
- [16] Ruth Schmidt, Viviane Cordovez, Wietse De Boer, Jos Raaijmakers, and Paolina Garbeva. Volatile affairs in microbial interactions. *The ISME journal*, 9(11):2329–2335, 2015.
- [17] Joachim Becker, Nico Eisenhauer, Stefan Scheu, and Alexandre Jousset. Increasing antagonistic interactions cause bacterial communities to collapse at high diversity. *Ecology Letters*, 15(5):468–474, 2012.
- [18] Judith Feichtmayer, Li Deng, and Christian Griebler. Antagonistic microbial interactions: contributions and potential applications for controlling pathogens in the aquatic systems. *Frontiers in microbiology*, 8:281883, 2017.
- [19] John R Tagg, Liam K Harold, Rohit Jain, and John DF Hale. Beneficial modulation of human health in the oral cavity and beyond using bacteriocin-like inhibitory substance-producing streptococcal probiotics. *Frontiers in Microbiology*, 14:1161155, 2023.

[20] Luke McNally, Eryn Bernardy, Jacob Thomas, Arben Kalziki, Jennifer Pentz, Sam P Brown, Brian K Hammer, Peter J Yunker, and William C Ratcliff. Killing by type vi secretion drives genetic phase separation and correlates with increased cooperation. *Nature communications*, 8(1):14371, 2017.

[21] Andrea Giometto, David R Nelson, and Andrew W Murray. Antagonism between killer yeast strains as an experimental model for biological nucleation dynamics. *Elife*, 10:e62932, 2021.

[22] Raymond Copeland, Christopher Zhang, Brian K Hammer, and Peter J Yunker. Spatial constraints and stochastic seeding subvert microbial arms race. *PLOS Computational Biology*, 20(1):e1011807, 2024.

[23] Sean W Buskirk, Alecia B Rokes, and Gregory I Lang. Adaptive evolution of nontransitive fitness in yeast. *Elife*, 9:e62238, 2020.

[24] David M Gordon and Claire L O'Brien. Bacteriocin diversity and the frequency of multiple bacteriocin production in *Escherichia coli*. *Microbiology*, 152(11):3239–3244, 2006.

[25] Rick Durrett and Simon Levin. Allelopathy in spatially distributed populations. *Journal of theoretical biology*, 185(2):165–171, 1997.

[26] Benjamin Kerr, Margaret A Riley, Marcus W Feldman, and Brendan JM Bohannan. Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature*, 418(6894):171–174, 2002.

[27] Carey D Nadell, Knut Drescher, and Kevin R Foster. Spatial structure, cooperation and competition in biofilms. *Nature Reviews Microbiology*, 14(9):589–600, 2016.

[28] Joseph H Connell. Diversity in tropical rain forests and coral reefs: high diversity of trees and corals is maintained only in a nonequilibrium state. *Science*, 199(4335):1302–1310, 1978.

[29] Michael Huston. A general hypothesis of species diversity. *The American Naturalist*, 113(1):81–101, 1979.

[30] Richard Levins. Coexistence in a variable environment. *The American Naturalist*, 114(6):765–783, 1979.

[31] Peter Chesson. Multispecies competition in variable environments. *Theoretical population biology*, 45(3):227–276, 1994.

[32] Daniel R Amor, Christoph Ratzke, and Jeff Gore. Transient invaders can induce shifts between alternative stable states of microbial communities. *Science Advances*, 6(8):eaay8676, 2020.

[33] Christopher P Mancuso, Hyunseok Lee, Clare I Abreu, Jeff Gore, and Ahmad S Khalil. Environmental fluctuations reshape an unexpected diversity-disturbance relationship in a microbial community. *Elife*, 10:e67175, 2021.

[34] Michael J Behrenfeld, Yongxiang Hu, Robert T O'Malley, Emmanuel S Boss, Chris A Hostetler, David A Siegel, Jorge L Sarmiento, Jennifer Schulien, Johnathan W Hair, Xiaomei Lu, et al. Annual boom–bust cycles of polar phytoplankton biomass revealed by space-based lidar. *Nature Geoscience*, 10(2):118–122, 2017.

[35] Flora Vincent, Matti Gralka, Guy Schleyer, Daniella Schatz, Miguel Cabrera-Brufau, Constanze Kuhlisch, Andreas Sichert, Silvia Vidal-Melgosa, Kyle Mayers, Noa Barak-Gavish, et al. Viral infection switches the balance between bacterial and eukaryotic recyclers of organic matter during coccolithophore blooms. *Nature communications*, 14(1):510, 2023.

[36] Manoshi S Datta, Elzbieta Sliwerska, Jeff Gore, Martin F Polz, and Otto X Cordero. Microbial interactions lead to rapid microscale successions on model marine particles. *Nature communications*, 7(1):11965, 2016.

[37] Tami L Swenson, Ulas Karaoz, Joel M Swenson, Benjamin P Bowen, and Trent R Northen. Linking soil biology and chemistry in biological soil crust using isolate exometabolomics. *Nature communications*, 9(1):19, 2018.

[38] Kamilla S Sjøgaard, Thomas B Valdemarsen, and Alexander H Treusch. Responses of an agricultural soil microbiome to flooding with seawater after managed coastal realignment. *Microorganisms*, 6(1):12, 2018.

[39] Racheal N Upton, Elizabeth M Bach, and Kirsten S Hofmockel. Spatio-temporal microbial community dynamics within soil aggregates. *Soil Biology and Biochemistry*, 132:58–68, 2019.

[40] Samuel E Barnett, Nicholas D Youngblut, Chantal N Koechli, and Daniel H Buckley. Multisubstrate dna stable isotope probing reveals guild structure of bacteria that mediate soil carbon cycling. *Proceedings of the National Academy of Sciences*, 118(47):e2115292118, 2021.

[41] David A Relman. The human microbiome: ecosystem resilience and health. *Nutrition reviews*, 70(suppl_1):S2–S9, 2012.

[42] Emily SC Rittershaus, Seung-Hun Baek, and Christopher M Sassetti. The normalcy of dormancy: common themes in microbial quiescence. *Cell host & microbe*, 13(6):643–651, 2013.

[43] Caroline Vincent, Mark A Miller, Thaddeus J Edens, Sudeep Mehrotra, Ken Dewar, and Amee R Manges. Bloom and bust: intestinal microbiota dynamics in response to hospital exposures and *Clostridium difficile* colonization or infection. *Microbiome*, 4:1–11, 2016.

[44] Christoph A Thaiss, Maayan Levy, Tal Korem, Lenka Dohnalová, Hagit Shapiro, Diego A Jaitin, Eyal David, Deborah R Winter, Meital Gury-BenAri, Evgeny Tatirovsky, et al. Microbiota diurnal rhythmicity programs host transcriptome oscillations. *Cell*, 167(6):1495–1510, 2016.

[45] Jason J Flowers, Tracey A Cadkin, and Katherine D McMahon. Seasonal bacterial community dynamics in a full-scale enhanced biological phosphorus removal plant. *Water research*, 47(19):7019–7031, 2013.

[46] Orr H Shapiro, Ariel Kushmaro, and Asher Brenner. Bacteriophage predation regulates microbial abundance and diversity in a full-scale bioreactor treating industrial wastewater. *The ISME journal*, 4(3):327–336, 2010.

[47] Marc W Van Goethem, Tami L Swenson, Gareth Trubl, Simon Roux, and Trent R Northen. Characteristics of wetting-induced bacteriophage blooms in biological soil crust. *MBio*, 10(6):10–1128, 2019.

[48] MR Brown, JC Baptista, M Lunn, DL Swan, SJ Smith, RJ Davenport, BD Allen, WT Sloan, and TP Curtis. Coupled virus-bacteria interactions and ecosystem function in an engineered microbial system. *Water Research*, 152:264–273, 2019.

[49] Sean M Gibbons, Monika Scholz, Alan L Hutchison, Aaron R Dinner, Jack A Gilbert, and Maureen L Coleman. Disturbance regimes predictably alter diversity in an ecologically complex bacterial system. *MBio*, 7(6):10–1128, 2016.

[50] Jen Nguyen, Juanita Lara-Gutiérrez, and Roman Stocker. Environmental fluctuations and their effects on microbial communities, populations and individuals. *FEMS microbiology reviews*, 45(4):fuaa068, 2021.

[51] Belinda Barnes, Harvinder Sidhu, and David M Gordon. Host gastro-intestinal dynamics and the frequency of colicin production by escherichia coli. *Microbiology*, 153(9):2823–2827, 2007.

[52] Marina V Kuznetsova, Veronika S Mihailovskaya, Natalia B Remezovskaya, and Marjanca Starčić Erjavec. Bacteriocin-producing escherichia coli isolated from the gastrointestinal tract of farm animals: Prevalence, molecular characterization and potential for application. *Microorganisms*, 10(8):1558, 2022.

[53] Antonio Maldonado-Barragán and Stuart A West. The cost and benefit of quorum sensing-controlled bacteriocin production in lactobacillus plantarum. *Journal of evolutionary biology*, 33(1): 101–111, 2020.

[54] Diego Gonzalez and Despoina AI Mavridou. Making the best of aggression: the many dimensions of bacterial toxin regulation. *Trends in microbiology*, 27(11):897–905, 2019.

[55] Rene Niehus, Nuno M Oliveira, Aming Li, Alexander G Fletcher, and Kevin R Foster. The evolution of strategy in bacterial warfare via the regulation of bacteriocins and antibiotics. *Elife*, 10:e69756, 2021.

[56] Eric L Miller, Morten Kjos, Monica I Abrudan, Ian S Roberts, Jan-Willem Veening, and Daniel E Rozen. Eavesdropping and crosstalk between secreted quorum sensing peptide signals that regulate bacteriocin production in streptococcus pneumoniae. *The ISME journal*, 12(10):2363–2375, 2018.

[57] Nazanin Saeidi, Choon Kit Wong, Tat-Ming Lo, Hung Xuan Nguyen, Hua Ling, Susanna Su Jan Leong, Chueh Loo Poh, and Matthew Wook Chang. Engineering microbes to sense and eradicate pseudomonas aeruginosa, a human pathogen. *Molecular systems biology*, 7(1):521, 2011.

[58] H Bussey. K1 killer toxin, a pore-forming protein from yeast. *Molecular microbiology*, 5(10):2339–2343, 1991.

[59] Stefanie Gier, Matthias Lermen, Manfred J Schmitt, and Frank Breinig. Substitution of cysteines in the yeast viral killer toxin k1 precursor reveals novel insights in heterodimer formation and immunity. *Scientific reports*, 9(1):13127, 2019.

[60] Avika Ruparell, Jean Frederic Dubern, Cartherine A Ortori, Freya Harrison, Nigel M Halliday, Abigail Emtage, MM Ashawesh, Charles A Laughton, Stephen P Diggle, Paul Williams, et al. The fitness burden imposed by synthesizing quorum sensing signals. *Scientific Reports*, 6(1):33101, 2016.

[61] Anna S Weiß, Alexandra Götz, and Madeleine Opitz. Dynamics of colicin2 production and release determine the competitive success of a toxin-producing bacterial population. *Scientific reports*, 10(1):4052, 2020.

[62] Jason Pintar and William T Starmer. The costs and benefits of killer toxin production by the yeast *pichia kluyveri*. *Antonie van Leeuwenhoek*, 83:89–97, 2003.

[63] Lin Chao and Bruce R Levin. Structured habitats and the evolution of anticompetitor toxins in bacteria. *Proceedings of the National Academy of Sciences*, 78(10):6324–6328, 1981.

[64] Zhijun Wang, Li Xiang, Junjie Shao, Alicja Węgrzyn, and Grzegorz Węgrzyn. Effects of the presence of cole1 plasmid dna in escherichia coli on the host cell metabolism. *Microbial Cell Factories*, 5:1–18, 2006.

[65] Wendell H Fleming and Raymond W Rishel. *Deterministic and stochastic optimal control*, volume 1. Springer, 2012.

[66] Suzanne Lenhart and John T Workman. *Optimal control applied to biological models*. Chapman and Hall/CRC, 2007.

[67] M. H. A. Davis. Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society. Series B (Methodological)*, 46(3):353–388, 1984. ISSN 00359246.

[68] June Andrews and Alexander Vladimirsny. Deterministic control of randomly-terminated processes. *Interfaces and Free Boundaries*, 16(1):1–40, 2014.

[69] Alexandru Hening and Dang H Nguyen. The competitive exclusion principle in stochastic environments. *Journal of mathematical biology*, 80:1323–1351, 2020.

[70] Ami Taitelbaum, Robert West, Michael Assaf, and Mauro Mobilia. Population dynamics in a changing environment: random versus periodic switching. *Physical Review Letters*, 125(4):048105, 2020.

[71] Zihan Wang, Akshit Goyal, Veronika Dubinkina, Ashish B George, Tong Wang, Yulia Fridman, and Sergei Maslov. Complementary resource preferences spontaneously emerge in diauxic microbial communities. *Nature communications*, 12(1): 6661, 2021.

[72] Yulia Fridman, Zihan Wang, Sergei Maslov, and Akshit Goyal. Fine-scale diversity of microbial communities due to satellite niches in boom and bust environments. *PLoS computational biology*, 18(12):e1010244, 2022.

[73] Bihter Bayramoglu, David Toubiana, Simon Van Vliet, R Fredrik Inglis, Nadav Shnerb, and Osnat Gillor. Bet-hedging in bacteriocin producing escherichia coli populations: the single cell perspective. *Scientific reports*, 7(1):42068, 2017.

[74] Laurent Philippot, Bryan S Griffiths, and Silke Langenbeder. Microbial community resilience across ecosystems and multiple disturbances. *Microbiology and Molecular Biology Reviews*, 85(2):10–1128, 2021.

[75] Bernt Øksendal and Agnes Sulem. Stochastic control of jump diffusions. In *Applied Stochastic Control of Jump Diffusions*, pages 93–155. Springer, 2019.

- [76] Steven E Shreve et al. *Stochastic calculus for finance II: Continuous-time models*, volume 11. Springer, 2004.
- [77] Magdalena D Pieczynska, Dominika Wloch-Salamon, Ryszard Korona, and J Arjan GM de Visser. Rapid multiple-level coevolution in experimental populations of yeast killer and nonkiller strains. *Evolution*, 70(6):1342–1353, 2016.
- [78] Michael Feldgarden and Margaret A Riley. The phenotypic and fitness effects of colicin resistance in escherichia coli k-12. *Evolution*, 53(4):1019–1027, 1999.
- [79] Elyse A Hope, Clara J Amorosi, Aaron W Miller, Kolena Dang, Caiti Smukowski Heil, and Maitreya J Dunham. Experimental evolution reveals favored adaptive routes to cell aggregation in yeast. *Genetics*, 206(2):1153–1167, 2017.
- [80] Olga E Petrova and Karin Sauer. Sticky situations: key components that control bacterial surface attachment. *Journal of bacteriology*, 194(10):2413–2425, 2012.
- [81] WF Van Zyl, SM Deane, and LMT Dicks. Bacteriocin production and adhesion properties as mechanisms for the anti-listerial activity of lactobacillus plantarum 423 and enterococcus mundtii st4sa. *Beneficial microbes*, 10(3):329–349, 2019.
- [82] Wenjing Zhao, Florence Caro, William Robins, and John J Mekalanos. Antagonism toward the intestinal microbiota and its effect on vibrio cholerae virulence. *Science*, 359(6372):210–213, 2018.
- [83] Raquel Perruca-Foncillas, Johan Davidsson, Magnus Carlquist, and Marie F Gorwa-Grauslund. Assessment of fluorescent protein candidates for multi-color flow cytometry analysis of saccharomyces cerevisiae. *Biotechnology Reports*, 34:e00735, 2022.
- [84] Nicholas T Ingolia and Andrew W Murray. Positive-feedback loops as a flexible biological module. *Current biology*, 17(8):668–677, 2007.

Supporting Information (SI) Appendix

SI Appendix provides additional details for the main text. We continue to use reference numbers of equations and figures previously introduced in the main text. The numbering of figures in SI text is prefixed with an “S”, and bibliographic references used here are listed and numbered separately.

