

Rethinking Population Bottlenecks: Intrinsic Fluctuations, Mutation and Dynamical Demographic Phases

Emanuele Crosato,^{1,*} Jeffrey N. Philippson,^{2,*} Shashi Thutupalli,^{2,3,†} and Richard G. Morris^{1,2,†}

¹*School of Physics and EMBL Australia node in Single Molecule Science, School of Medicine, UNSW, Sydney, Australia*

²*Simons Centre for the Study of Living Machines,*

National Centre for Biological Sciences, Tata Institute of Fundamental Research,

GKVK Campus, Bellary Road, Bengaluru 560065, India

³*International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Survey no 151, Shivakote, Hesarahatta Hobli, Bengaluru 560089, India*

By abruptly changing the size and composition of a population, bottlenecks can dramatically alter evolutionary trajectories. In the traditional picture, the faster a population recovers from a bottleneck, the more rapidly finite-size intrinsic fluctuations are suppressed and therefore the greater the likelihood of fixation to the attractor within whose basin the bottleneck constrained the population, initially. We now argue that this intuition is misleading; demonstrating that precisely the opposite behaviour is also possible. Depending critically on the rate of mutation, increased population growth can drive fixation to attractors that are *different* to that from whose basin it started. These findings are explained in terms of statistically distinct regimes of demographic behaviour, drawing parallels with the notion of non-equilibrium phase transitions. Such dynamical demographic phases are delimited by sharp transitions in time, as a population grows, and ultimately result from a time-dependent antagonism between mutation and the stochastically-induced effects of frequency-dependent birth.

I. INTRODUCTION

Population bottlenecks are widely accepted as an important facet in the modern understanding of evolution. They involve a steep decline in population size due to exogenous events, including disease, changes to the climate, or population fracture (often referred-to as a founder event) [1, 2]. Crucially, bottlenecks typically also alter a population’s composition, and have been implicated in the reduction of both genetic and phenotypic variation in organisms across a range of scales, such as viruses [3, 4], song sparrows [5, 6], tropical surgeonfish [7], elephant seals [8] and humans [9].

However, despite such putative importance, the question of *how* a bottleneck influences evolution is still largely open; there is currently a lack of both representative models and/or classifying phenomenology. This is particularly relevant since post-bottleneck recovery is often characterised by rapid and/or sustained population growth, which several pioneering theoretical works have now shown can lead to a variety of counter-intuitive behaviours. These include driving novel demographic transients [10]; facilitating cooperation [11, 12], even when deterministic selection favours defection [13]; and altering the success of invading variants [14].

An important factor common to these studies, and indeed all growing populations that are well-mixed, is that the standard deviation of intrinsic fluctuations decreases in time with the inverse square root of the growing population size, N . In the context of bottlenecks, and their

evolutionary role, this is naively understood to give rise to three monotonic effects: faster rates of recovery will more rapidly suppress fluctuations, increasing the likelihood that a given population will fixate to the attractor within whose basin it started its recovery; a larger population at the start of recovery will similarly result in an increased likelihood of such fixation, and; the nearer a population is, at the start of recovery, to an attractor, the greater the likelihood of fixation to that attractor.

We now challenge all three aspects of this picture using a growing variant of the otherwise well-studied evolutionary game—the Iterated Prisoner’s Dilemma (IPD) under replicator-mutator dynamics [15, 16]. Specifically, we show that the above phenomena are only true whilst mutation is greater than a function of the population’s growth rate. When mutation rates are below this threshold, decreasing intrinsic fluctuations leads to a variety of novel non-monotonic behaviours, including those that *reverse* the aforementioned understanding of how a bottleneck might impact evolutionary outcomes. For example, increasing the rate of growth can drive fixation to an attractor that is *different* to that from whose basin it started.

Ultimately, our findings can be traced back to an otherwise overlooked antagonistic relationship between two generic aspects of biological populations; birth rates that are frequency-dependent (*i.e.*, they depend on relative fitness and abundance) and mutation. The former leads to intrinsic fluctuations that are *state-dependent*, generically resulting in stochastically-induced ‘forces’ that drive the population towards homogeneity. The latter, by contrast, drives a population towards heterogeneity. In this sense, despite being typically seen as a source of stochasticity, mutation can act to suppress the effects of another source of stochasticity: intrinsic fluctuations due

* These authors contributed equally to the manuscript.

† To whom correspondence should be addressed: shashi@ncbs.res.in and r.g.morris@unsw.edu.au.

to replication.

Importantly, the relative magnitude of these effects changes in time, since stochastically-induced effects reduce as a population grows, whilst the effects of mutation do not. We show that this results in three sequential, statistically distinct regimes of demographic behaviour. These can be explained in terms of *effective* non-equilibrium phases of fixed- N systems, which are reminiscent of the classical study of non-equilibrium phase transitions [17, 18], previously implicated across a number of areas, including directed percolation [19], self-organisation of particle suspensions [20], surface growth [21], epidemiology [22] and even hard sphere packing [23]. These are shown to be delimited by abrupt transitions at critical values of a control parameter $\sim \log N$, implying sharp transitions in time when populations grow exponentially.

When seen through this lens, the rates of population growth and mutation together alter the character and duration of such regimes. Considering that *i*) the stochastically-induced forces demonstrated here are a general feature of frequency-dependent replication, and *ii*) demographic phases of this kind are likely to be characteristic of any well-mixed population with asymmetric attractors, we believe that our framework may prove an important paradigm for understanding the coupling between population growth, mutation and evolution, more broadly. A pertinent example is the study of quasispecies, such as RNA viruses, which are synonymous with frequent bottlenecks interspersed with periods of rapid growth.