S1 Experiments

S1.1 Additional experimental results.

Fig. S1 shows cell density estimates obtained from Optical Density (OD600) measurements taken throughout the experiment of Fig. 2, showing that populations diluted every 24 h (teal curves) did not reach carrying capacity, whereas those diluted every 48 h (pink curves) did.

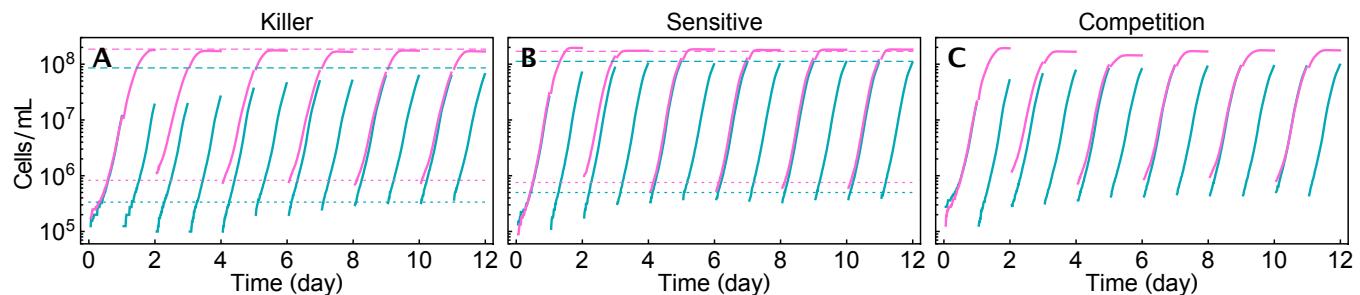


Figure S1. Mean cell density values recorded throughout the experiment of Fig. 2. The mean cell densities in the 24 h cycle replicates are shown in teal, and the 48 h cycle ones are in pink. (A) Mean cell density across three replicates initialized with 100% killer. (B) Mean cell density across three replicates initialized with 0% killer. (C) Mean cell density across four replicates initialized with 44% killer. Cell densities were estimated from Optical Density (OD600) data via a cubic calibration curve. Dashed and dotted lines are predictions of asymptotic pre- and post-dilution cell densities predictions based on Eq. (S3.1.3) with non-normalized population sizes.

S1.2 Strains, oligos, plasmids and additional methods.

The Complete Synthetic Medium (CSM) composition was, per liter: 0.59 g of amino acid mix, 1.706 g of Yeast Nitrogen Base without amino acids, carbohydrates, ammonium sulfate, ferric chloride and copper sulfate (USBiological Y2030-03), 20 g of dextrose, 4 mL of a 5 g/L L-histidine hydrochloride monohydrate stock solution, 20 mL of a 5 g/L L-leucine stock solution, 8 mL of a 2.5 g/L uracil stock solution, 20 mL of a stock solution containing 5 g/L adenine hydrochloride and 5 g/L L-tryptophan, 11.2 g of succinic acid, 5 g of ammonium sulfate and NaOH to reach pH 4.5. The amino acid mix was made by mixing 10 g L-arginine hydrochloride, 16 g L-aspartic acid, 10 g L-isoleucine, 10 g L-lysine hydrochloride, 4 g L-methionine, 10 g L-phenylalanine, 20 g L-threonine, 10 g L-tyrosine and 28 g L-valine.

Tables 1, 2, and 3 report oligos, plasmids and strains used for the experiments. We substituted ymCherry in pAG11 with CyOFP1opt from pRP008 using Gibson assembly, producing pAG134, whose sequence was verified via long-read sequencing. Strain yAG171 was obtained by digesting pAG134 with PpuMI and transforming it into yAG75 [1] for integration of pAG134 in T_{CYC1} . Strain yAG177 was obtained by digesting pAG5 [1] with PpuMI and transforming it into yAG74 [1] for integration of pAG5 in P_{ACT1} .

Oligo name	Oligo sequence
oAG187	cgctgaggactTGACCACACCTCTACCGG
oAG188	gctgaggcatTCGACCTGCAGCGTACGAAG
oAG189	tgcaggtcgaATGCCTCAGCACTAGTCCTG
oAG190	tgcaggtcgaATGCCTCAGCACTAGTCCTG

Table 1: DNA oligos used to assemble pAG134 via Gibson assembly. Bases in capital letters represent homology to the PCR template (pAG11 for yAG187/188 and pRP008 for yAG189/190) Lowercase bases represent homology to the backbone or fragment for Gibson assembly.

Plasmid name	Relevant transcriptional units	Sourced from
pAG5	P_{ACT1} -ymCitrine- T_{ADH1} P_{TEF} -KanMX6- T_{TEF}	[1]
pAG11	P_{ACT1} -ymCherry- T_{ADH1} P_{TEF} -KanMX6- T_{TEF} P_{GAL1} -K1- T_{CYC1}	[1]
pRP008	P_{TEF1} -CyOFP1opt- T_{ADH1}	[2]
pAG134	P_{ACT1} -CyOFP1opt- T_{ADH1} P_{TEF} -KanMX6- T_{TEF} P_{GAL1} -K1- T_{CYC1}	This study

Table 2: Plasmids used for strain construction.

Strain name	Genotype
yAG171	can1-100, his3-11,15, ura3Δ0, BUD4-S288C, gal1/10Δ::LEU2, prGAL3Δ::His3MX6- P_{ACT1} -GAL3, hxx2Δ::HphMX4, P_{TEF1} -CyOFP1Opt- T_{ADH1} , P_{TEF1} -KanMX6- T_{TEF} , prGAL1-K1- T_{CYC1}
yAG177	can1-100, his3-11,15, ura3Δ0, BUD4-S288C, gal1/10Δ::LEU2, P_{GAL3} Δ::His3MX6- P_{ACT1} -GAL3, P_{ACT1} -ymCitrine- T_{ADH1} , P_{TEF1} -KanMX6- T_{TEF}

Table 3: Strains used for the experiments.

S2 Comparison with the Durrett-Levin model

In the main text, all of our results are based on a basic competition model between the toxin-producing “killer” strain (K) and the toxin-sensitive strain (S) (Eq. (1) of the main text). This model assumes the growth of both strains is logistic with respective intrinsic growth rates (r_K, r_S) and a shared carrying capacity C . The toxin-induced death rate of the sensitive cells is proportional to the product of the strain densities, $n_K n_S$. Focusing on *constitutive killers* who produce the toxin at the maximal rate $a = 1$, the original Eq. (1) reads

$$\begin{cases} \frac{dn_K}{dt}(t) = r_{KS}(1 - \varepsilon)(1 - n_K - n_S)n_K, \\ \frac{dn_S}{dt}(t) = (1 - n_K - n_S)n_S - \gamma n_K n_S, \end{cases} \quad (\text{S2.1})$$

where $r_{KS} = r_K/r_S < 1$ is the intrinsic growth rates ratio, γ is the killing rate rescaled by r_S , and ε is the cost associated with producing toxin at the maximal rate ($a = 1$). We note three equilibria (fixed points) of Eq. (S2.1):

1. $(n_K, n_S) = (0, 0)$ – a nodal source
2. $(n_K, n_S) = (1, 0)$ – a nodal sink
3. $(n_K, n_S) = (0, 1)$ – a degenerate node

It follows that a competitive exclusion of the sensitive by the killer will be observed starting from any non-equilibrium initial condition. Despite this eventual dominance, the sensitive strain may transiently increase its relative abundance depending on the initial configuration; see Fig. S2. This occurs because the growth rate benefit of not producing the toxin initially outweighs the deficit caused by toxin-induced death. Fig. S3D further illustrates this phenomenon in a phase plane, showing that the sensitive population begins to decline once the total population reaches a sufficiently large size.

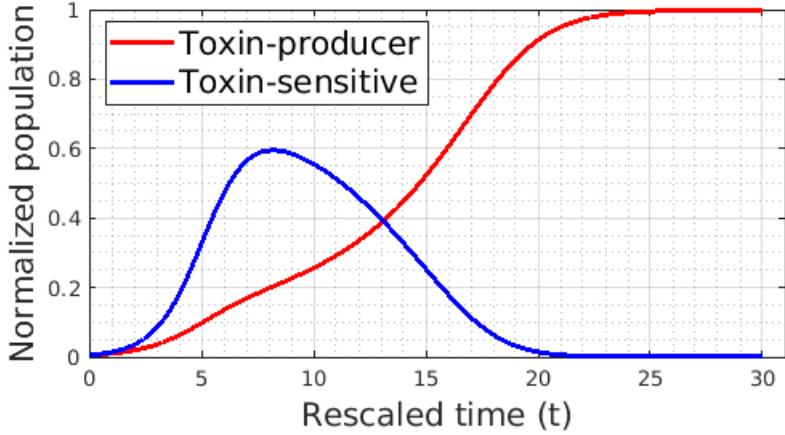


Figure S2. Population trajectories for a constitutive toxin-producing strain (based on Eq. (S2.1)) competing against a toxin-sensitive one. Starting with 50% of each strain and a total population equal to 1% of the carrying capacity, sensitives initially grow faster due to their higher intrinsic growth rate. However, as the population size approaches the carrying capacity, the intrinsic growth advantage of the sensitive strain diminishes, leading to eventual domination by the toxin producer due to the toxin’s action. Parameter values: $r_{KS} = 0.85$, $\varepsilon = 0.2$, and $\gamma = 1$.

One natural extension of model Eq. (S2.1) is to incorporate natural (toxin-unrelated) death rates (δ_K, δ_S) for both strains. With both of these rates rescaled by r_S^{-1} , this yields a system of ODEs equivalent to a competition model previously studied by Durrett and Levin [3] :

$$\begin{cases} \frac{dn_K}{dt}(t) = r_{KS}(1 - \varepsilon)(1 - n_K - n_S)n_K - \delta_K n_K, \\ \frac{dn_S}{dt}(t) = (1 - n_K - n_S)n_S - \gamma n_K n_S - \delta_S n_S. \end{cases} \quad (\text{S2.2})$$

By setting both equations to zero, we find four fixed points of this system:

1. $(n_K, n_S) = (0, 0)$
2. $(n_K, n_S) = \left(1 - \frac{\delta_K}{r_{KS}(1 - \varepsilon)}, 0\right)$
3. $(n_K, n_S) = (0, 1 - \delta_S)$
4. $(n_K, n_S) = \left(\frac{\delta_K}{\gamma r_{KS}(1 - \varepsilon)} - \frac{\delta_S}{\gamma}, \frac{\delta_S + \gamma}{\gamma} - \frac{\delta_K(1 + \gamma)}{\gamma r_{KS}(1 - \varepsilon)}\right)$

Stability analysis shows that the origin is a *nodal source*, the two boundary equilibria are both *nodal sinks*, and the interior fixed point is a *hyperbolic saddle* if

$$\delta_K < r_{KS}(1 - \varepsilon), \quad \delta_S < 1, \quad \delta_S < \frac{\delta_K}{r_{KS}(1 - \varepsilon)} < \frac{\delta_S + \gamma}{1 + \gamma}. \quad (\text{S2.3})$$

For this range of parameters, the system exhibits “bi-stability”: starting above the saddle’s stable manifold (the magenta dash-dotted line in Fig. S3A) leads to competitive exclusion of the killer by the sensitive, while starting below it results in competitive exclusion of the sensitive by the killer. See Fig. S3A for a detailed phase portrait. It is worth noting that, when the natural death rates approach zero (i.e., as $\delta_S, \delta_K \rightarrow 0$), Eq. (S2.2) reduces to Eq. (S2.1). In the meantime,

$$\begin{aligned} 1 - \frac{\delta_K}{r_{KS}(1 - \varepsilon)} &\rightarrow 1, \\ 1 - \delta_S &\rightarrow 1, \\ \frac{\delta_K}{\gamma r_{KS}(1 - \varepsilon)} - \frac{\delta_S}{\gamma} &\rightarrow 0, \\ \frac{\gamma + \delta_S}{\gamma} - \frac{\delta_K(\gamma + 1)}{\gamma r_{KS}(1 - \varepsilon)} &\rightarrow 1. \end{aligned}$$

Consequently, the last two equilibria collapse to $(n_K, n_S) = (0, 1)$ while the second fixed point moves to $(n_K, n_S) = (1, 0)$. This means the system loses its “bi-stability” structure as the “sensitive-winning” region shrinks when the natural death rates

approach zero (see the transitions in Fig. S3). Schink et al. [4] estimated the death rate of an *E. coli* strain K-12 at 0.018 h^{-1} in laboratory experiments, whereas the max growth rate was 0.7 h^{-1} . In our time scale, this corresponds to $\delta_K = \delta_S \approx 0.0257$. We see from Fig. S3C that in this scenario, the hyperbolic saddle is extremely close to the “all sensitives” equilibrium and both of them are close to $(n_K, n_S) = (0, 1)$. This extremely small basin of attraction for the “all sensitives” equilibrium makes it harder to explain why sensitive strains are often found in the natural environment. This issue is one of the motivations for the current paper, and for our conjecture that, with dilutions, the sensitives can win even in the $\delta_K = \delta_S = 0$ limit, where the bi-stability disappears; Fig. S3D.

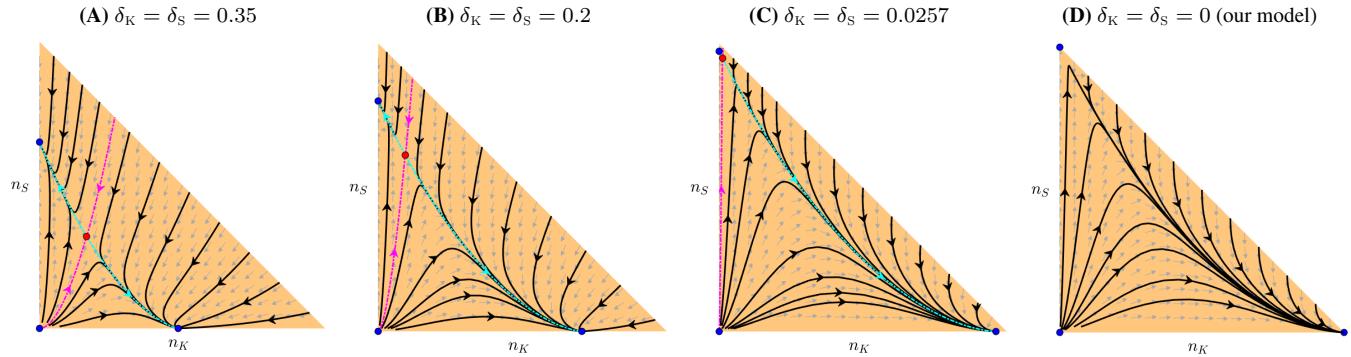


Figure S3. Phase portraits of the Durrett-Levin model (S2.2) with decreasing death rates. The “bi-stability” is noticeable when the death rates (δ_K, δ_S) are comparable in magnitude to the growth rates (panels (A,B)). Using laboratory-estimated death rates values [4], the hyperbolic saddle moves toward the “all sensitives” stable node, with both converging toward $(n_K, n_S) = (1, 0)$ (panel C). When $(\delta_K, \delta_S) = (0, 0)$, the Durrett-Levin model reduces to our model (S2.1), and “bi-stability” disappears (panel D). In (A-C), the hyperbolic saddle is plotted with a red dot while other equilibria are plotted with a blue dot. The stable manifold is plotted with a magenta dotted-dashed line while the unstable manifold is plotted with a cyan dotted-dashed line. In (D), all equilibria are plotted with a blue dot. In all of them, gray arrows denote the vector field directions corresponding to Eq. (S2.2). Parameter values: $r_{KS} = 0.85$, $\varepsilon = 0.2$, and $\gamma = 1$.

S3 Effect of dilutions on a single-strain logistic growth model

This section presents theoretical and numerical results relevant for a *single* strain subjected to either regular or randomly timed dilutions. These findings provide background information for the main text and are relevant after one of the strains becomes strongly dominant.

S3.1 Effect of regular dilutions.

In this subsection, we prove the theoretical pre-dilution population limit for a population growing according to the logistic model and undergoing regular dilutions.

Theorem S3.1. *Consider the following rescaled logistic growth model*

$$\dot{q} = rq(1 - q), \quad q(0) = x \in (0, 1]. \quad (\text{S3.1.1})$$

If the system undergoes regular dilutions with a fundamental period of T , and after each dilution, a deterministic fraction ρ of the population survives, i.e.,

$$q((mT)^+) = \rho q((mT)^-), \quad \text{for } m = 1, 2, 3, \dots, \quad (\text{S3.1.2})$$

then

$$\lim_{m \rightarrow \infty} q((mT)^-) = \begin{cases} \frac{1}{\rho + \frac{1-\rho}{1-\exp(-(rT+\ln\rho))}}, & \text{if } rT + \ln \rho > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S3.1.3})$$

Proof. We begin the proof of Theorem S3.1 by proving the following lemma.

Lemma S3.2. *Given the rescaled logistic model Eq. (S3.1.1), the pre-dilution population size up to the m -th cycle is*

$$\begin{cases} q((mT)^-) = \frac{1}{\rho + (1-\rho) \sum_{k=0}^{m-1} \exp(-(rT+\ln\rho)k) + \frac{1-x}{x} \exp(-(rT+\ln\rho)(m-1)-rT)}, \\ q(0) = x \in (0, 1]. \end{cases} \quad (\text{S3.1.4})$$

Proof. We prove the above lemma by mathematical induction. Starting with the first dilution ($m = 1$), from Eq. (S3.1.4) we have

$$q(T^-) = \frac{1}{\rho + (1-\rho) + \frac{1-x}{x} \exp(-rT)} = \frac{1}{1 + \frac{1-x}{x} \exp(-rT)}, \quad q(0) = x \in (0, 1]. \quad (\text{S3.1.5})$$

We show it is true by solving Eq. (S3.1.1) analytically. By separation of variables, one can show that the general solution to Eq. (S3.1.1) is

$$q(t) = \frac{1}{1 + C e^{-rt}}. \quad (\text{S3.1.6})$$

With $q(0) = x$, we have $C = \frac{1-x}{x}$. Consequently,

$$q(T^-) = \frac{1}{1 + \frac{1-x}{x} \exp(-rT)}, \quad q(0) = x.$$

Now assume Eq. (S3.1.4) holds true for some integer $j > 1$. I.e.,

$$q((jT)^-) = \frac{1}{\rho + (1-\rho) \sum_{k=0}^{j-1} \exp(-(rT+\ln\rho)k) + \frac{1-x}{x} \exp(-(rT+\ln\rho)(j-1)-rT)}, \quad q(0) = x.$$

We now show it holds for $j + 1$. Notice that $q([(j+1)T]^-)$ with $q(0) = x$ is equivalent to $q(T^-)$ with $q(0) = q((jT)^+) = \rho q((jT)^-)$ under the assumption of proportional dilutions. Given the general solution to Eq. (S3.1.1) in Eq. (S3.1.6), we

compute the arbitrary constant C by imposing the new initial condition $q(0) = \rho q((jT)^-)$. It follows that

$$\begin{aligned} \frac{1}{1+C} = q(0) &= \frac{\rho}{\rho + (1-\rho) \sum_{k=0}^{j-1} \exp(-(rT + \ln \rho)k) + \frac{1-x}{x} \exp(-(rT + \ln \rho)(j-1) - rT)} \\ &= \frac{1}{1 + (1-\rho) \sum_{k=0}^{j-1} \frac{\exp(-(rT + \ln \rho)k)}{\rho} + \frac{1-x}{x} \frac{\exp(-(rT + \ln \rho)(j-1) - rT)}{\rho}} \\ &= \frac{1}{1 + (1-\rho) \sum_{k=0}^{j-1} \frac{\exp(-(rT + \ln \rho)k)}{\exp(\ln \rho)} + \frac{1-x}{x} \frac{\exp(-(rT + \ln \rho)(j-1) - rT)}{\exp(\ln \rho)}} \\ C &= (1-\rho) \sum_{k=0}^{j-1} \frac{\exp(-(rT + \ln \rho)k)}{\exp(\ln \rho)} + \frac{1-x}{x} \frac{\exp(-(rT + \ln \rho)(j-1) - rT)}{\exp(\ln \rho)}. \end{aligned}$$

Substituting it back into Eq. (S3.1.6), we have

$$\begin{aligned} q(T^-) &= \frac{1}{1 + \left[(1-\rho) \sum_{k=0}^{j-1} \frac{\exp(-(rT + \ln \rho)k)}{\exp(\ln \rho)} + \frac{1-x}{x} \frac{\exp(-(rT + \ln \rho)(j-1) - rT)}{\exp(\ln \rho)} \right] \exp(-rT)} \\ &= \frac{1}{1 + (1-\rho) \sum_{k=0}^{j-1} \frac{\exp(-(rT + \ln \rho)k)}{\exp(rT + \ln \rho)} + \frac{1-x}{x} \frac{\exp(-(rT + \ln \rho)(j-1) - rT)}{\exp(rT + \ln \rho)}} \\ &= \frac{1}{1 + (1-\rho) \sum_{k=0}^{j-1} \exp(-(rT + \ln \rho)(k+1)) + \frac{1-x}{x} \exp(-(rT + \ln \rho)j - rT)} \\ &= \frac{1}{[\rho + (1-\rho)] + (1-\rho) \sum_{k=0}^{j-1} \exp(-(rT + \ln \rho)(k+1)) + \frac{1-x}{x} \exp(-(rT + \ln \rho)j - rT)} \\ &= \frac{1}{\rho + (1-\rho) \left[1 + \sum_{k=0}^{j-1} \exp(-(rT + \ln \rho)(k+1)) \right] + \frac{1-x}{x} \exp(-(rT + \ln \rho)j - rT)} \\ &= \frac{1}{\rho + (1-\rho) \sum_{k=0}^j \exp(-(rT + \ln \rho)k) + \frac{1-x}{x} \exp(-(rT + \ln \rho)j - rT)}, \\ &= q([(j+1)T]^-), \quad q(0) = x. \end{aligned}$$

Consequently, Eq. (S3.2) holds for all $j \in \mathbb{N}$. □

Now, we take the limit of Eq. (S3.2) as m approaches infinity. Notice that the second term in the denominator is a geometric series. It follows that

$$\lim_{m \rightarrow \infty} (1-\rho) \sum_{k=0}^{m-1} \exp(-(rT + \ln \rho)k) = \begin{cases} \frac{1-\rho}{1 - \exp(-(rT + \ln \rho))}, & \text{if } \rho > e^{-rT}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (\text{S3.1.7})$$

With the third term in the denominator of Eq. (S3.1.4) approaching 0 as $m \rightarrow \infty$, we thus conclude that

$$\lim_{m \rightarrow \infty} q((mT)^-) = \begin{cases} \frac{1}{\rho + \frac{1-\rho}{1 - \exp(-(rT + \ln \rho))}}, & \text{if } \rho > e^{-rT}, \\ 0, & \text{otherwise.} \end{cases}$$

□

Therefore, for a population consisting of sensitives only, i.e.,

$$\dot{n}_S = n_S(1 - n_S),$$

the limiting pre-dilution population size is

$$\lim_{m \rightarrow \infty} n_S((mT)^-) = \begin{cases} \frac{1}{\rho + \frac{1-\rho}{1 - \exp(-(T + \ln \rho))}}, & \text{if } \rho > e^{-T}, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S3.1.8})$$

For a population consisting of toxin-producing killers only, i.e.,

$$\dot{n}_K = r_{KS}(1 - \varepsilon)n_K(1 - n_K),$$

the limiting pre-dilution population size is

$$\lim_{m \rightarrow \infty} n_K((mT)^-) = \begin{cases} \frac{1}{\rho + \frac{1-\rho}{1 - \exp(-(r_{KS}(1 - \varepsilon)T + \ln \rho))}}, & \text{if } \rho > e^{-r_{KS}(1 - \varepsilon)T}, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S3.1.9})$$

Fig. S4 shows the numerical values derived in Eq. (S3.1.8) and Eq. (S3.1.9) in the (ρ, T) phase plane. It is evident that the larger the proportion of the population surviving after dilution (higher ρ) and the longer the period (higher T), the larger the pre-dilution population size as the number of dilutions approaches infinity. For non-zero limits, sensitives can maintain a higher limiting population than any killers due to their faster reproduction rate.

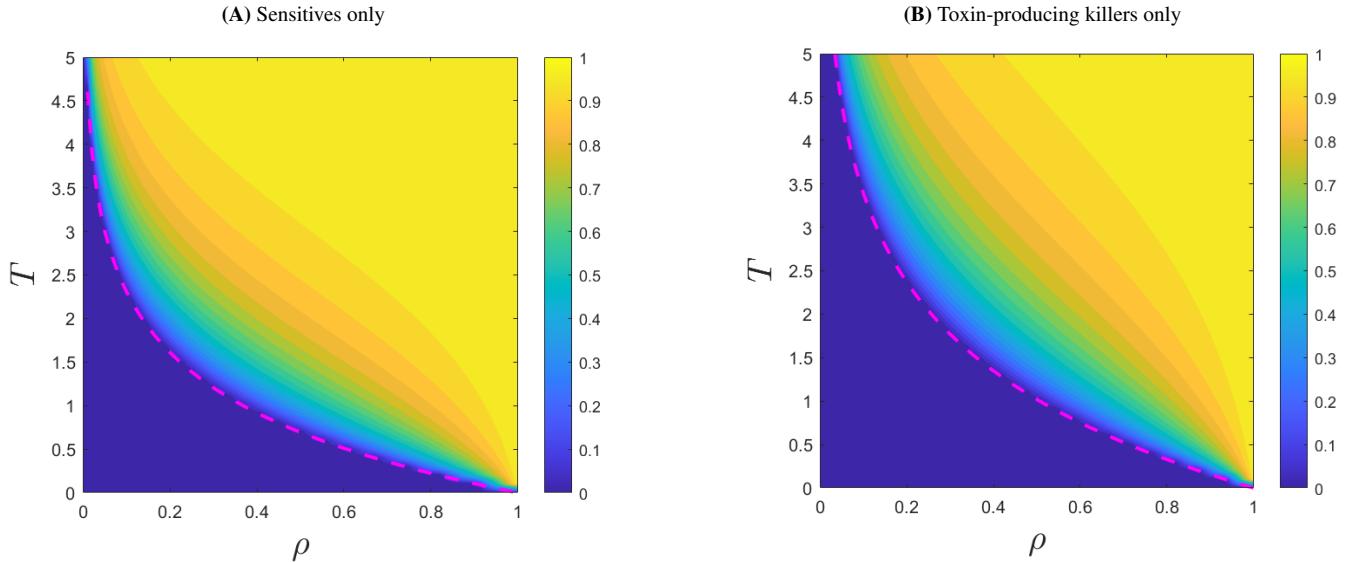


Figure S4. Pre-dilution single population limits under regular dilutions in (ρ, T) phase plane. (A) sensitives only ($r = 1$); (B) Toxin-producing killers only ($r = r_{KS}(1 - \varepsilon) = 0.68$). In both of them, the magenta-dashed line corresponds to $\rho = \exp(-rT)$ with their respective intrinsic growth rate r . The limiting pre-dilution population is zero below this line.

S3.2 Effect of randomly-timed dilutions.

Next, we consider randomly timed dilutions, following a Poisson process with rate λ . In this case, any inter-dilution time T is exponentially distributed with the expected value $\mathbb{E}[T] = 1/\lambda$. To compare with the previous results, we focus on the *expected pre-dilution population size* as the number of dilutions approaches infinity, with the population growing according to Eq. (S3.1.1). We still assume that a fixed fraction ρ of the population survives after each dilution. As an analytical limit is unlikely to be obtained, we conducted Monte Carlo simulations with 200 dilutions starting from the initial population $q(0) = 0.5$. Fig. S5 shows that the general trend remains unchanged: weaker dilution strength (higher ρ) and slower arrival rate (lower λ) result in higher expected pre-dilution population size in the limit. However, with randomly timed dilutions, the region where the averaged population is nearly 1 in the limit is significantly reduced in both cases. Surprisingly, the previous “deterministic boundary” $\rho = \exp(-r/\lambda)$ (the magenta dashed line in Fig. S5) still appears to accurately predict where the population goes extinct in the limit in the $(\rho, 1/\lambda)$ phase plane.

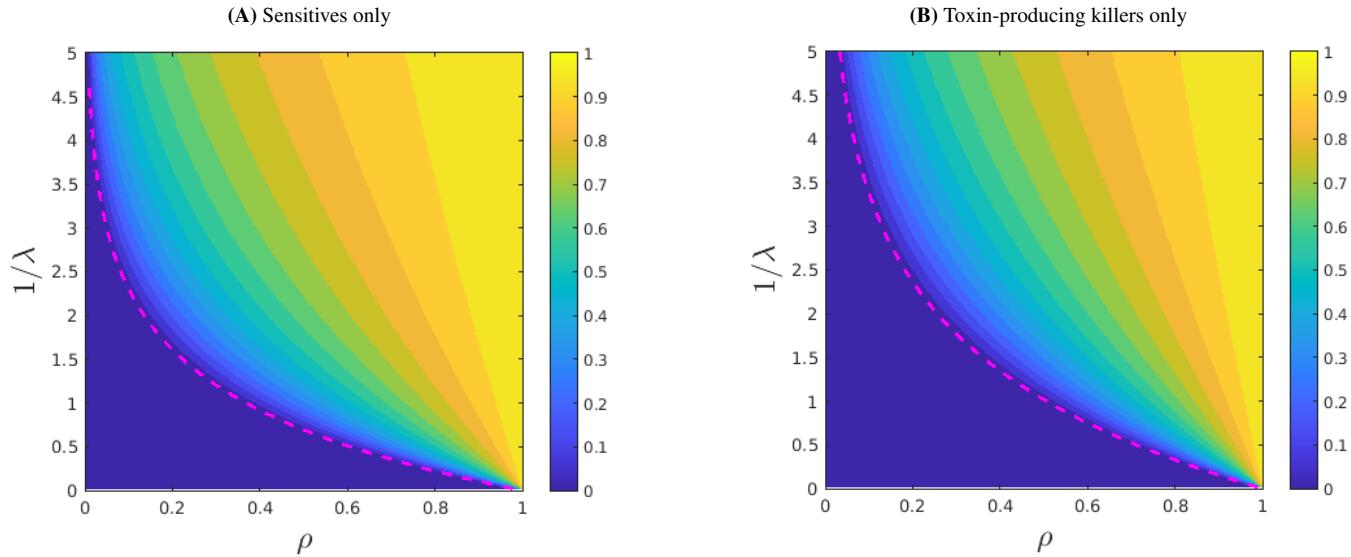


Figure S5. Randomly-timed dilutions: mean empirical population just before the 201st dilution shown in the $(\rho, 1/\lambda)$ phase plane. (A) sensitives only ($r = 1$); (B) Toxin-producing killers only ($r = r_{KS}(1 - \epsilon) = 0.68$). In both panels, the magenta-dashed line corresponds to $\rho = \exp(-r/\lambda)$ with their respective intrinsic growth rate r . The population below this line will most likely go extinct. Both panels are produced by Monte Carlo simulations on a uniform grid with gridpoints $(\rho_i, (1/\lambda)_j) = (i/100, 0.01 + j/20)$ where $i, j = 0, \dots, 100$. At each gridpoint (representing a particular environment), we used 10^5 random population trajectories starting from the same initial population size of 0.5.

S4 Who benefits from randomness in dilution times?