II. RESULTS

Our results are organised as follows. In Section II A we introduce a model for understanding recovery from bottlenecks; a growing variant of the well-studied Iterated Prisoner’s Dilemma (IPD) under replicator-mutator dynamics [15, 16]. We demonstrate that there are only three possible outcomes in the long-time limit: fixation on one of the two attractors or extinction. In Section II B we compute the likelihood of such fixation as a function of bottleneck (initial conditions), and the rates of growth and mutation. This shows that growth and mutation can critically dictate the impact of a given bottleneck on evolutionary outcomes (long-term fixation probabilities), and illustrates both traditional *monotonic* fixation, and novel *non-monotonic* regimes. To understand this non-monotonicity, Section II C outlines how dynamical behaviour can be characterised by one of three effective non-equilibrium phases, dependent on the population size. These are:

1. a *stochastically-induced* phase, at small population sizes, where state-dependent fluctuations due to frequency-dependent births dominate, driving the system, temporarily, towards homogeneity;

2. an *asymmetric* phase at intermediate population sizes, where crossings between the two basins of attraction are overwhelmingly likely to be in one direction only, and;
3. a *locked-in* phase, at large population sizes, where escape from either basins of attraction is extremely unlikely.

Section II D then shows how a growing population exhibits these phases sequentially, in a manner controlled by the rates of growth and mutation. Finally, Section II E validates this conjecture by satisfactorily reconstructing the outcome statistics of Section II B using a decomposition of conditional probabilities motivated by the observed dynamical phases.

A. Growing frequency-dependent IPD

The frequency-dependent IPD is a well-established evolutionary game involving three strategies. It abstracts the key facets of biological evolution and population dynamics (e.g., birth, death, and mutation) into a tractable framework whereby the fitness of individuals is determined by their relative successes and/or failures when playing each other in the repeated game, the IPD.

In the classical formulation, each player of the repeated IPD assumes one of three strategies: “always cooperate” (AllC), who cooperate in every round; “always defect” (AllD), who defect in every round, or; “tit-for-tat” (TFT), who default to cooperation for the first round and then, at a small complexity cost, copy their opponents’ moves thereafter. In a single repeated game, players accumulate payoff over the rounds according to the standard Prisoner’s Dilemma rules: if both players cooperate, they receive a larger pay-off than if they both defect, but if one player cooperates and the other defects, then the defector receives the highest possible pay-off whilst the cooperator gets the lowest payoff. The accumulated pay-offs for an m -round repeated game are encoded by the matrix [15]:

$$\begin{array}{c} \text{AllC} \\ \text{AllD} \\ \text{TFT} \end{array} \begin{pmatrix} \text{AllC} & \text{AllD} & \text{TFT} \\ Rm & Sm & Rm \\ Tm & Pm & T + P(m-1) \\ Rm - c & S + P(m-1) - c & Rm - c \end{pmatrix}$$

where $T > R > P > S$ and $R > (T + S)/2$. We use $T = 5$, $R = 3$, $P = 1$, $S = 0.1$, and $c = 0.2$.

This game is then set against a backdrop of birth, death and mutation (Fig. 1). With a rate b , the fraction of players born into a given strategy is proportional to the fraction of the total pay-off accumulated by that strategy when all players play each other. This is given by $f_i N^i / \varphi$, where

$$f_i = \frac{\sum_{j=1}^3 a_{ij} N^j - a_{ii}}{N - 1}, \quad (1)$$

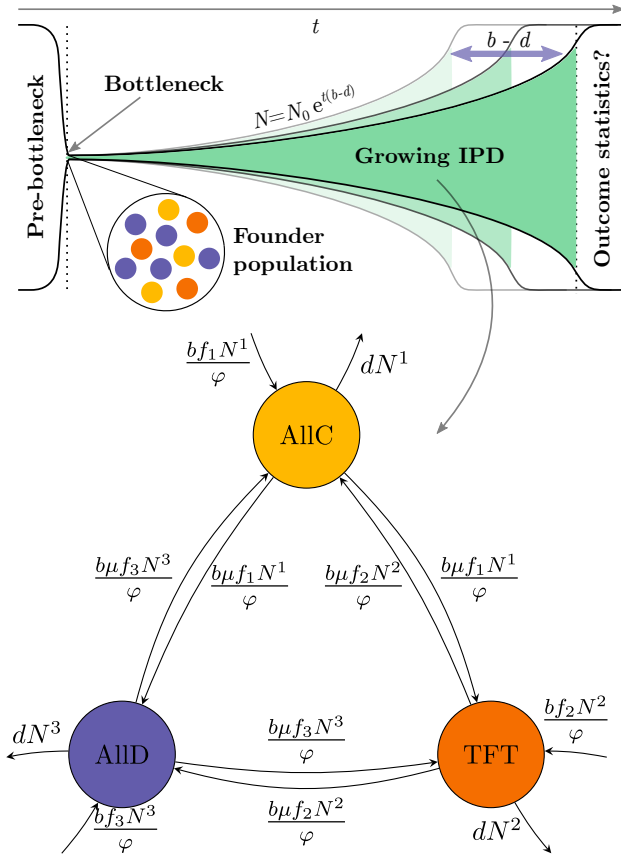


FIG. 1. The growing frequency-dependent IPD: a toy model for bottleneck recovery via population growth. By de-coupling the rates of birth and death in the frequency-dependent IPD [15], we adopt a model that combines the notions of frequency-dependence and mutation with that of population growth (see main text for definitions of rates). We ask: what is the long-term effect of a population bottleneck? Specifically, how does it influence evolutionary outcomes, and how does this change with the rate at which the recovering population grows?

is the average fitness of a given strategy when played against the whole population (including individuals of the same strategy). Here, Latin indices denote the three different strategies, i.e., AllC \rightarrow 1, AllD \rightarrow 2 and TFT \rightarrow 3 such that the a_{ij} are the components of the above payoff matrix and N^i is the number of individuals playing each strategy. In our analysis, but not in simulations, we use the simplification that

$$f_i = \frac{\sum_{j=1}^3 a_{ij} N^j}{N}, \quad (2)$$

which, although including self-interactions, still retains all the relevant features associated with the IPD [16]. In either case, the mean fitness is $\varphi = \sum_{i=1}^3 f_i N^i / N$. As a result, the higher the relative fitness associated with a strategy, the more likely that individuals are born with that strategy. A small fraction μ of births further undergo mutation and are assigned a different strategy.