To check whether the randomness in dilution times and the values of the surviving fraction ρ and dilution frequency λ affect the performance of constitutive and myopic/tactically-optimal toxin production policies, we introduce a new metric of competitive advantage and use it across a range of (ρ, λ) values. Assuming that the initial population size N_0 is fixed while $f_0 \in [0, 1]$ is selected uniformly at random, we examine the probability of killers’ winning. For regular/deterministically-timed dilutions with period $T = 1/\lambda$, this probability (denoted by $\mathcal{L}(N_0)$) is simply the width of the killer’s “deterministically-winning” (red) region on the horizontal line $N = N_0$, e.g., in Fig. 4C-D. With random dilution times, this probability is $\mathcal{P}(N_0) = \int_0^1 \hat{w}^\alpha(x, N_0) dx$. To quantify the impact of randomness, we define $\bar{Q}(N_0) = \mathcal{L}(N_0) - \mathcal{P}(N_0)$. Focusing on $N_0 = 0.5$, Fig. S6 presents a heat map of $\bar{Q}(0.5)$ for various (ρ, λ) values. Panel A shows that, for constitutive killers, the randomness is beneficial in the upper left half of the parameter space (where dilutions are more severe and frequent) but actually slightly detrimental in the bottom right half (where dilutions are more moderate and rare). In contrast, for the population-sensing (myopic, optimized for a specific T or λ) killers, the randomness in dilution times appears to be beneficial across all (ρ, λ) values.

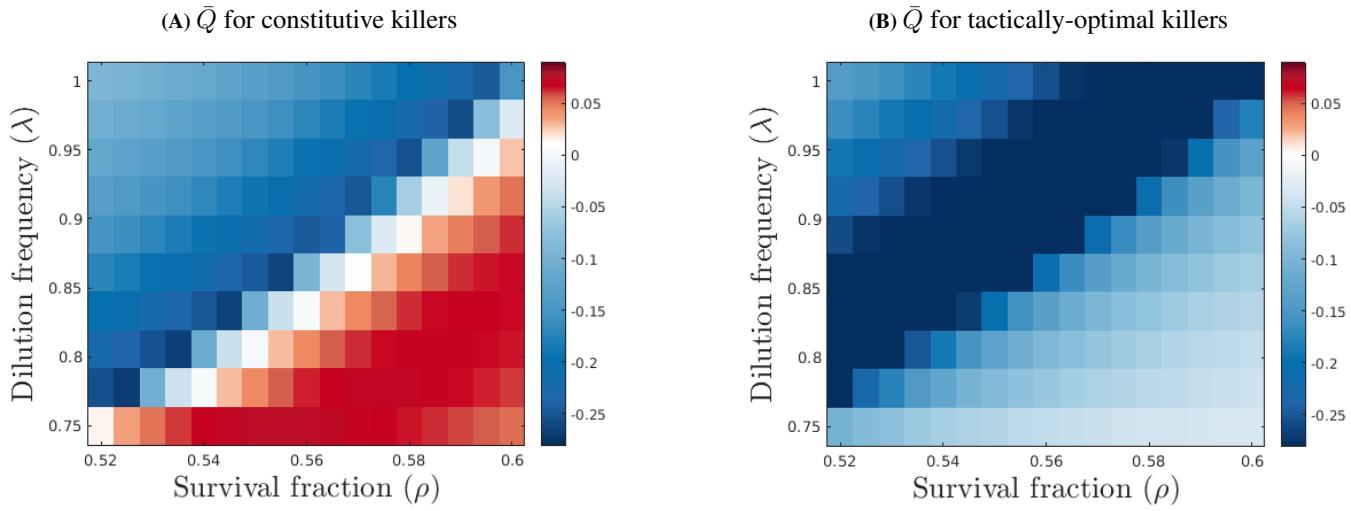


Figure S6. The impact of dilution-time randomness and intensity on the performance of constitutive and tactically-optimal killers across a range of (ρ, λ) values. The performance metric \bar{Q} (defined in the text) is shown in red wherever the toxin-producers have better chances of winning with regular/periodic dilutions and in blue wherever their chances are better with randomly-timed dilutions (assuming the same average frequency: $\lambda = 1/T$). (A) For constitutive killers, the randomness is beneficial when dilutions are more severe and frequent, but it is slightly detrimental when dilutions happen more rarely and are less drastic. (B) For tactically-optimal killers (with policies optimized for each λ and $T = 1/\lambda$), the randomness is beneficial across all tested parameters.

S5 Comparison of strategically and tactically-optimal killers and constitutive ones

Fig. S7B shows that the strategically-optimal policy α_∞ prescribes producing toxin slightly more conservatively than the tactically-optimal α_1 , but in the end its performance is only marginally better for our chosen parameter values (Fig. S7D). However, both of these population-sensing-enabled policies have a very significant advantage over the constitutive $\alpha_0 = 1$; see Fig. S7C and Fig. 8 in the main text.

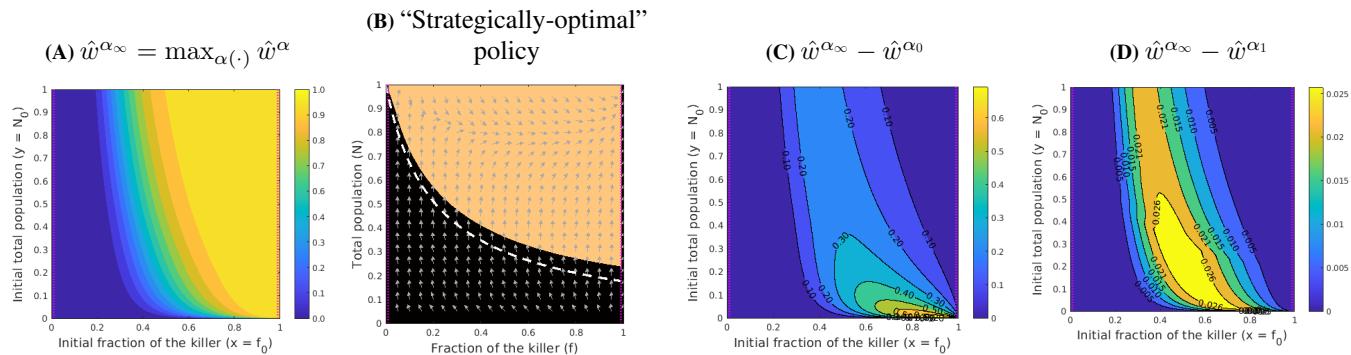


Figure S7. “Strategically-optimal” killers: performance (A), policy (B), and comparison with constitutive (C) and tactically-optimal (D) killers. The optimal toxin-on region for the “strategically-optimal” killers (orange in B) slightly shrinks compared to that of the “tactically-optimal” killers (the boundary of which is shown by a white-dashed line). As a result, the maximized probability of winning (\hat{w}^{α_∞} in panel A) is only marginally better than \hat{w}^{α_1} , with a maximum difference of just 0.025 (see the absolute difference map in D). However, compared to constitutive killers, the advantage is significant: on a large part of the domain, the improvement in the chances of winning is above 20% (panel C). For really small N and relatively large f , this advantage is even above 60% – this is the set of initial conditions where constitutive killers grow much slower and are thus more affected by occasional short inter-dilutions intervals. In all panels, the victory and defeat barriers (γ_v and γ_d , respectively) are plotted with a magenta dotted line. In (B), gray arrows denote the vector field directions corresponding to Eq. (2) in the main text with $a = \alpha_\infty(f, N)$. In both (C) & (D), the contour lines are labeled with their respective probability values.

S6 Derivation of Hamilton-Jacobi-Bellman equations

In this section, we derive the HJB equations for “tactically-optimal” killers (Eq. (11) in Box 4) and the “strategically-optimal” killers (Eq. (14) in Box 4) via tools of dynamic programming. For the former, see also [5, 6] for more details. The derivation of the time-dependent HJB PDE for the finite-horizon problem (Eq. (3) in the main text) is omitted here, as it can be easily found in classical literature, such as [7].

Recall that we model the random dilution events as a Poisson process with a fixed rate $\lambda > 0$. Thus, any inter-arrival time T is exponentially distributed with rate λ . We assume that after each dilution, the fractions (relative abundances) are preserved while only a fraction ρ of the total population survives.

S6.1 “Tactically-optimal” killers.

Recall the value function for the “tactically-optimal” killer in Box 4 in the main text:

$$v(x, y) = \sup_{a(\cdot)} \mathbb{E}[f(\mathcal{T}^-) \mid f(0) = x, N(0) = y, \text{ following policy } a(\cdot)].$$

Since $\mathcal{T} \sim \text{Exp}(\lambda)$, the expectation is defined as

$$v(x, y) = \sup_{a(\cdot)} \int_0^\infty \lambda e^{-\lambda t} f(t; a(t)) dt = \int_0^\infty \lambda e^{-\lambda t} f(t; a_*(t)) dt,$$

with an assumption that a maximizing/optimal policy $a_*(\cdot)$ exists. Rewriting this formula as

$$v(x, y) = \int_0^\infty e^{-\lambda t} [\lambda f(t; a_*(t))] dt,$$

we can now re-interpret it as an *infinite horizon* problem with running cost $[\lambda f(t; a_*(t))]$ and a discounting factor λ .

For a sufficiently small $h > 0$, by Bellman’s Optimality Principle, we have

$$\begin{aligned} v(x, y) &= \sup_{a(\cdot)} \left\{ \int_0^h e^{-\lambda t} [\lambda f(t; a(t))] dt + e^{-\lambda h} v(f(h; a(h)), N(h; a(h))) \right\} \\ &= \int_0^h e^{-\lambda t} [\lambda f(t; a_*(t))] dt + e^{-\lambda h} v(f(h; a_*(h)), N(h; a_*(h))) \\ &= h[\lambda x] + (1 - \lambda h)[v(x, y) + h v_x(x, y) \cdot F(x, y, a_*(0)) + h v_y(x, y) \cdot G(x, y, a_*(0))] + o(h). \\ 0 &= h \lambda x - h \lambda v(x, y) + h v_x(x, y) \cdot F(x, y, a_*(0)) + h v_y(x, y) \cdot G(x, y, a_*(0)) + o(h). \end{aligned}$$

Now dividing both sides by h and sending h to 0, we obtain

$$0 = \lambda(x - v(x, y)) + v_x(x, y) \cdot F(x, y, a_*(0)) + v_y(x, y) \cdot G(x, y, a_*(0)),$$

or, more generally,

$$0 = \lambda(x - v(x, y)) + \sup_{a(\cdot)} \{v_x(x, y) \cdot F(x, y, a(0)) + v_y(x, y) \cdot G(x, y, a(0))\}.$$

Notice that the above equation involves $a_*(0)$ only. It is then natural to switch to a state-dependent optimal control in feedback form. The HJB equation that v satisfies is then obtained by maximizing over $a = a(0) \in [0, 1]$. By demanding the above equation holds for all $(x, y) \in [0, 1]^2$, the PDE can be written as:

$$0 = \lambda(x - v(x, y)) + \max_{a \in [0, 1]} \{v_x(x, y) \cdot F(x, y, a) + v_y(x, y) \cdot G(x, y, a)\}, \quad (\text{S6.1.1})$$

with the boundary condition

$$v(x, y) = \begin{cases} 1, & \text{if } x = 1 \text{ and } y \neq 0, \\ 0, & \text{if } y = 0 \text{ or } x = 0. \end{cases} \quad (\text{S6.1.2})$$

Substituting the actual definitions of F and G , we obtain Eq. (S6.1.1) in the specific form:

$$\begin{aligned} 0 = & \lambda[x - v(x, y)] + \max_{a \in [0, 1]} \left\{ \left(\nabla v(x, y) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} \right) a \right. \\ & \left. + \nabla v(x, y) \cdot \begin{bmatrix} x(1-x)(1-y)(r_{\text{KS}} - 1) \\ y(1-y)[1 + (r_{\text{KS}} - 1)x] \end{bmatrix} \right\}. \end{aligned} \quad (\text{S6.1.3})$$

The linear dependence on a yields the *bang-bang* property:

$$\alpha_1(x, y) = a_*(x, y) = \begin{cases} 1, & \text{if } \nabla v(x, y) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S6.1.4})$$

S6.2 “Strategically-optimal” killers.

Recall the corresponding value function Eq. (13) from the main text:

$$w(x, y) = \mathbb{P}\left(T_v(x, y, a_*(\cdot)) < T_d(x, y, a_*(\cdot))\right). \quad (\text{S6.2.1})$$

Let $\{\tau_i\}_{i=1}^{\infty}$ be an infinite sequence of random dilution times and $t_m = m\Delta t$, $m = 0, 1, 2, 3, \dots$ be a uniform time discretization. We again assume the optimal policy a_* exists, and let

$$\mathbb{E}^0 = \mathbb{E}[\cdot \mid f(0) = x, N(0) = y, \text{ following } a_*].$$

Thus, by Bellman’s Optimality Principle and the *law of total expectation*, we have

$$\begin{aligned} w(x, y) &= \mathbb{E}^0\left[w(f(t_1), N(t_1)) \mid \tau_1 \leq t_1\right] \mathbb{P}(\tau_1 \leq t_1) + \mathbb{E}^0\left[w(f(t_1), N(t_1)) \mid \tau_1 > t_1\right] \mathbb{P}(\tau_1 > t_1) \\ &= \left(1 - e^{-\lambda\Delta t}\right) w\left(f(\Delta t), \rho N(\Delta t)\right) + e^{-\lambda\Delta t} w\left(f(\Delta t), N(\Delta t)\right) + o(\Delta t). \end{aligned}$$

A first-order approximation around $t_0 = 0$ gives

$$\begin{aligned} w\left(f(\Delta t), N(\Delta t)\right) &= w(x, y) + w_x(x, y) \cdot F(x, y, a_*(0))\Delta t + w_y(x, y) \cdot G(x, y, a_*(0))\Delta t + O(\Delta t^2) \\ w\left(f(\Delta t), \rho N(\Delta t)\right) &= w(x, \rho y) + O(\Delta t). \end{aligned}$$

And hence

$$\begin{aligned} \underline{w(x, y)} &= \left(1 - e^{-\lambda\Delta t}\right) w\left(f(\Delta t), \rho N(\Delta t)\right) + e^{-\lambda\Delta t} w\left(f(\Delta t), N(\Delta t)\right) + o(\Delta t) \\ &= (\lambda\Delta t) \left[w(x, \rho y) + O(\Delta t) \right] \\ &\quad + (1 - \lambda\Delta t) \{w(x, y) + w_x(x, y) \cdot F(x, y, a_*(0))\Delta t + w_y(x, y) \cdot G(x, y, a_*(0))\Delta t\} + O(\Delta t^2) \\ &= \lambda\Delta t w(x, \rho y) + \underline{w(x, y)} - \lambda\Delta t w(x, y) \\ &\quad + w_x(x, y) \cdot F(x, y, a_*(0))\Delta t + w_y(x, y) \cdot G(x, y, a_*(0))\Delta t + O(\Delta t^2). \end{aligned}$$

Dividing it by Δt and taking $\Delta t \downarrow 0$, we have

$$0 = \lambda [w(x, \rho y) - w(x, y)] + w_x(x, y) \cdot F(x, y, a_*(0)) + w_y(x, y) \cdot G(x, y, a_*(0)).$$

Notice that the above equation involves $a_*(0)$ only. It is then natural to switch to a state-dependent optimal control in feedback form. The HJB equation for Eq. (13) is then obtained by maximizing over $a = a(0) \in [0, 1]$. By demanding the above equation holds for all $(x, y) \in [0, 1]^2 \setminus \Delta$ (where Δ is the terminal set, see Box 3), the PDE can be written as:

$$0 = \lambda [w(x, \rho y) - w(x, y)] + \max_{a \in [0, 1]} \left\{ w_x(x, y) \cdot F(x, y, a) + w_y(x, y) \cdot G(x, y, a) \right\}, \quad (\text{S6.2.2})$$

with the boundary condition

$$w(x, y) = \begin{cases} 1, & \text{if } (x, y) \in \Delta_v \text{ and } y \neq 0, \\ 0, & \text{if } (x, y) \in \Delta_d \text{ or } y = 0. \end{cases} \quad (\text{S6.2.3})$$

Substituting the actual definitions of F and G , we obtain Eq. (S6.2.2) in the specific form:

$$\begin{aligned} 0 = \lambda [w(x, \rho y) - w(x, y)] + \max_{a \in [0, 1]} &\left\{ \left(\nabla w(x, y) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} \right) a \right. \\ &\left. + \nabla w(x, y) \cdot \begin{bmatrix} x(1-x)(1-y)(r_{\text{KS}} - 1) \\ y(1-y)[1 + (r_{\text{KS}} - 1)x] \end{bmatrix} \right\}. \end{aligned} \quad (\text{S6.2.4})$$

The linear dependence on a yields the *bang-bang* property:

$$\alpha_\infty(x, y) = a_*(x, y) = \begin{cases} 1, & \text{if } \nabla w(x, y) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S6.2.5})$$

As mentioned in the main text, our system dynamics with randomly timed dilutions can be interpreted as a Piecewise-Deterministic Markov Process (PDMP). In general, the value function associated with a PDMP might not be smooth or even continuous. However, it can still often be interpreted as a unique (discontinuous) *viscosity solution* of the HJB equation [8].

Remark I: The linear non-local Eq. (9) in Box 3 can be derived in the same way but using a fixed policy \hat{a} instead of a_* .

Remark II: Let

$$\vec{T}(x, y) := \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix}.$$

For both the “tactically-optimal” killer and the “strategically-optimal” killer, the respective HJB equation is linear in a , which yields a *generally* bang-bang optimal policy. However, we note that singular controls may arise when either $\nabla v(x, y) \cdot \vec{T}(x, y)$ or $\nabla w(x, y) \cdot \vec{T}(x, y)$ is equal to 0. But by computing the vector field associated with both α_1 and α_∞ , we find no vector tangential to the boundary of $a_*(x, y) = 1$, thereby excluding the possibility of singular arcs in our models. The same situation applies to the optimal policy $a_*(x, y, t)$ for the regular/period dilutions as well.

S7 Numerical methods and implementation details

In this section, we provide the numerical schemes and implementation details of solving: (i) the finite horizon HJB PDE for $u(x, y, t)$ with regular/periodic dilutions in §S7.1; (ii) the limiting fractions \hat{u}^∞ and total population $\hat{\phi}^\infty$ under regular dilutions in §S7.2; and (iii) the non-local HJB equation for $w(x, y)$ for the “strategically-optimal” killers in §S7.3. For (i) and (iii), the optimal feedback policy is found by numerically solving the corresponding HJB equation.

S7.1 For the finite-horizon HJB.

Recall from the main text that we define the fraction-maximizing *value function* as

$$u(x, y, t) = \sup_{a(\cdot)} f(T^-), \quad \text{with } f(t) = x, \quad N(t) = y, \quad t \in [0, T], \quad (\text{S7.1.1})$$

where u satisfies a time-dependent HJB PDE

$$-\frac{\partial u}{\partial t}(x, y, t) = \max_{a \in [0, 1]} \left\{ \nabla u(x, y, t) \cdot \begin{bmatrix} F(x, y, a) \\ G(x, y, a) \end{bmatrix} \right\} \quad (\text{S7.1.2})$$

with the terminal condition

$$u(x, y, T^-) = \begin{cases} x, & \text{if } y > 0, \\ 0, & \text{if } y = 0. \end{cases}$$

Substituting the actual definitions of F and G , we obtain Eq. (S7.1.2) in the specific form:

$$\begin{aligned} -\frac{\partial u}{\partial t}(x, y, t) = \max_{a \in [0, 1]} & \left\{ \left(\nabla u(x, y, t) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} \right) a \right. \\ & \left. + \nabla u(x, y, t) \cdot \begin{bmatrix} x(1-x)(1-y)(r_{\text{KS}} - 1) \\ y(1-y)[1 + (r_{\text{KS}} - 1)x] \end{bmatrix} \right\}. \end{aligned} \quad (\text{S7.1.3})$$

The linear dependence on a yields the *bang-bang* property:

$$\alpha(x, y, t) = a_*(x, y, t) = \begin{cases} 1, & \text{if } \nabla u(x, y, t) \cdot \begin{bmatrix} x(1-x)[\gamma xy - \varepsilon r_{\text{KS}}(1-y)] \\ -xy[\gamma y(1-x) + \varepsilon r_{\text{KS}}(1-y)] \end{bmatrix} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S7.1.4})$$

The time-dependent total population, denoted by $\phi(x, y, t)$, satisfies a linear PDE

$$-\frac{\partial \phi}{\partial t}(x, y, t) = \nabla \phi(x, y, t) \cdot \begin{bmatrix} F(x, y, \alpha(x, y, t)) \\ G(x, y, \alpha(x, y, t)) \end{bmatrix} \quad (\text{S7.1.5})$$

with the terminal condition $\phi(x, y, T^-) = y$.

We approximate the solution to Eq. (S7.1.2) by a first-order semi-Lagrangian discretization [9] on a uniform rectangular grid over the (x, y, t) space. I.e., $(x_i, y_j, t_k) = (i\Delta x, j\Delta y, k\Delta t)$, where $\Delta x = 1/M_x$, $\Delta y = 1/M_y$, $\Delta t = T/M_t$, while $i = 0, \dots, M_x$, $j = 0, \dots, M_y$, and $k = 0, \dots, M_t$. We further simplify the notation for the spatial part as $\Xi = \{(i\Delta x, j\Delta y) \mid i = 0, \dots, M_x, j = 0, \dots, M_y\}$. We will use $U_{i,j}^k \approx u(x_i, y_j, t_k)$ to denote the discretized approximation at (x_i, y_j, t_k) , and similarly, $\Phi_{i,j}^k \approx \phi(x_i, y_j, t_k)$, $A_{i,j}^k \approx \alpha(x_i, y_j, t_k)$.

For a sufficiently small $\Delta t > 0$, a first-order approximation of $(f((k+1)\Delta t; a), N((k+1)\Delta t; a))$ starting from $(f(k\Delta t), N(k\Delta t)) = (x_i, y_j)$ for any $k \in \{0, \dots, M_t\}$, with a control value $a \in \{0, 1\}$ is

$$\begin{aligned}\tilde{f}_{i,a} &= x_i + \Delta t * F(x_i, y_j, a), \\ \tilde{N}_{j,a} &= y_j + \Delta t * G(x_i, y_j, a).\end{aligned}$$

Let $\tilde{U}_{i,j,a}^{k+1} \approx u(\tilde{f}_{i,a}, \tilde{N}_{j,a}, t_{k+1})$. Since Eq. (S7.1.1) is in a Mayer form [7], the discretized dynamic programming equation is

$$U_{i,j}^k = \max_{a \in \{0,1\}} \tilde{U}_{i,j,a}^{k+1} + o(\Delta t), \quad (\text{S7.1.6})$$

where $\tilde{U}_{i,j,a}^{k+1}$ is evaluated by a bi-linear interpolation using the U values from the 4 neighboring gridpoints surrounding $(\tilde{f}_{i,a}, \tilde{N}_{j,a})$. Note that $A_{i,j}^k$ is found as the argmax of Eq. (S7.1.6) at each gridpoint. This straightforward time-marching scheme is summarized in Algorithm S1. In all of our numerical experiments, we have used $M_x = M_y = 1600$ on each side of the unit fN -square, and $\Delta t = 6.25 \times 10^{-3}$. *To obtain the numerical solution for constitutive killers, one can simply apply Algorithm S1 with $a = 1$ without the maximization.*

Remark III: Despite the time-dependent nature of the problem, one is typically interested in the 0-th time slice of the value function (i.e., $u(x, y, 0)$), which predicts the value at $t = T$ starting from $t = 0$ for all possible initial states. We note that by setting a sufficiently large T in Algorithm S1, any k -th time slice serves as the 0-th slice for a reduced horizon of $T - t_k$. This allows us to obtain the prediction for a range of horizons simultaneously in a single sweep.

Algorithm S1: Finite-horizon value function computation

Initialize U, Φ, A at $t = T$ ($k = M_t$) using the terminal condition;

for $t_k = k\Delta t$, $k = M_t - 1, \dots, 0$ **do**

for every $(x_i, y_j) \in \Xi$ **do**

for $a \in \{0, 1\}$ **do**

$\tilde{f}_{i,a} = x_i + \Delta t * F(x_i, y_j, a);$

$\tilde{N}_{j,a} = y_j + \Delta t * G(x_i, y_j, a);$

$U_{i,j,a}^k \leftarrow u(\tilde{f}_{i,a}, \tilde{N}_{j,a}, t_{k+1})$ by interpolation;

end

$U_{i,j}^k \leftarrow \max_{a \in \{0,1\}} \{U_{i,j,a}^k\};$

$A_{i,j}^k \leftarrow \arg \max_{a \in \{0,1\}} \{U_{i,j,a}^k\};$

end

for every $(x_i, y_j) \in \Xi$ **do**

$\tilde{f}_i = x_i + \Delta t * F(x_i, y_j, A_{i,j}^k);$

$\tilde{N}_j = y_j + \Delta t * G(x_i, y_j, A_{i,j}^k);$

$\Phi_{i,j}^k \leftarrow \phi(\tilde{f}_i, \tilde{N}_j, t_{k+1})$ by interpolation;

end

end

S7.2 Approximating population limits under competitions.

When sensitives and killers compete under regular dilutions, we are interested in whether they can coexist or if one population will dominate the other. As mentioned in the main text, this can be numerically found by a repetitive mapping using $u(x, y, 0)$ and $\phi(x, y, 0)$ computed by Algorithm S1.

Let $\hat{u}(x, y) = u(x, y, 0)$ and $\hat{\phi}(x, y) = \phi(x, y, 0)$. We will use \hat{u}^n and $\hat{\phi}^n$ to denote the relative abundance of the killer, and the normalized total population, respectively, by the end of n -th cycle. Thus, by definition, $\hat{u}^1 = \hat{u}$ and $\hat{\phi}^1 = \hat{\phi}$. Based on our assumption, after the n -th dilution, $f((nT)^+) = f((nT)^-)$ and $N((nT)^+) = \rho N((nT)^-)$. It follows that

$$\begin{aligned}\hat{u}^{n+1}(x, y) &= \hat{u}\left(\hat{u}^n(x, y), \rho\hat{\phi}^n(x, y)\right), \\ \hat{\phi}^{n+1}(x, y) &= \hat{\phi}\left(\hat{u}^n(x, y), \rho\hat{\phi}^n(x, y)\right),\end{aligned}$$

with $f(0) = x$ and $N(0) = y$ and following policy α . We repeat the process until $\|\hat{u}^{n+1} - \hat{u}^n\|$ is small, and output $\hat{u}^\infty \approx \hat{u}^{n+1}$ and $\hat{\phi}^\infty \approx \hat{\phi}^{n+1}$ as the limits. Given our focus on this limit, we describe α as a “myopic” policy. This designation highlights that α is only optimal for a single period. The truly optimal policy for an infinite number of periods would typically adapt from one period to the next. Our full method of approximating \hat{u}^∞ with this “myopic” policy is summarized in Algorithm S2 with the usual notations of solution on the discretized grid: $\hat{U}_{i,j} \approx \hat{u}(x_i, y_j)$, $\hat{\Phi}_{i,j} \approx \hat{\phi}(x_i, y_j)$, $\hat{U}_{i,j}^n \approx \hat{u}^n(x_i, y_j)$, and $\hat{\Phi}_{i,j}^n \approx \hat{\phi}^n(x_i, y_j)$.

Algorithm S2: Limiting performance of the myopic policy

```

Initialize  $\hat{U}$ ,  $\hat{\Phi}$ ,  $\hat{U}^1$ ,  $\hat{\Phi}^1$  using the outputs from Algorithm S1;
 $n = 1$ ;
 $\text{err} = 1\text{e}6$ ;
while  $\text{err} > \text{tol}$  do
  for every  $(x_i, y_j) \in \Xi$  do
     $f_{\text{temp}} = \hat{U}_{i,j}^n$ ;
     $N_{\text{temp}} = \rho * \hat{\Phi}_{i,j}^n$ ;
     $\hat{U}_{i,j}^{n+1} \leftarrow \hat{u}(f_{\text{temp}}, N_{\text{temp}})$  by interpolation;
     $\hat{\Phi}_{i,j}^{n+1} \leftarrow \hat{\phi}(f_{\text{temp}}, N_{\text{temp}})$  by interpolation;
  end
   $\text{err} = \|\hat{U}^{n+1} - \hat{U}^n\|$ ;
   $n \leftarrow n + 1$ ;
end
 $\hat{U}^\infty = \hat{U}^{n+1}$ ;
 $\hat{\Phi}^\infty = \hat{\Phi}^{n+1}$ ;

```

Remark IV: Given our interest in the limit as the number of dilutions approaches infinity, we can significantly accelerate Algorithm S2. Rather than updating just one cycle with \hat{u} and $\hat{\phi}$ per iteration, we can exponentially increase the number of cycles updated at each iteration. Specifically, at the n -th iteration, we can update 2^n cycles by using the results from the previous iteration. Let \tilde{u}^n and $\tilde{\phi}^n$ represent the values at the n -th iteration of this accelerated algorithm, starting with $\tilde{u}^1 = \hat{u}$ and $\tilde{\phi}^1 = \hat{\phi}$, we now have

$$\begin{aligned}\tilde{u}^{n+1}(x, y) &= \tilde{u}^n\left(\tilde{u}^n(x, y), \rho\tilde{\phi}^n(x, y)\right), \\ \tilde{\phi}^{n+1}(x, y) &= \tilde{\phi}^n\left(\tilde{u}^n(x, y), \rho\tilde{\phi}^n(x, y)\right),\end{aligned}$$

with $\tilde{u}^n = \hat{u}^{2^n}$ and $\tilde{\phi}^n = \hat{\phi}^{2^n}$. This accelerated algorithm is summarized in Algorithm S3.

Algorithm S3: Accelerated computation of \hat{u}^∞

```

Initialize  $\tilde{U}^1, \tilde{\Phi}^1$  using the outputs from Algorithm S1;
 $m = 1$  ;
 $\text{err} = 1\text{e}6$  ;
while  $\text{err} > \text{tol}$  do
  for every  $(x_i, y_j) \in \Xi$  do
     $f_{\text{temp}} = \tilde{U}_{i,j}^m$  ;
     $N_{\text{temp}} = \rho * \tilde{\Phi}_{i,j}^m$  ;
     $\tilde{U}_{i,j}^{m+1} \leftarrow \tilde{u}^m(f_{\text{temp}}, N_{\text{temp}})$  by interpolation;
     $\tilde{\Phi}_{i,j}^{m+1} \leftarrow \tilde{\phi}^m(f_{\text{temp}}, N_{\text{temp}})$  by interpolation;
  end
   $\text{err} = \|\tilde{U}^{m+1} - \tilde{U}^m\|$  ;
   $m \leftarrow m + 1$  ;
end
 $\hat{U}^\infty = \tilde{U}^{m+1}$  ;
 $\hat{\Phi}^\infty = \tilde{\Phi}^{m+1}$  ;

```

Remark V: Either Algorithm S2 or Algorithm S3 can be applied to each time-slice in Algorithm S1 to compute the \hat{u}^∞ and $\hat{\phi}^\infty$ for a range of horizons simultaneously in a single sweep. This can be achieved by defining $\hat{u} = u(x, y, t_k)$ and $\hat{\phi} = \phi(x, y, t_k)$ at each k -th slice, and thus obtaining the limits for each reduced horizon $T - t_k$.

S7.3 For the “strategically-optimal” killer.

The HJB PDE (S6.1.1) associated with “tactically-optimal” killers can be numerically computed via standard Value [10–12] (or Value-Policy [13]) Iterations with a semi-Lagrangian discretization [5, 6, 9]. Here, we propose a similar Value-Policy Iterations (VPI) scheme to compute Eq. (S6.2.2) for the “strategically-optimal” killers.

Assuming the same spatial discretization Ξ as in §S7.1, we denote the approximate solution to the value function as $W_{i,j} \approx w(x_i, y_j)$. Similarly to §S7.1 (but now with k always fixed at 0), the foot of the characteristics starting from a gridpoint (x_i, y_j) reaches a new state $(\tilde{f}_{i,a}, \tilde{N}_{j,a})$ for a sufficiently small Δt with control value a .