Death also occurs at random, with a rate d .

Choosing a population birth rate b that is greater than the death rate d gives rise to unbounded exponential (*i.e.*, Malthusian) population growth (Fig. 2a). This rapidly suppresses $O(1/\sqrt{N})$ fluctuations and converges to well-studied deterministic behaviour (represented on the unit simplex in Fig. 2a, inset). In particular, for values of mutation rate μ in the interval $10^{-7.5} \leq \mu \leq 10^{-2.5}$, there are two stable deterministic attractors [15]: a stable ‘AllD’ fixed point, where a small fraction of TFT players (who mutually cooperate) are exploited by a large population of AllD defectors, and; a stable limit cycle around an unstable ‘mixed-strategy’ fixed point (Fig. 2b & Supplementary Information, Sec. 2). The latter is characterised by a three-phase cycle whose handedness is anti-clockwise in the traditional presentation of the state-space simplex; players of TFT can out-compete those playing AllD due to their capacity for mutual cooperation, however: they are then susceptible to invasion by players of AllC due to the complexity cost, whereby; AllC players can be exploited by those playing AllD, completing the cycle.

As a consequence, only three outcomes are possible as $t \rightarrow \infty$. Either *i*) the population goes extinct in the early stages due to finite size fluctuations (Figs. 2a-black & c) or, its demographic mix converges on *ii*) the mixed-strategy limit cycle (Figs. 2a-blue & d) or *iii*) the AllD fixed point (Figs. 2a-red & e). This justifies our choice of the growing IPD as representative model: it exhibits both a fixed point and a limit cycle—two of the most common features of any model of population dynamics—and an asymmetric simplex, which is arguably a feature of any real world setting.

B. Statistics of evolutionary outcomes

Like most models of population dynamics, the growing IPD is both non-linear and time-inhomogeneous, and thus resists most standard approaches to probabilistic analysis. Computing the statistics of the aforementioned $t \rightarrow \infty$ outcomes—*i.e.*, the fixation probabilities—therefore involves using a high performance computing facility [24] to perform stochastic simulations.

Specifically, we employ a hybrid Gillespie-Itô approach (see Supplementary Information, Sec. 1) to approximate the following fixation probability:

$$p_{\text{AllD}} = \lim_{t \rightarrow \infty} \Pr \{ \mathbf{n}_t = \text{AllD-fp} \mid \mathbf{N}_0, N_{t'} > 0 \forall t' \leq t \}, \quad (3)$$

where $\mathbf{n}_t N_t = \mathbf{N}_t = \{N_t^{\text{AllC}}, N_t^{\text{AllD}}, N_t^{\text{TFT}}\}$. This is the limiting probability, after long times, that the system converges to the AllD fixed point, given specified post-bottleneck initial conditions, and conditioned on populations that do not become extinct.

The results demonstrate several interesting features. For instance, choosing \mathbf{n}_0 to be the unstable mixed-strategy fixed point, we see that p_{AllD} depends on both the rate of population growth and the rate of muta-