Therefore, from the Dynamic Programming Principle (DPP), we have

$$w(x_i, y_j) = \max_{a \in \{0,1\}} \left\{ \underbrace{(1 - e^{-\lambda \Delta t}) w(\tilde{f}_{i,a}, \rho \tilde{N}_{j,a})}_{\text{prob of arrival}} + \underbrace{e^{-\lambda \Delta t}}_{\text{prob of not arrival}} w(\tilde{f}_{i,a}, \tilde{N}_{j,a}) \right\} + o(\Delta t) \quad (\text{S7.3.1})$$

yielding the discretized version

$$W_{i,j} = \max_{a \in \{0,1\}} \left\{ (1 - e^{-\lambda \Delta t}) \check{W}_{i,j,a} + e^{-\lambda \Delta t} \tilde{W}_{i,j,a} \right\}, \quad (\text{S7.3.2})$$

where $\check{W}_{i,j,a} \approx w(\tilde{f}_{i,a}, \rho \tilde{N}_{j,a})$ and $w(\tilde{f}_{i,a}, \tilde{N}_{j,a})$ are computed through a *bi-linear* interpolation of the W values from the four neighboring gridpoints surrounding $(\tilde{f}_i, \rho \tilde{N}_j)$ and $(\tilde{f}_i, \tilde{N}_j)$, respectively. The optimal feedback policy $A_{i,j} \approx \alpha(x_i, y_j)$ is recovered as an argmax in Eq. (S7.3.2).

We start with value iterations where we solve the nonlinear Eq. (S6.2.2) by a Gauss-Seidel relaxation. Let $W_{i,j}^n \approx w^n(x_i, y_j)$ and $A_{i,j}^n \approx \alpha^n(x_i, y_j)$ be the discretized solution/policy at the n -th iteration at gridpoint (x_i, y_j) . We use err to denote the L_∞ -norm of W -change in the current value iteration. Whenever err stagnates, we proceed to the “policy-evaluation” (PE) step.

In the PE step, we compute the value function by solving a system of linear equations with a fixed policy $\hat{\alpha}$ (recovered from the most recent value iteration). A first-order approximation of the system Eq. (2) starting from (x_i, y_j) with policy $\hat{A}_{i,j} \approx \hat{\alpha}(x_i, y_j)$ is $(\tilde{f}_{i,\hat{A}}, \tilde{N}_{j,\hat{A}}) = (x_i + \Delta t * F(x_i, y_j, \hat{A}_{i,j}), y_j + \Delta t * G(x_i, y_j, \hat{A}_{i,j}))$. We thus solve a linear system of equations

$$\hat{W}_{i,j}^{\hat{A}} = (1 - e^{-\lambda \Delta t}) \check{W}_{i,j}^{\hat{A}, \rho} + e^{-\lambda \Delta t} \tilde{W}_{i,j}^{\hat{A}}, \quad (\text{S7.3.3})$$

where $\check{W}_{i,j}^{\hat{A}, \rho} \approx \hat{w}^{\hat{\alpha}}(\tilde{f}_{i,\hat{A}}, \rho \tilde{N}_{j,\hat{A}})$ and $\tilde{W}_{i,j}^{\hat{A}} \approx \hat{w}^{\hat{\alpha}}(\tilde{f}_{i,\hat{A}}, \tilde{N}_{j,\hat{A}})$ are again computed through a bi-linear interpolation.

After obtaining the solution to Eq. (S7.3.3), we return to the value iteration part and repeat the process until $\text{err} < \text{tol}$, where tol is a preset tolerance of convergence. In all of our numerical experiments, we have used $M_x = M_y = 1600$ on each side of the unit fN -square, $\Delta t = 0.025$, and $\text{tol} = 10^{-6}$. Our full method is summarized in Algorithm S4.

Under mild technical assumptions, Kushner and Dupuis [14, Chapters 10 & 16] showed that the discretized solution derived from a general jump-diffusion process converges to the value function using standard iterative methods. Our model forms a PDMP, which is just a specific case of jump-diffusion processes.

Algorithm S4: Value-Policy Iterations for the non-local HJB equation Eq. (S6.2.2)

```

Initialize  $W^1$  and  $A^1$  based on the boundary condition Eq. (15);
Prob_not_arrival =  $\exp(-\lambda\Delta t)$ ;
Prob_arrival =  $1 - \text{Prob\_not\_arrival}$ ;
n = 1 ;
err = 1e6 ;
while err > tol do
    for every  $(x_i, y_j) \in \Xi$  do
        for  $a \in \{0, 1\}$  do
             $\tilde{f}_{i,a} = x_i + \Delta t * F(x_i, y_j, a)$ ;
             $\tilde{N}_{j,a} = y_j + \Delta t * G(x_i, y_j, a)$ ;
             $\check{W}_{\text{temp}} \leftarrow w^n(\tilde{f}_{i,a}, \tilde{N}_{j,a})$  by interpolation;
             $\check{V}_{\text{temp}} \leftarrow w^n(\tilde{f}_{i,a}, \rho\tilde{N}_{j,a})$  by interpolation;
             $W_{i,j,a}^{n+1,\text{temp}} \leftarrow \text{Prob\_arrival} * \check{W}_{\text{temp}} + \text{Prob\_not\_arrival} * \check{V}_{\text{temp}}$ ;
        end
         $W_{i,j}^{n+1,\text{temp}} \leftarrow \max_{a \in \{0,1\}} \{W_{i,j,a}^{n+1,\text{temp}}\}$ ;
         $A_{i,j}^{n+1,\text{temp}} \leftarrow \arg \max_{a \in \{0,1\}} \{W_{i,j,a}^{n+1,\text{temp}}\}$ ;
        if  $W_{i,j}^{n+1,\text{temp}} > W_{i,j}^n$  then
             $W_{i,j}^{n+1} \leftarrow W_{i,j}^{n+1,\text{temp}}$ ;
             $A_{i,j}^{n+1} \leftarrow A_{i,j}^{n+1,\text{temp}}$ ;
        else
             $W_{i,j}^{n+1} \leftarrow W_{i,j}^n$ ;
             $A_{i,j}^{n+1} \leftarrow A_{i,j}^n$ ;
        end
    end
    err =  $\|W^{n+1} - W^n\|$ ;
    if err stagnates then
        | “Policy-Evaluation” with  $A^{n+1}$ 
    end
    n  $\leftarrow n + 1$ ;
end
 $W = W^{n+1}$ ;
 $A = A^{n+1}$ ;

```

Remark VI: This is a semi-Lagrangian-based method for solving the linear non-local probabilistic performance metric Eq. (9) in the main text, we will apply it to assess the probability performance of α_0 and α_1 .

S8 Population-dependent (hyperbolic) win and defeat boundaries

In the main text, we have focused on *fraction-dependent (vertical)* boundary arising from the definitions of killers' victory and defeat:

$$T_v(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t) > \gamma_v; f(0) = x, N(0) = y, \text{ with many dilutions} \right\}, \quad (\text{S8.1})$$

$$T_d(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t) < \gamma_d; f(0) = x, N(0) = y, \text{ with many dilutions} \right\},$$

where both stopping criteria depended on the *fraction* of the killer strain, $f(t)$, only. In the main text, we have used $\gamma_v = 0.99$ and $\gamma_d = 0.01$.

However, there are many other suitable ways of defining killers' victory/defeat. In this section, we explore *population-dependent (hyperbolic)* win and defeat boundaries and demonstrate that the results remain qualitatively similar to those obtained in the main text.

Mathematically, we now define the (random) victory/defeat time as

$$T_v(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t)N(t) > \gamma_v^h; f(0) = x, N(0) = y, \text{ with many dilutions} \right\}, \quad (\text{S8.2})$$

$$T_d(x, y, \alpha(\cdot)) = \inf \left\{ t > 0 \mid f(t)N(t) < \gamma_d^h; f(0) = x, N(0) = y, \text{ with many dilutions} \right\},$$

so that the stopping criteria depend on the *population size* of the killer strain, $n_K(t) = f(t)N(t)$.

As a result, while the equations for the "strategically-optimal" killer ($w(x, y)$) and the probabilistic performance metric ($\hat{w}^\alpha(x, y)$) remain unchanged, their respective boundary conditions are now specified on two *hyperbolas*:

$$w(x, y) = \begin{cases} 1, & \text{if } xy > \gamma_v^h, \\ 0, & \text{if } xy < \gamma_d^h. \end{cases} \quad (\text{S8.3})$$

$$\hat{w}^\alpha(x, y) = \begin{cases} 1, & \text{if } xy > \gamma_v^h, \\ 0, & \text{if } xy < \gamma_d^h. \end{cases} \quad (\text{S8.4})$$

Let w_v denote the value function with vertical boundaries and w_h the one with hyperbolic boundaries. Using the same parameter values as in Fig. S7 in the main text but with $\gamma_d^h = 0.005$ and $\gamma_v^h = 0.95$, we find that α_∞ computed with hyperbolic boundaries (Fig. S8B) is mostly the same as in Fig. S7B, except for the bottom right corner of the orange region. Consequently, the value function remains qualitatively the same too. To quantify the difference between w_v (Fig. S7A) and w_h (Fig. S8A), we calculate the mean absolute difference

$$\bar{D}(y) = \int_0^1 |w_v(x, y) - w_h(x, y)| dx. \quad (\text{S8.5})$$

and plot it for all initial populations $y \in [0, 1]$. Fig. S8C shows that \bar{D} is only large when $y < 0.05$ due to a more abrupt transition from zero to a positive winning probability in w_v compared to w_h when x is close to 1 and y is small. This is not surprising since these initial conditions are much closer to the hyperbolic defeat boundary $xy = \gamma_d^h$ than they are to the vertical defeat boundary $x = \gamma_d$. For $y > 0.05$, this x -averaged difference is very close to zero, suggesting that for most initial conditions the probability of killers' winning is largely insensitive to the type of boundary used.

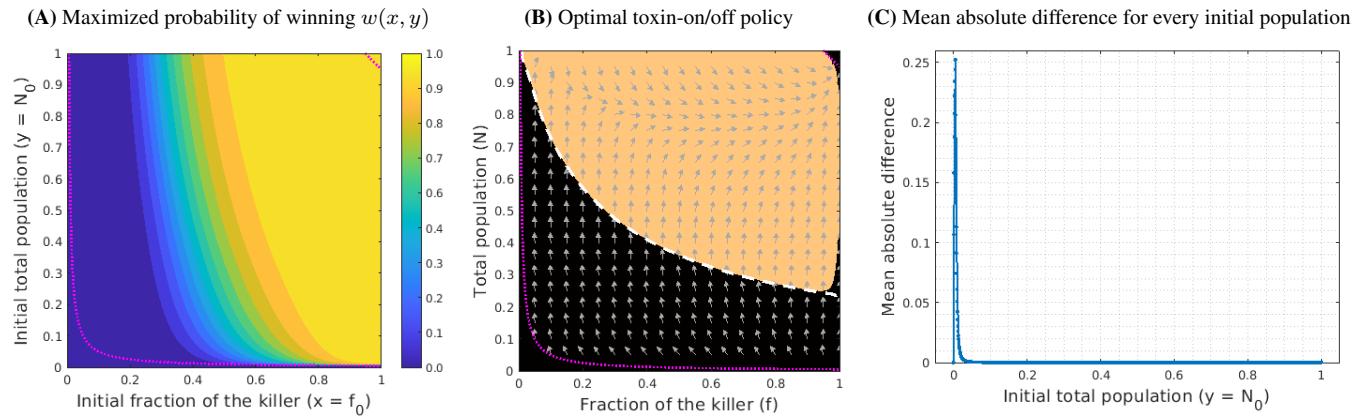


Figure S8. "Strategically-optimal" killers with hyperbolic win/defeat boundaries. The optimal toxin-on/off region (orange in panel B) is almost the same as the one computed with vertical boundaries in Fig. S7B. (The toxin-on/off switch curve from the latter is shown here as a white-dashed line). As a result, the maximized probability of winning (panel A) is also very similar to the one computed with vertical boundaries in Fig. S7A. Panel C shows the x -averaged mean absolute difference between panel A and Fig. S7A across all initial populations $N_0 = y \in [0, 1]$. This difference is only noticeable when $y < 0.05$. In (B), gray arrows denote the vector field directions corresponding to Eq. (2) in the main text with $a = \alpha_\infty(f, N)$ computed using hyperbolic win/defeat boundaries. In all panels, the victory and defeat barriers (γ_v and γ_d , respectively) are plotted with a magenta dotted line. All parameter values are the same as in Fig. S7.

This conclusion also generally holds true when using the above hyperbolic boundaries to compute \hat{w} with α_0 and α_1 . Analogously to Fig. 8 in the main text, we again focus on the initial condition $(f_0, N_0) = (0.5, 0.1)$ and compare the performance of these toxin-production policies for a range of (ρ, λ) values. Fig. S9 shows that these heat maps are both qualitatively and quantitatively similar to those in Fig. 8.

We similarly quantify the boundary-related difference in the probabilistic performance of constitutive and tactically-optimal killers in Fig. S10. Let \hat{w}_v and \hat{w}_h denote the probabilistic performance with vertical and hyperbolic victory/defeat boundaries, respectively. Focusing on the same initial condition $(f_0, N_0) = (0.5, 0.1)$, we observe that $|\hat{w}_v - \hat{w}_h|$ is negligible across half of the heat map ($\rho \geq 0.6$) for both α_∞ and α_1 . In these two cases, a noticeable difference (with a maximum of approximately 0.017) is observed when the dilutions are strong ($\rho \leq 0.55$). For α_0 , the maximum difference is slightly lower, around 0.014, while the region with noticeable differences is larger ($\rho \leq 0.6$). This is expected, as a stronger dilution more likely leads to a defeat under hyperbolic boundaries (due to a significantly larger Δ_d) compared to vertical boundaries. These results indicate that our observations and conclusions in the main text remain largely unaffected by whether the victory or defeat of the killer is defined by its fraction or population.

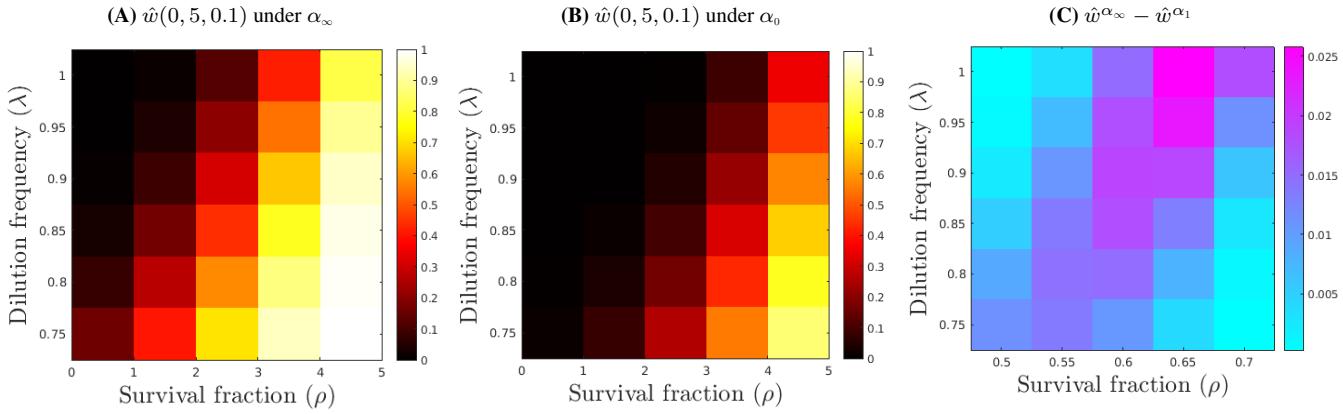


Figure S9. Hyperbolic boundaries: comparison of probabilistic performance for different types toxin-production policies starting from $(f_0, N_0) = (0.5, 0.1)$ for a range of dilution strengths and frequencies. Policy α_1 is recomputed for each λ , while policy α_∞ is recomputed for each (ρ, λ) combination. The results remain both qualitatively and quantitatively similar to Fig. 8 in the main text. A stronger survival rate (larger ρ) combined with less frequent dilutions (smaller λ) increases the chances of toxin-producers winning for all three policies. It is clear that the strategically optimal (panel A) and tactically optimal killers significantly outperform the constitutives (panel B). The differences in $\hat{w}(0.5, 0.1)$ between α_∞ and α_1 are still small, with the discrepancy increasing toward the upper right corner (panel C).