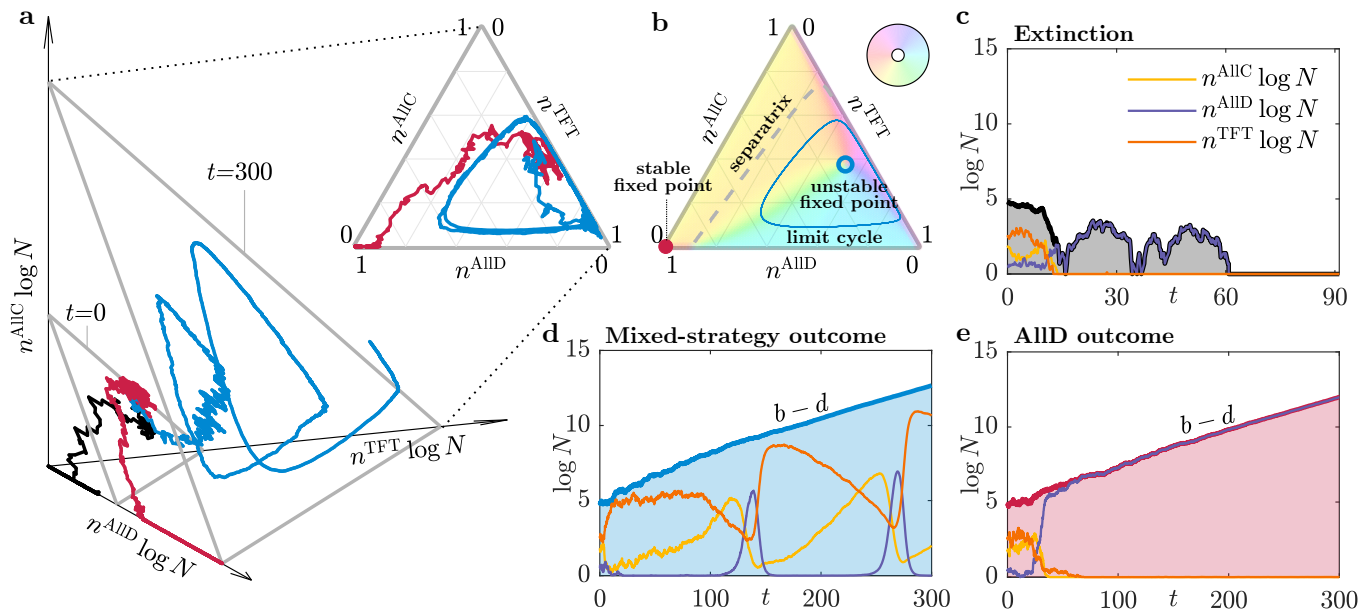


FIG. 2. **Barring extinction, recovering populations fixate on one of two deterministic attractors.** Projecting the dynamics of the growing IPD (panel a) onto the unit simplex (panel a, inset) demonstrates $O(1/\sqrt{N})$ intrinsic fluctuations that decrease as the population grows in time. Since overall population growth is exponential (with rate $b-d$) the system converges on canonical deterministic dynamics as $t \rightarrow \infty$ (panel b, colours represent the direction of deterministic flows). This results in only three different evolutionary outcomes. Either, the population goes extinct (panel a in black, panel c), or the system fixates on one of the two deterministic attractors: the mixed-strategy limit cycle (panels a in blue, panel d) or the AIID fixed point (panels a in red, panel e). The statistics of these two latter outcomes are dictated not only by the founding population, but by fluctuations. These determine the likelihood of crossing the separatrix that marks the boundary between the two basins of the stable attractors (dashed, panel b). Lower-case font represents population fractions, e.g., $n^{\text{AIID}} = N^{\text{AIID}}/N$.

tion (Fig. 3a). Similarly, fixing the rates of population growth and mutation reveals sensitivity to initial conditions, where p_{AIID} depends on both initial demographic mix, \mathbf{n}_0 , and initial population size, N_0 (Figs. 3b-g).

The parameter space of mutation and growth rates can moreover be divided into two qualitative regions (panel a, red dashed line).

In the upper region, behaviour agrees with the monotonic expectation that faster growth rates reduce $O(1/\sqrt{N})$ intrinsic fluctuations more rapidly and therefore increase the likelihood of fixation within the same basin of attraction that the system started (panel a). Similarly, as the initial population size, N_0 , increases, the system experiences fewer large fluctuations at early times and hence this also increases the likelihood of fixation within the starting basin of attraction (Figs. 3b-d).

In the lower region, however, behaviour is more complex. There are high likelihoods of an AIID outcome at both low ($b-d \lesssim 0.01$) and intermediate ($0.03 \lesssim b-d \lesssim 0.07$) rates of growth (panel a). This results in a non-monotonic dependence of p_{AIID} as a function of population growth. That is, despite more rapidly reducing the fluctuations that are ostensibly required to cross the separatrix between mixed-strategy and AIID basins, growth can actually *increase* the likelihood of fixating on the AIID fixed point. Similarly, rather than decreasing the likelihood of crossing the separatrix and

fixating there, increasing N_0 can actually increase this likelihood for certain initial states, confounding expectations regarding the role of fluctuations at early times when populations remain small (Figs. 3e-g).

In the context of population bottlenecks, this demonstrates that the long-term ramifications of reducing a population to a particular size and demographic mix can depend, critically, on the rates of post-bottleneck growth and mutation.

C. Effective non-equilibrium phases

To understand the non-trivial behaviour in Fig. 3, we repeatedly simulate the *fixed* population size IPD (Supplementary Information, Secs. 1 & 7), computing a so-called empirical distribution [25]. Specifically, we calculate the mean fraction of time spent in the AIID basin,

$$F = \left\langle \frac{1}{T} \int_0^T 1_{\mathbf{n}_t \in \text{AIID-b}} dt \right\rangle_{\mathcal{N}}, \quad (4)$$

where 1 is the indicator function, \mathcal{N} denotes the size of the ensemble over which the average is taken, and the integration is understood in the Itô sense. The time $T > (\log N_{\text{max}})/(b-d)$ exceeds the entire duration of our growing simulations (which we stop at $N_{\text{max}} = 10^{10}$)

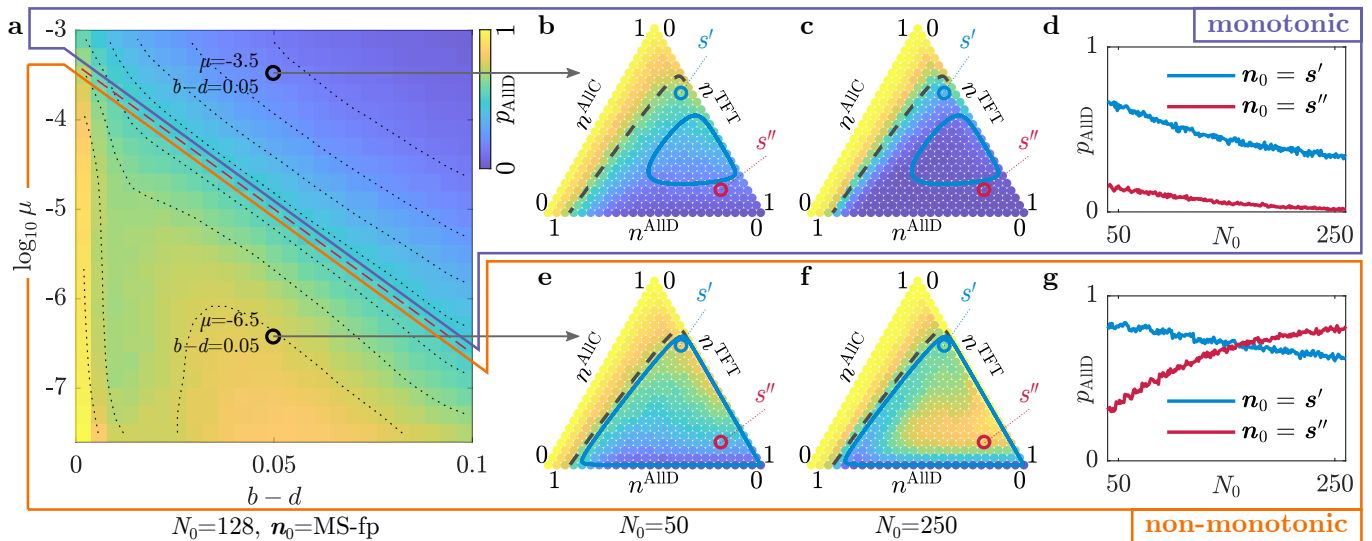


FIG. 3. Rates of growth and mutation critically alter how a given bottleneck translates into fixation probabilities. For a range of initial conditions, growth and mutation rates, we use repeated Gillespie-Itô simulations to compute p_{AIIID} , the probability that a growing population will fixate on the AIIID fixed point. Choosing the initial demographic mix as the mixed-strategy fixed point, we see that p_{AIIID} depends on both the rate of population growth and the rate of mutation (panel a). Conversely, fixing the rates of growth and mutation, p_{AIIID} also depends on both the initial demographic mix \mathbf{n}_0 (panels b, c, e & f) and the initial population size, N_0 (panel d & g). Behaviour can be qualitatively classified into two regions separated by the red dashed line (panel a). In the upper region, behaviour agrees with expectations: higher rates of growth more rapidly suppress $O(1/\sqrt{N})$ intrinsic fluctuations and therefore increase the likelihood of fixation within the basin of attraction that the system started (panel a). Increasing the initial population size reduces the likelihood of large fluctuations at early times, therefore exacerbating this effect (panels b-d). In the lower region, behaviour is more complex: there are high likelihoods of an AIIID outcome at both low ($b-d \lesssim 0.01$) and intermediate ($0.03 \lesssim b-d \lesssim 0.07$) rates of growth (panel a). This results in a non-monotonic dependence of p_{AIIID} on $b-d$, which implies that population growth can actually *increase* the likelihood of crossing the separatrix between the two basins of attraction, despite more rapidly suppressing $O(1/\sqrt{N})$ fluctuations. There is a similarly non-trivial structure to the effects of initial demographic mix on p_{AIIID} where, for certain \mathbf{n}_0 , an increase in N_0 actually causes p_{AIIID} to increase rather than decrease (representative initial states \mathbf{s}'' and \mathbf{s}' , respectively, panels e-g).

and represents an effective cutoff, so that the statistics of fixed-size simulations are not skewed by events that are highly unlikely to occur in the growing simulations (*i.e.*, with characteristic rates $\ll 1/T$). Such a large but finite T therefore aims to capture the average transient behaviour of a growing population at a particular N .

Our results (Fig. 4a) are suggestive of a large deviation principle, such that $p(F|\mathcal{N}) \asymp \exp[-\mathcal{N} I_\mu(F_N)]$, where I_μ is a convex rate function. Although determining the precise functional form of I_μ is considered out-of-scope for this article, our data suggests that it has only three zeros, despite varying N over ten orders of magnitude:

$$\arg \min I_\mu(F) \approx \begin{cases} 0.5, & \forall N \leq 10^{2.7} \\ 0, & \forall 10^{2.7} < N < 1/\mu \\ 1, & \forall N > 1/\mu \end{cases} \quad (5)$$

This means that, depending on the population size, there are three statistically distinct types of characteristic demographic behaviour. Due to the finite size of T , we call these *effective* non-equilibrium phases. The three effective phases are characterised as follows:

1. *Stochastically-induced phase*— Demographic trajectories at small population sizes are characterised

by large intrinsic fluctuations and an intermediate value of F (Fig. 4a in magenta, Figs. 4b & c). Fluctuations are both correlated and state-dependent; features that are captured by the symmetric 3×3 correlation matrix, B_{ij}^\dagger , that can be obtained by performing a Van Kampen system-size expansion [26] and projecting the results onto the unit simplex using Itô's lemma (Fig. 4b & c, orange crosses, and Supplementary Information, Sec. 5). The frequency-dependent nature of births means that fluctuations at the centre of the simplex are large and isotropic, whilst the components normal to the boundaries decrease rapidly as the edges and corners are approached. This gives rise to stochastically-induced effects [27, 28], where fluctuation gradients bias stochastic trajectories, driving them towards the simplex edges and corners, on average (Fig. 4b & c and Supplementary Information Secs. 3-5). Despite such overall behaviour, B_{ij}^\dagger is not symmetric under the interchange of n^{AIIID} , n^{AIIIC} and n^{TFT} , and stochastic trajectories retain characteristics encoded by the payoff matrix; including a bias for anti-clockwise motion, and a compara-