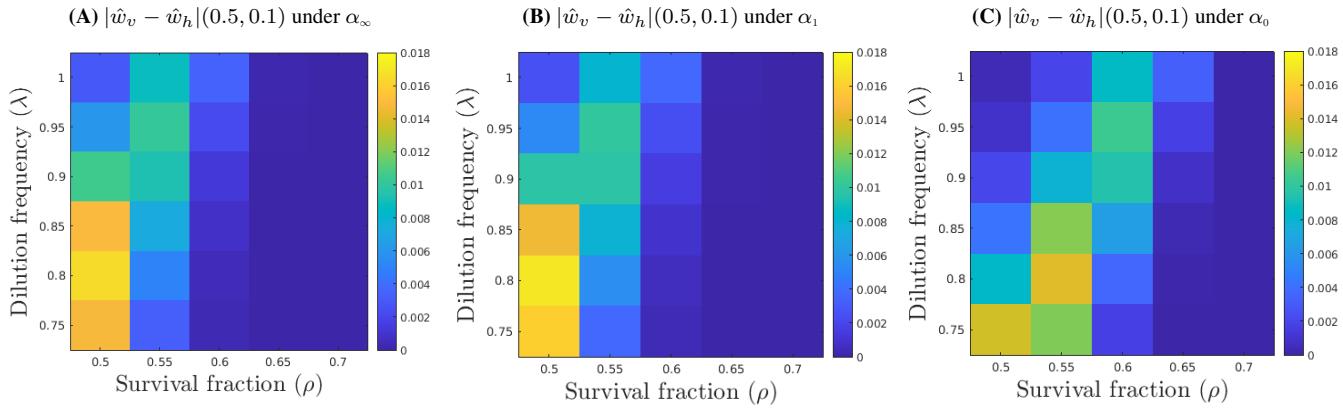


Figure S10. Hyperbolic boundaries: absolute difference in $\hat{w}(0.5, 0.1)$ resulting from two types of boundaries computed for a range of ρ and λ values. The differences under α_∞ (panel A) and α_1 (panel B) are again similar, with a maximum difference of around 0.017 when $\rho = 0.5$. For α_0 in panel C, the region of noticeable differences is slightly larger ($\rho \leq 0.6$) although the maximum difference remains relatively small (≈ 0.014). All three panels share the same colorbar.

S9 Monte Carlo simulations with “Binomial dilutions”

The results in the main text are all produced under “deterministic” dilution outcomes. That is, after each dilution, the relative abundances are preserved while only a ρ proportion of the total population survives. In this section, we present results under a specific form of *random* dilution outcomes and demonstrate, using Monte Carlo (MC) simulations, that they are qualitatively similar to the previous results.

In particular, we adopt a “Binomial Sampling” strategy to produce random dilution outcomes, which we refer to as “Binomial dilutions.” Let f^- be the pre-dilution fraction of the killers, and N^- be the pre-dilution total population. Accordingly, the actual pre-dilution number of killer cells is $n_K^- = f^- N^- C$, and the pre-dilution number of sensitive cells is $n_S^- = (1 - f^-)N^- C$. We assume each cell has an independent survival probability ρ after each dilution. As a result, the post-dilution number of cells is a *Binomial random variable*:

- Post-dilution number of killer cells: $n_K^+ = \text{Bi}(n_K^-, \rho)$;
- Post-dilution number of sensitive cells: $n_S^+ = \text{Bi}(n_S^-, \rho)$.

Consequently, the random post-dilution (normalized) total population is $N^+ = \frac{n_K^+ + n_S^+}{C}$, and the random post-dilution fraction of killers is $f^+ = \frac{n_K^+}{n_K^+ + n_S^+}$, which will serve as the initial condition for the next cycle.

We first conduct Monte Carlo simulations on a uniform grid. Starting from each $(x_i, y_j) = (i/10, j/10)$ with $i, j = 1, 2, \dots, 9$, each sample was simulated with $n = 200$ dilutions. Fig. S12 shows the empirical distributions of $f((nT)^-)$ with $T = 1$ and Fig. S11A shows their respective means. We observe that most distributions are unimodal, being either nearly 0 (all sensitives) or nearly 1 (all killers). The two exceptions with a bimodal distribution, peaking at 0 or 1, intersect precisely with the boundary (black-dashed line in Fig. S11A) that separates the initial conditions leading to a *deterministic* victory of the killers under proportional dilutions. Additionally, the killer-winning region (dark-red background in Fig. S11A) under “Binomial dilutions” aligns well with this “deterministically-killer-winning” region shown in Fig. 4C in the main text. This is not surprising since the stochastic fluctuations introduced by “Binomial dilutions” can be sufficiently large to cause samples starting near the boundary to drift towards either competitive exclusion over successive dilutions. See Fig. S11C for such an example starting from $(f_0, N_0) = (x, y) = (0.5, 0.4)$. However, these fluctuations are typically insufficient to alter the fate of samples starting further from the boundary, where initial conditions strongly favor one strain.

Considering that focusing on a single initial condition for all samples might not capture enough information, we conducted additional Monte Carlo simulations using “Binomial dilutions” with *uniformly random in a cell* initial conditions. Specifically, for each grid cell centered at (x_i, y_j) , the initial condition for each sample was chosen uniformly at random from the square $(f_0, N_0) = (x, y) \in [x_i - 0.05, x_i + 0.05] \times [y_j - 0.05, y_j + 0.05]$. This strategy diversifies the range of initial conditions, increasing the likelihood of intersecting the boundary of the “deterministically-killer-winning” region. As a result, we see from Fig. S11B that almost all cells intersecting or near the dashed-line boundary now exhibit intermediate mean values. Moreover, Fig. S13 shows that these cells again have a bimodal distribution, with random dilution outcomes pushing the dynamics towards one of the two competitive exclusions. Comparing Fig. S13 with Fig. S12, we further observe that the median at each grid cell remains unchanged. This consistency underscores the robustness of our simulation results, indicating that the conclusions in the main text would largely remain valid even under “Binomial dilutions.”

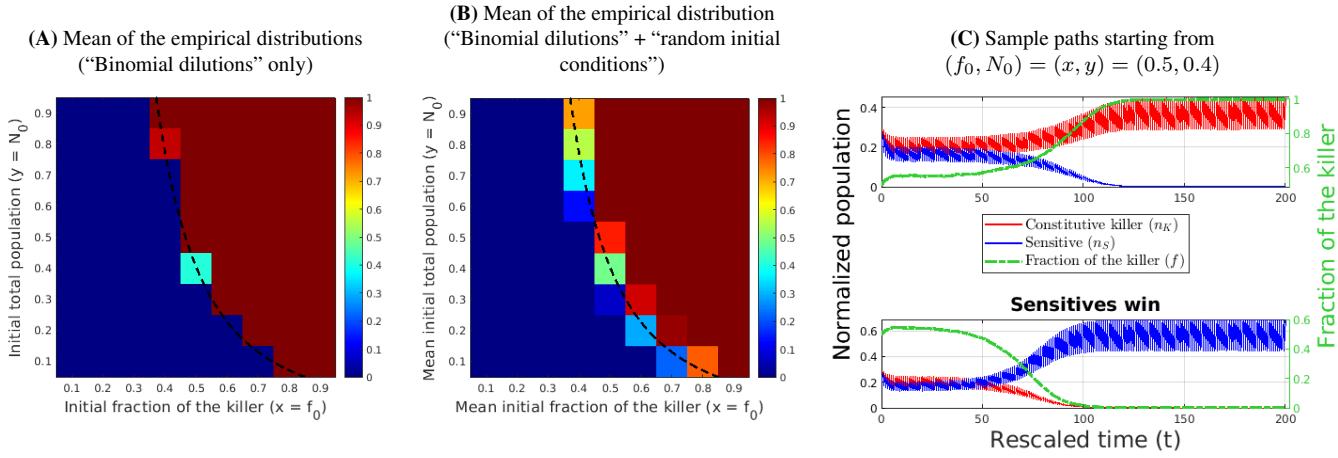


Figure S11. Monte Carlo simulations with “Binomial dilutions” on a uniform cell grid. (A) The mean of the empirical distribution, sampled with “Binomial dilutions” starting from the center of each cell in $(f_0, N_0) = (x, y)$ space. The resulting distribution is almost always unimodal (dark blue - all sensitives; dark red - all killers). The exceptions are seen in only two cells among those intersected by the boundary (shown by a black dashed line) that separates the initial conditions leading to a deterministic victory of the killers under proportional dilutions (cf. Fig. 4C in the main text). (B) Most means of the empirical distribution, sampled with both “Binomial dilutions” and “uniformly random in a cell” initial conditions, are also close to 0 or close to 1 in most cells. However, most cells that intersect or are close to that dashed line boundary now have more diverse intermediate mean values. In both cases, such cells exhibit a bimodal distribution with peaks at 0 and 1; see Figs. S12&S13 for the actual distributions. Panel (C) shows two sample trajectories starting from $(f_0, N_0) = (x, y) = (0.5, 0.4)$ resulting in different competitive exclusion outcomes due to the randomness in Binomial dilutions. All Monte Carlo simulations were conducted with 10^5 samples and 200 dilutions using parameter values $T = 1$, $\varepsilon = 0.2$, $r_{KS} = 0.85$, $\gamma = 1$, and $\rho = 0.65$. In (A), all samples for each grid cell start from the same initial condition $(x_i, y_j) = (i/10, j/10)$ with $i, j = 1, \dots, 9$. In (B), for each grid cell centered at (x_i, y_j) , the initial condition for each sample was chosen uniformly at random from the square $(f_0, N_0) = (x, y) \in [x_i - 0.05, x_i + 0.05] \times [y_j - 0.05, y_j + 0.05]$.

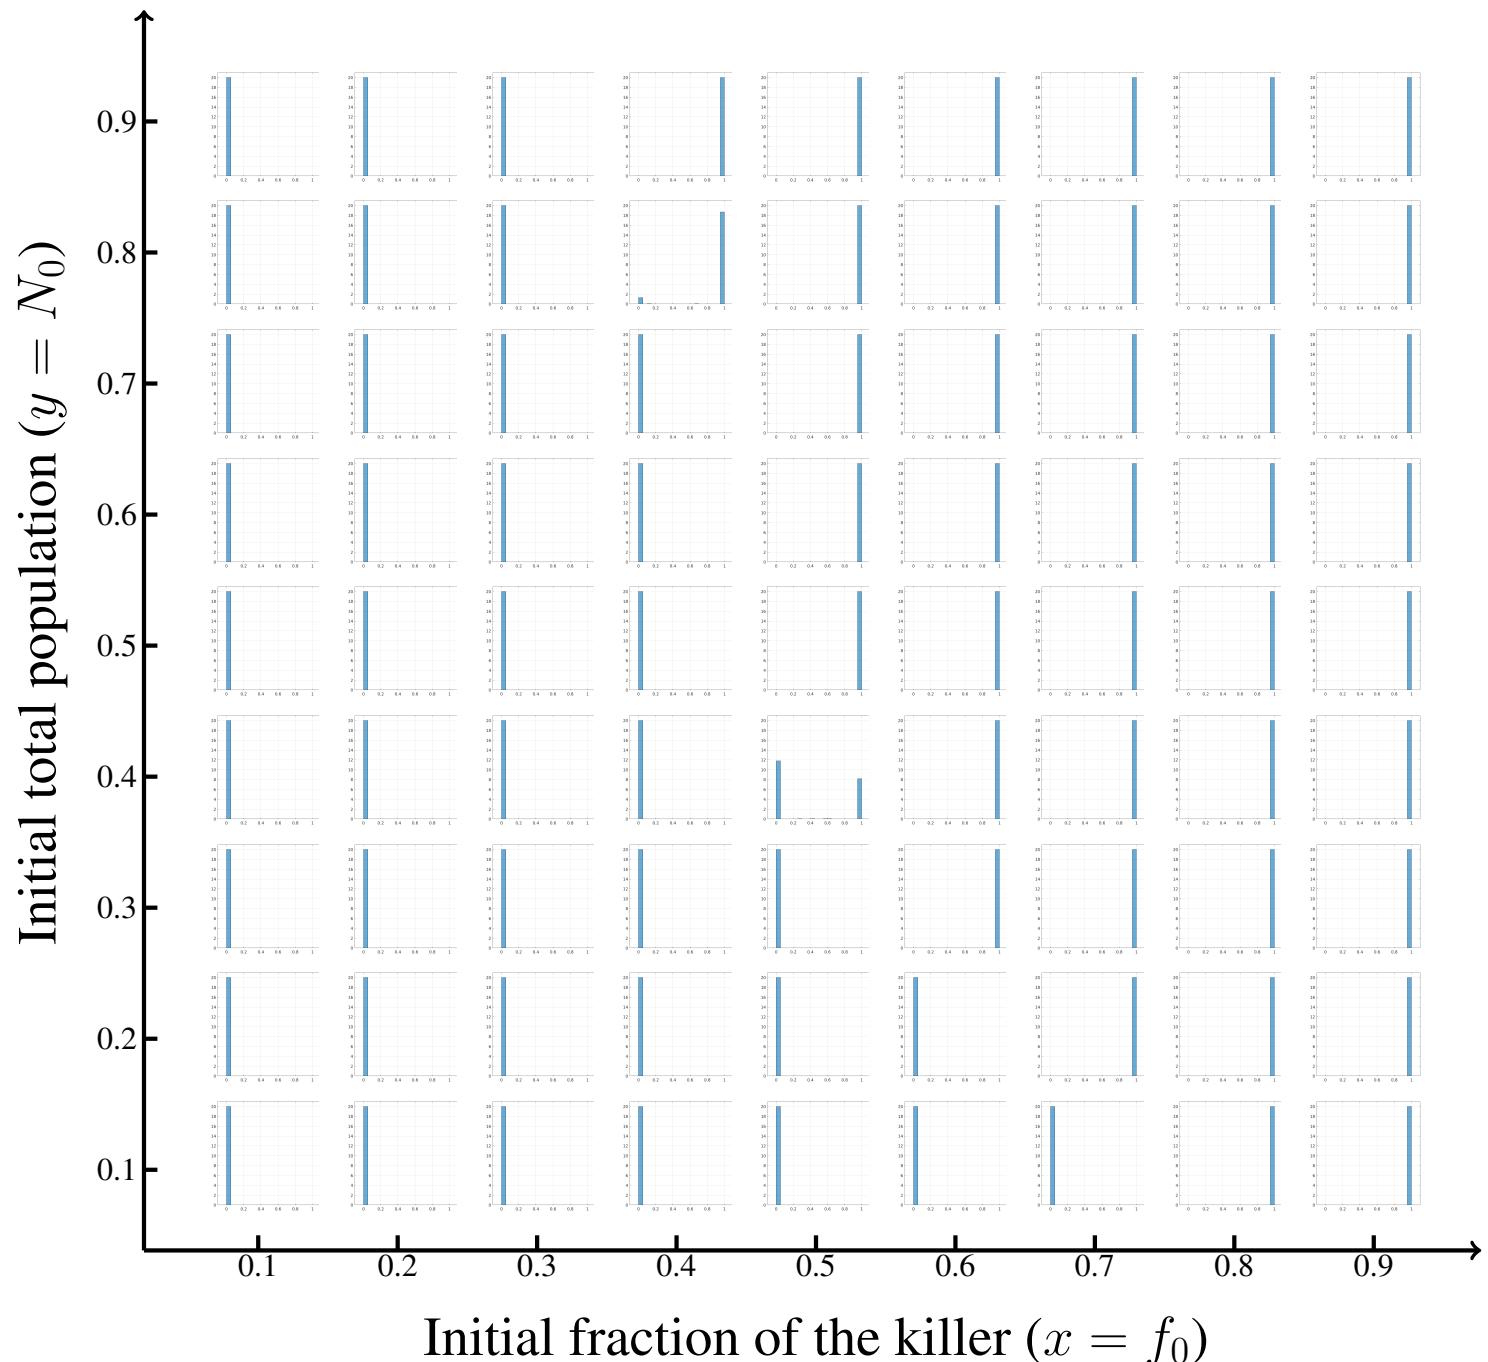


Figure S12. Empirical distributions of the fraction of killers with “Binomial dilutions” on a uniform grid. Most of the distributions are unimodal (either almost entirely 0 or almost entirely 1), except for two that are bimodal. The horizontal axis (of the entire figure) represents the initial fraction of the killer while the vertical axis encodes the initial total population for the simulations. For each panel, an empirical distribution of f (starting from the same $(x_i, y_j) = (i/10, j/10)$ with $i, j = 1, \dots, 9$) after 200 dilutions is shown by a histogram. All panels share the same horizontal and vertical axes. The parameter values are the same as in Fig. S11A.