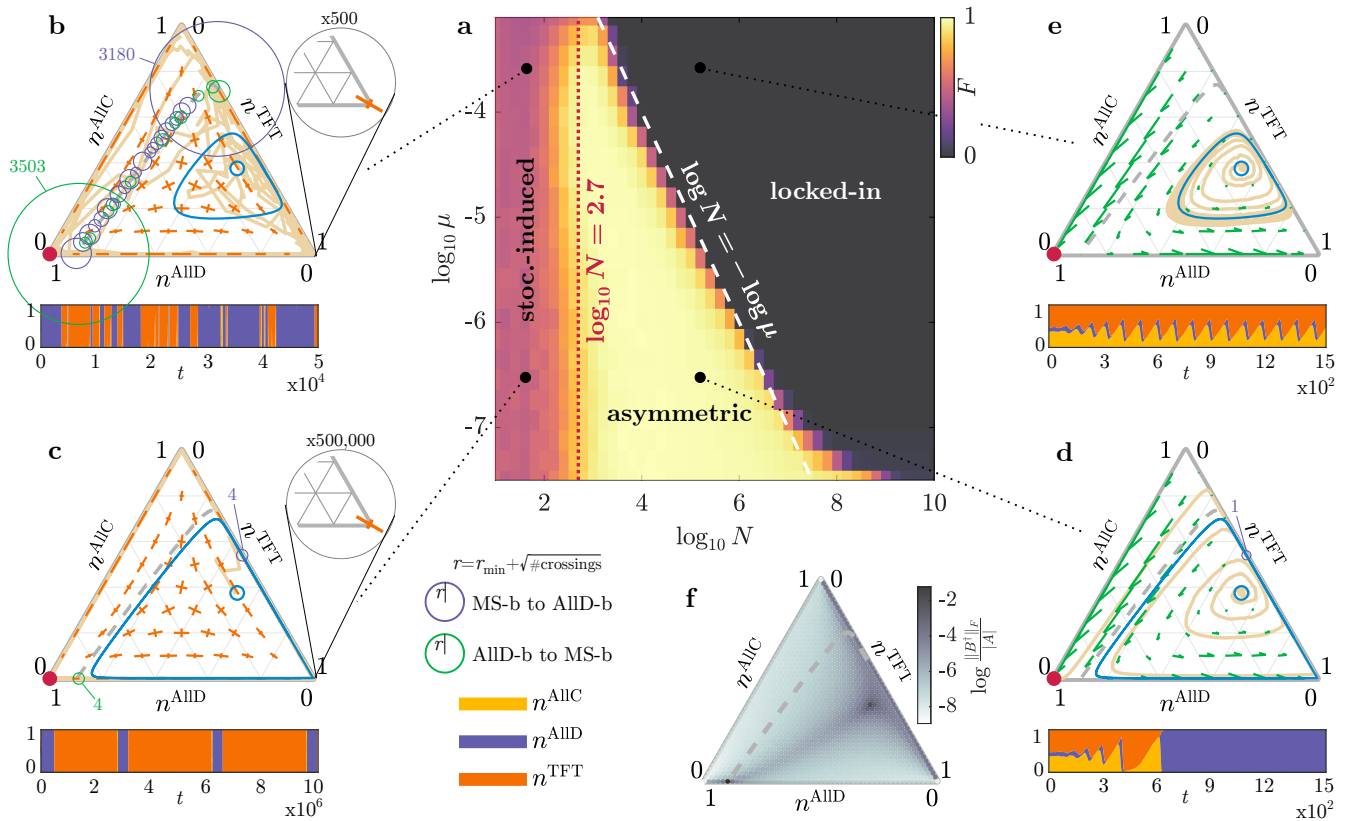


FIG. 4. Effective non-equilibrium phases. When characterised by the fraction F (of a large but finite time) that trajectories spend in the basin of the AllD fixed point, repeated simulations of the IPD at fixed population sizes demonstrate three distinct regimes of demographic behaviour (panel a). When population sizes are small, fluctuations are dominant, and the separatrix can be crossed in both directions (circles, panels b & c). Eigenvalues of the projected correlation matrix, B_{ij}^\dagger indicate that the magnitude and bias of the fluctuations depend on demographic mix (orange crosses, panels b & c). Generally, this gives rise to fluctuation gradients that drive the system towards the simplex edges and corners. Although, due to the structure of the payoff matrix there is also an anti-clockwise bias and a lower probability of finding the AllC corner. At the simplex boundaries, normal fluctuations are proportional to $\sqrt{\mu}$ (panels b & c, magnified inset). Characteristic trajectories (beige) are therefore increasingly confined to the boundaries and corners as μ decreases. Since the residence times associated with the corners are $O(1/\mu)$, trajectories spend a disproportionate fraction of their time in the AllD and TFT corners (bars, panels b & c). Beyond this regime, behaviour becomes increasingly deterministic (green half-arrows, panels e-d). For intermediate population sizes, behaviour is *asymmetric*: demographic trajectories can cross the separatrix from the mixed-strategy limit cycle, but not from the AllD fixed point (panel d). The upper critical population size of this regime scales as $\sim 1/\mu$ (white, panel a). For large enough N , behaviour becomes increasingly deterministic, and trajectories remain locked-in to the mixed-strategy limit cycle (beige, panel e). At low μ , the separatrix is typically crossed where it intersects either the TFT or AllD edges, which is where the magnitude of stochastic effects, $\|B^\dagger\|_F$, are largest relative to deterministic flows, $|A_i|$, with $\|\cdot\|_F$ and $|\cdot|$ denoting Frobenius- and ℓ^2 -norms, respectively (panels c, d & f). Representative trajectories and population fractions (bars) only show a fraction of the total time simulated (see bar legend) in order to aid visualisation. Crossing statistics (circles) are taken from single simulations lasting 10^6 s.

tively low likelihood of reaching the AllC corner (when compared to AllD and TFT corners). These behaviours are crucial to understanding the precise μ -dependent mechanisms that underpin the value $F \approx 0.5$.

2. *Asymmetric phase*— Increasing N , the magnitude of fluctuations decreases, and the relative geometry of the underlying attractors becomes increasingly important. In particular, the system enters an asymmetric regime at populations above

$N \approx 10^{2.7}$, for which $F \approx 1$ (Fig. 4a, yellow). Here, state-dependent fluctuations permit the system to cross from the mixed-strategy limit cycle to the AllD basin, but not from the AllD fixed point to the basin of the limit cycle (Fig. 4d). In other words, once the separatrix has been crossed, trajectories are extremely unlikely to come back within the time T .

3. *Locked-in phase*— Once N is sufficiently large, fluctuations are small and demographic trajectories are

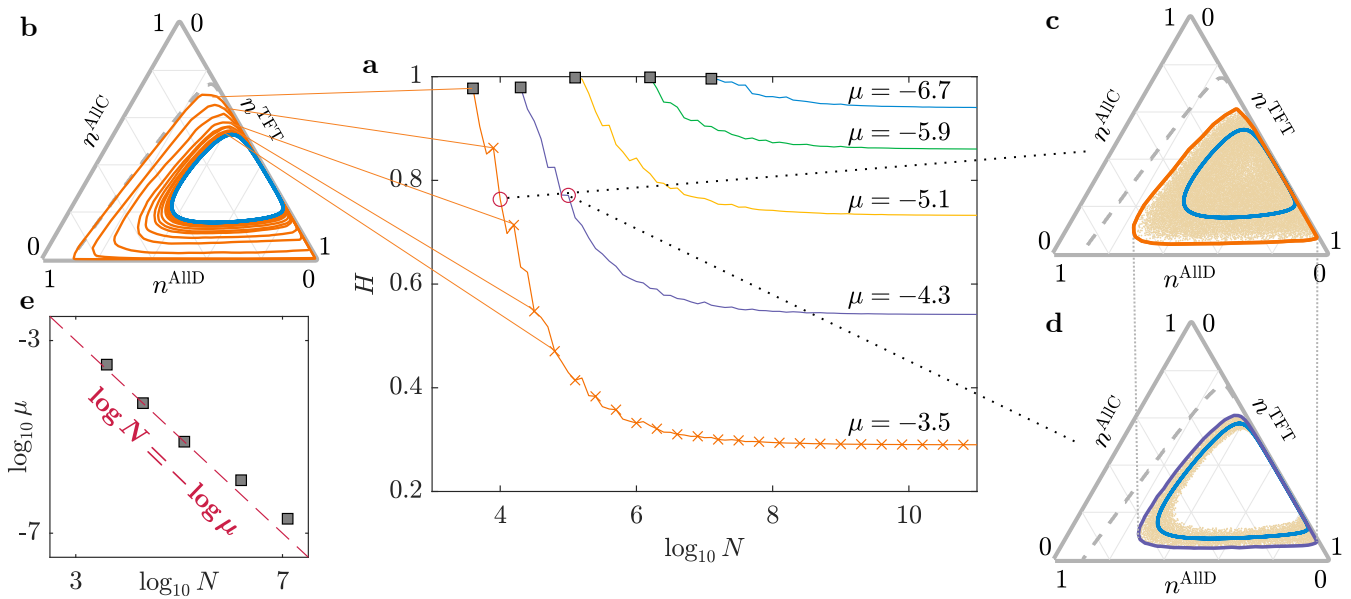


FIG. 5. Understanding the mutation-dependent transition between asymmetric and locked-in phases. At finite population sizes (and for times $t < T$) demographic trajectories in the locked-in regime repeatedly circumnavigate the limit cycle, leading to a well-defined ‘footprint’ of visited points due to the state-dependence of the noise correlations. We compare such footprints via the convex hull, $\text{Hull}(\mathbf{n}_{t < T})$, of all points sampled by a trajectory $\mathbf{n}_{t < T}$ within a specified (large) time, T (panel a). The area of this convex hull is written as a fraction of the area A of the mixed-strategy basin of attraction, $H = \text{Area}[\text{Hull}(\mathbf{n}_{t < T})]/A$. For a given μ , the fraction H converges exponentially to the value associated with the deterministic limit cycle, as N increases (panels a & b). As a result, there are approximate equivalences between the footprint of trajectories with a large μ and small N and those with a correspondingly smaller μ and larger N (panels a, c & d). This leads to a family of scaling relations in the locked-in regime (panel a). The *critical* scaling, at which the footprint occupies the entire mixed-strategy basin, is $N \sim 1/\mu$, which agrees with ensemble statistics (panel e). Since this corresponds to onset of the locked-in phase, the corollary is that the asymmetric phase has a duration that is μ -dependent (*cf.* Fig. 4).

effectively locked into the mixed-strategy basin for times $< T$, implying $F \approx 0$ (Fig. 4e, dark gray).

The aforementioned effective phases moreover couple to mutation, which alters both the character of the stochastically-induced phase (Figs. 4b & c), and the population size at which the system transitions from asymmetric to locked-in phases (Fig. 4a, white dashed line).

The origin of the former is that the standard deviation of fluctuations normal to the boundaries is $O(\sqrt{\mu}/N)$ (Supplementary Information, Sec. 6). Evolutionary trajectories therefore become increasingly confined to the boundaries as μ decreases (Figs. 4b & c). This not only exacerbates stochastically-induced effects, but also increases the mean residence times associated with the corners. In particular, at comparatively high levels of mutation, residence times are less than T , which results in a stochastic cycling between TFT and AIID corners (recall that there is a lower likelihood of finding the AIC corner) (Fig. 4b). Since these two corners have comparable mean rates of escape to the opposite basin of attraction (Supplementary Information, Sec. 8), F takes a value of approximately 0.5. By contrast, lower rates of mutation imply dwell times greater than T (Fig. 4c). On average, therefore, the system will find either the AIID or TFT corner and then remain there. Here, the value

of ≈ 0.5 results from the position of the mixed strategy fixed point, and the correspondingly equal probability that trajectories are expelled to either the AIID or TFT corners (typically, via the AIC-TFT edge, see Fig. 4f).