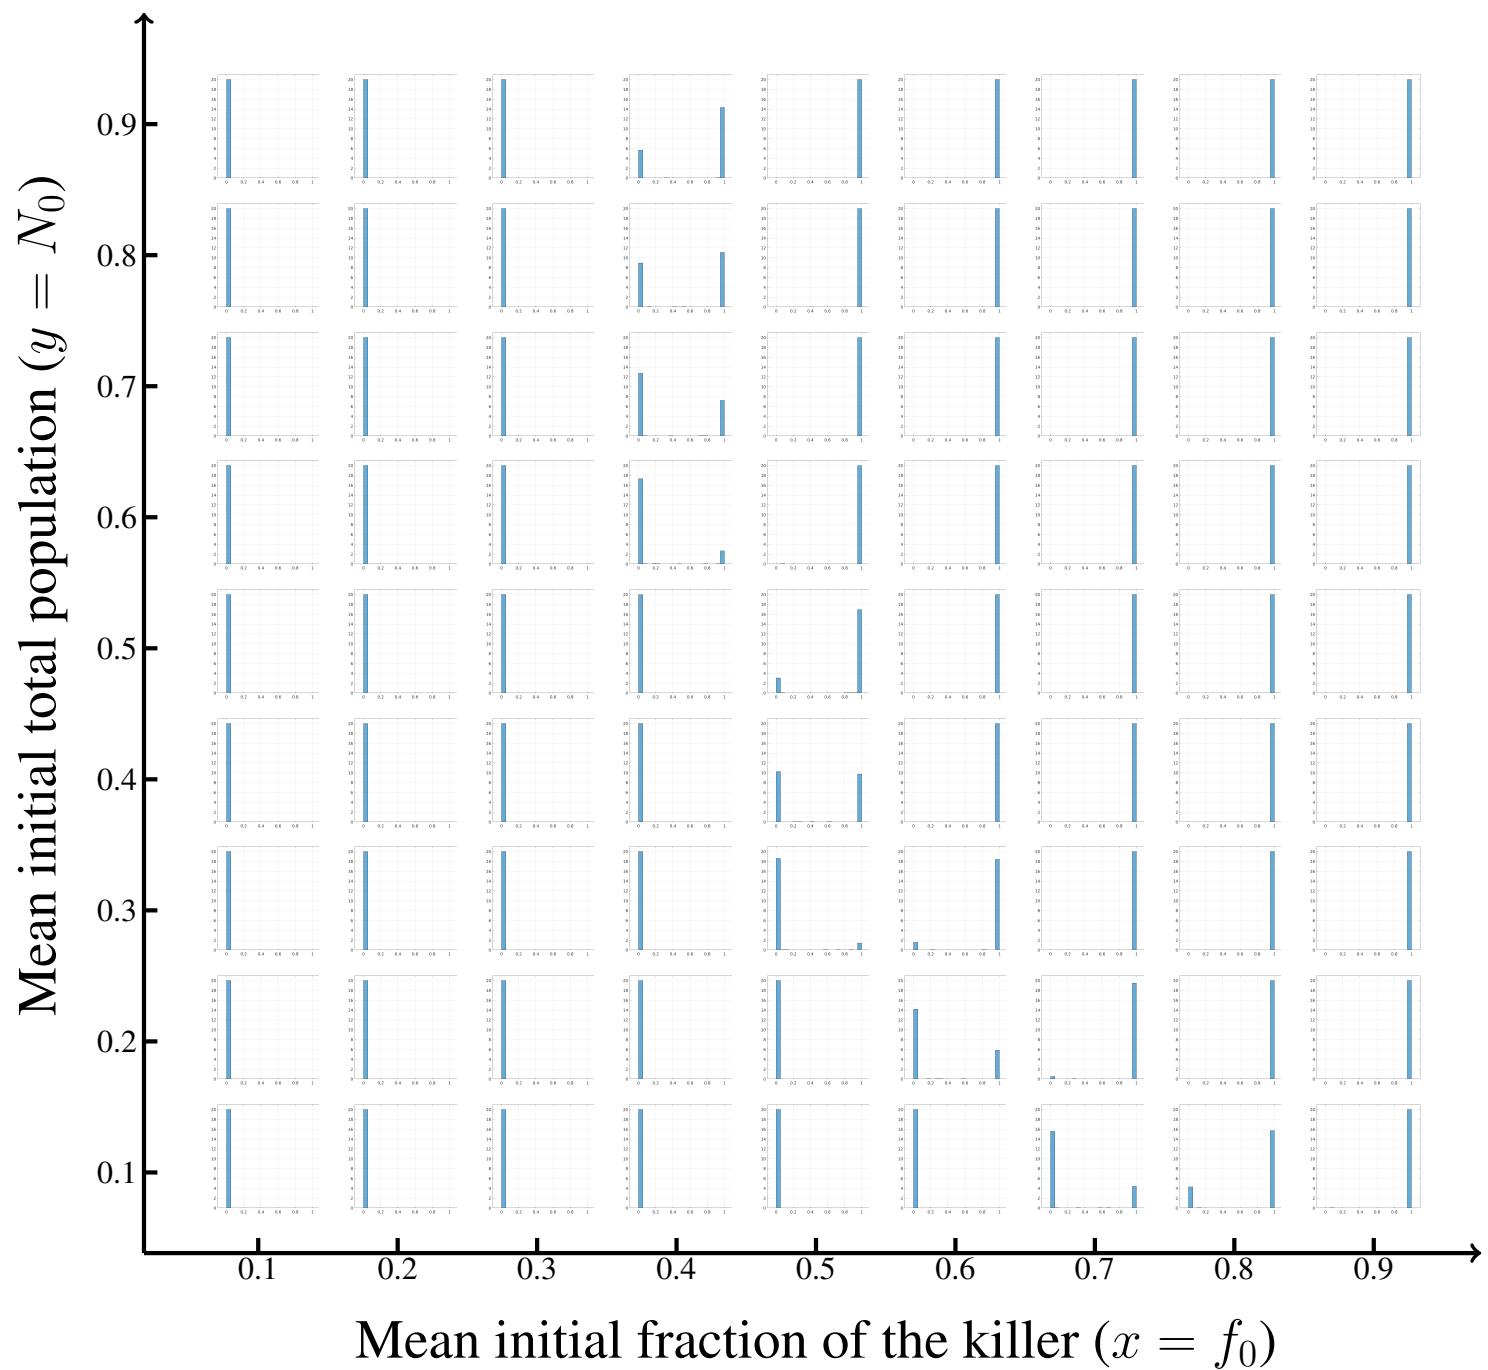


Figure S13. Empirical distributions of the fraction of killers with “Binomial dilutions” and “uniformly random in a cell” initial conditions on a cell grid. Most of the distributions are unimodal (either almost entirely 0 or almost entirely 1). However, the ones near the boundary of the “deterministically-winning” region (depicted as a black-dashed line) are *bimodal*. The horizontal axis (of the entire figure) represents the mean initial fraction of the killer while the vertical axis encodes the mean initial total population for the simulations. For each panel, an empirical distribution of $f(t)$ after 200 dilutions, with the initial condition chosen uniformly at random within the grid cell centered at $(x_i, y_j) = (i/10, j/10)$ with $i, j = 1, \dots, 9$, is shown by a histogram. All panels share the same horizontal and vertical axes. The parameter values are the same as in Fig. S11B.

The results presented above appear to be qualitatively similar for all tested parameter values. As the penalty for the ability to produce toxin and the penalty for actually producing it decreases (i.e., $r_{KS} \rightarrow 1$ and $\varepsilon \rightarrow 0$), it becomes more difficult for the sensitives to win and the dark blue region in Fig. S11A significantly shrinks; e.g., Fig. S14A. Nevertheless, when starting in the overwhelming majority, the sensitives can occasionally win under “Binomial dilutions” even if $r_{KS} = 1$; see Fig. S14A. I.e., Binomial dilutions might wipe out even super-competitive killers if they are added in small numbers to an established population of sensitives. If the Monte-Carlo simulations are started from randomly chosen initial conditions, the part of the $(f_0, N_0) = (x, y)$ domain where both strains have a chance to win increases; see Fig. S14B. This observation aligns well with the results obtained for $r_{KS} < 1$ (Fig. S11B), confirming the strong dependence of competitive exclusion on initial conditions, even when $r_{KS} = 1$.

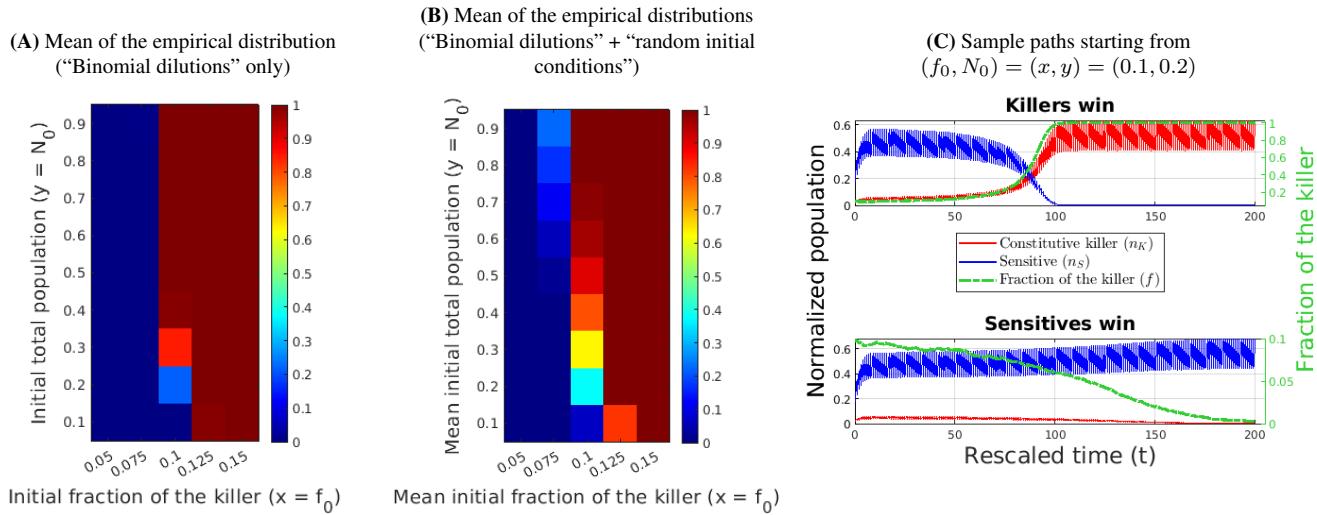


Figure S14. Monte Carlo simulations with “Binomial dilutions” on a uniform cell grid with $r_{KS} = 1$ and $\varepsilon = 0.1$. (A) The mean of the empirical distribution, sampled with “Binomial dilutions” starting from the center of each cell in $(f_0, N_0) = (x, y)$ space. The killers will always win (dark red) if $f_0 > 0.1$ while the sensitives will always win (dark blue) if $f_0 < 0.1$. Bimodal distributions, peaking at 0 or 1, occur for $f_0 = 0.1$. (B) When sampling with both “Binomial dilutions” and “uniformly random in a cell” initial conditions, more bimodal distributions with more diverse intermediate mean values appear near $f_0 = 0.1$. Panel (C) shows two sample trajectories starting from $(f_0, N_0) = (x, y) = (0.1, 0.2)$ resulting in different competitive exclusion outcomes due to the randomness in Binomial dilutions. All Monte Carlo simulations were conducted with 10^5 samples and 200 dilutions using parameter values $T = 1$, $\varepsilon = 0.1$, $r_{KS} = 1$, $\gamma = 1$, and $\rho = 0.65$. In (A), all samples for each grid cell start from the same initial condition $(x_i, y_j) = (i/40, j/10)$, with $i = 2, \dots, 6$, and $j = 1, 2, \dots, 9$. In (B), for each grid cell centered at (x_i, y_j) , the initial condition for each sample was chosen uniformly at random from the square $(f_0, N_0) = (x, y) \in [x_i - 0.0125, x_i + 0.0125] \times [y_j - 0.05, y_j + 0.05]$.

References

- [1] Andrea Giometto, David R Nelson, and Andrew W Murray. Antagonism between killer yeast strains as an experimental model for biological nucleation dynamics. *Elife*, 10:e62932, 2021.
- [2] Raquel Perruca-Foncillas, Johan Davidsson, Magnus Carlquist, and Marie F Gorwa-Grauslund. Assessment of fluorescent protein candidates for multi-color flow cytometry analysis of *Saccharomyces cerevisiae*. *Biotechnology Reports*, 34: e00735, 2022.
- [3] Rick Durrett and Simon Levin. Allelopathy in spatially distributed populations. *Journal of theoretical biology*, 185(2): 165–171, 1997.
- [4] Severin J Schink, Elena Biselli, Constantin Ammar, and Ulrich Gerland. Death rate of *E. coli* during starvation is set by maintenance cost and biomass recycling. *Cell systems*, 9(1):64–73, 2019.
- [5] June Andrews and Alexander Vladimirs. Deterministic control of randomly-terminated processes. *Interfaces and Free Boundaries*, 16(1):1–40, 2014.
- [6] Dongping Qi, Adam Dhillon, and Alexander Vladimirs. Optimality and robustness in path-planning under initial uncertainty. *arXiv preprint arXiv:2106.11405*, 2021.
- [7] Wendell H Fleming and Raymond W Rishel. *Deterministic and stochastic optimal control*, volume 1. Springer, 2012.

- [8] Mark H. A. Davis and Mohammad Farid. *Piecewise-Deterministic Processes and Viscosity Solutions*, pages 249–268. Birkhäuser Boston, Boston, MA, 1999. ISBN 978-1-4612-1784-8. doi: 10.1007/978-1-4612-1784-8_15.
- [9] Maurizio Falcone and Roberto Ferretti. *Semi-Lagrangian approximation schemes for linear and Hamilton–Jacobi equations*. SIAM, 2013.
- [10] Dimitri P Bertsekas. *Abstract dynamic programming*. Athena Scientific, 2022.
- [11] Dimitri P Bertsekas. *Dynamic Programming and Optimal Control: Volume II; Approximate Dynamic Programming*. Athena Scientific optimization and computation series. Athena Scientific, 2012. ISBN 9781886529441.
- [12] Dimitri P Bertsekas. Value and policy iterations in optimal control and adaptive dynamic programming. *IEEE transactions on neural networks and learning systems*, 28(3):500–509, 2015.
- [13] Lars Grüne and Willi Semmler. Using dynamic programming with adaptive grid scheme for optimal control problems in economics. *J. Econ. Dyn. Control*, 28(12):2427 – 2456, 2004. ISSN 0165-1889. doi: <https://doi.org/10.1016/j.jedc.2003.11.002>.
- [14] Harold J. Kushner. *Numerical methods for stochastic control problems in continuous time*. Applications of mathematics. Springer, New York, second edition, 2001. ISBN 9781461265313.