For the latter, the $\sim 1/\mu$ dependence of the asymmetric to locked-in transition can be understood in terms of the stochastic ‘footprint’ of evolutionary trajectories in the locked-in regime—*i.e.*, those that repeatedly (and stochastically) navigate the limit cycle (Fig. 5). In particular, due to the finite nature of T , the convex hull of this footprint is well-defined, reflecting the shape of the limit cycle at a *different* value of μ (Fig. 5a-d). The result is a family of scaling relations, where the footprint of small populations with high levels of mutation (*i.e.*, high noise, small limit cycle, Fig. 5c) is approximately equivalent to that of large populations with low levels of mutation (*i.e.*, low noise, large limit cycle, Fig. 5d). The *critical* scaling that defines the onset of the regime occurs when the stochastic footprint fills the mixed-strategy basin (Figs. 5a & e), therefore facilitating the crossing of the separatrix. Here, $O(\sqrt{\mu}/N)$ fluctuations must be equivalent to the $O(\mu)$ deterministic repulsion in the direction normal to the AIC-TFT edge (Supplementary Information, Sec. 6), implying $N \sim 1/\mu$, which agrees with both ensemble statistics and convex hull analysis

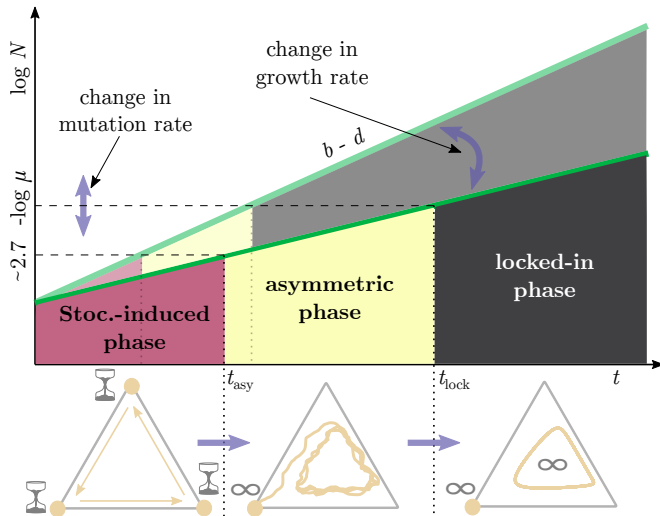


FIG. 6. Growth and mutation control the relative duration of dynamical regimes. The three effective non-equilibrium phases in fixed-size systems are associated with a control parameter, $\log N$ (Fig. 4). In an exponentially growing system, this is directly proportional to time. As such, a growing population exhibits three distinct dynamical regimes in sequence. Changing the overall rate of growth therefore re-scales the duration of the stochastically-induced and asymmetric phases. Changing the rate of mutation, by contrast, alters only the asymmetric phase. The duration of the asymmetric phase is important to overall fixation probabilities, since a longer duration implies a higher likelihood of crossing the separatrix, from which there is no return.

(cf. Figs. 4a & 5e).

This behaviour can be recast as a size-dependent antagonistic relationship between mutation and intrinsic noise, where the former favours population heterogeneity (attracting towards the centre of the simplex) and the latter homogeneity (expelling towards the simplex boundaries). This is particularly important in growing populations, since the balance between the two factors changes over time.

D. Dynamical regimes in growing populations

Since the control parameter of the effective non-equilibrium phases is $\log N$, an exponentially growing population exhibits the three phases in sequence, with transitions occurring abruptly in time (Fig. 6).

Seen through this lens, it is clear that changing the rate of growth changes the duration of the stochastically-induced and asymmetric phases by an overall factor (the locked-in phase is, in principle, open-ended). By contrast, the rate of mutation changes only the duration of the asymmetric phase (by virtue of the μ -dependent transition to the locked-in phase).

The asymmetric phase is especially important for determining fixation probabilities. If the separatrix is

crossed from the mixed-strategy basin during this phase, then the system remains in the basin of the AllD fixed point until the onset of the locked-in phase, where it remains for all until the population reaches $N_{\max} = 10^{10}$. Therefore, the longer the duration of the asymmetric phase, the greater the likelihood of crossing the separatrix.

However, despite the appeal of this heuristic, we must also account for the fact that this likelihood is also conditioned on the state at which the system enters the asymmetric regime. This is determined by the duration of preceding stochastically-induced regime, which is set by the rate of growth. It is also determined by the character of that regime, which is set by mutation. Despite the presence of two μ -dependent mechanisms in the fixed- N ensemble (see previous Section) it should be stressed that, in a *growing* system the duration of the stochastically-induced regime is sufficiently short in comparison to corner dwell times that, for all but pathologically slow growth rates or very large μ , the separatrix is only ever likely to be crossed once. In turn, this puts greater emphasis on the role of the system's initial conditions.

To test our understanding of this complex interplay, we construct an approximation to the full outcome statistics. This is based on a decomposition in terms of conditional probabilities associated with each phase, and four 'equivalence classes' of states at the start of the asymmetric phase.

E. Equivalence classes

Consider the conditional probability $p_{\text{AllD}}^{\text{asy}}(\mathbf{s}) = \Pr \{ \mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s} \}$, where avoidance of extinction is now assumed implicitly. That is, the likelihood of being in the AllD basin at the onset of the locked-in phase, t_{lock} , given that the system was in a state \mathbf{s} at the onset of the asymmetric phase, t_{asy} . Computing this probability via stochastic simulation demonstrates the existence of three approximate equivalence classes, \mathcal{S}_i , such that $p_{\text{AllD}}^{\text{asy}}(\mathbf{s}_i)$ is agnostic to the demographic mix $\mathbf{s}_i \in \mathcal{S}_i$, at the onset of the asymmetric phase (Fig. 7a and Supplementary Information, Fig. 11). These are: those states in the AllD basin, \mathcal{S}_1 ; those along the AllD-TFT edge that stretch from the separatrix to the TFT corner, \mathcal{S}_2 , and; those along the TFT-AllC edge (excluding the TFT corner) on the mixed-strategy side of the separatrix, \mathcal{S}_3 . The remaining states of the mixed-strategy basin are labelled \mathcal{S}_4 . Whilst these do not form an equivalence class, we assume (and later show) that they only minimally contribute to overall fixation probabilities.

The existence of equivalence classes prompts the following simplification (Supplementary Information, Sec. 9):

$$p_{\text{AllD}} \approx \sum_{i=1}^4 p_{\text{AllD}}^{\text{asy}}(\mathbf{s}_i) p_{\mathcal{S}_i}^{\text{stoc}}, \quad (6)$$

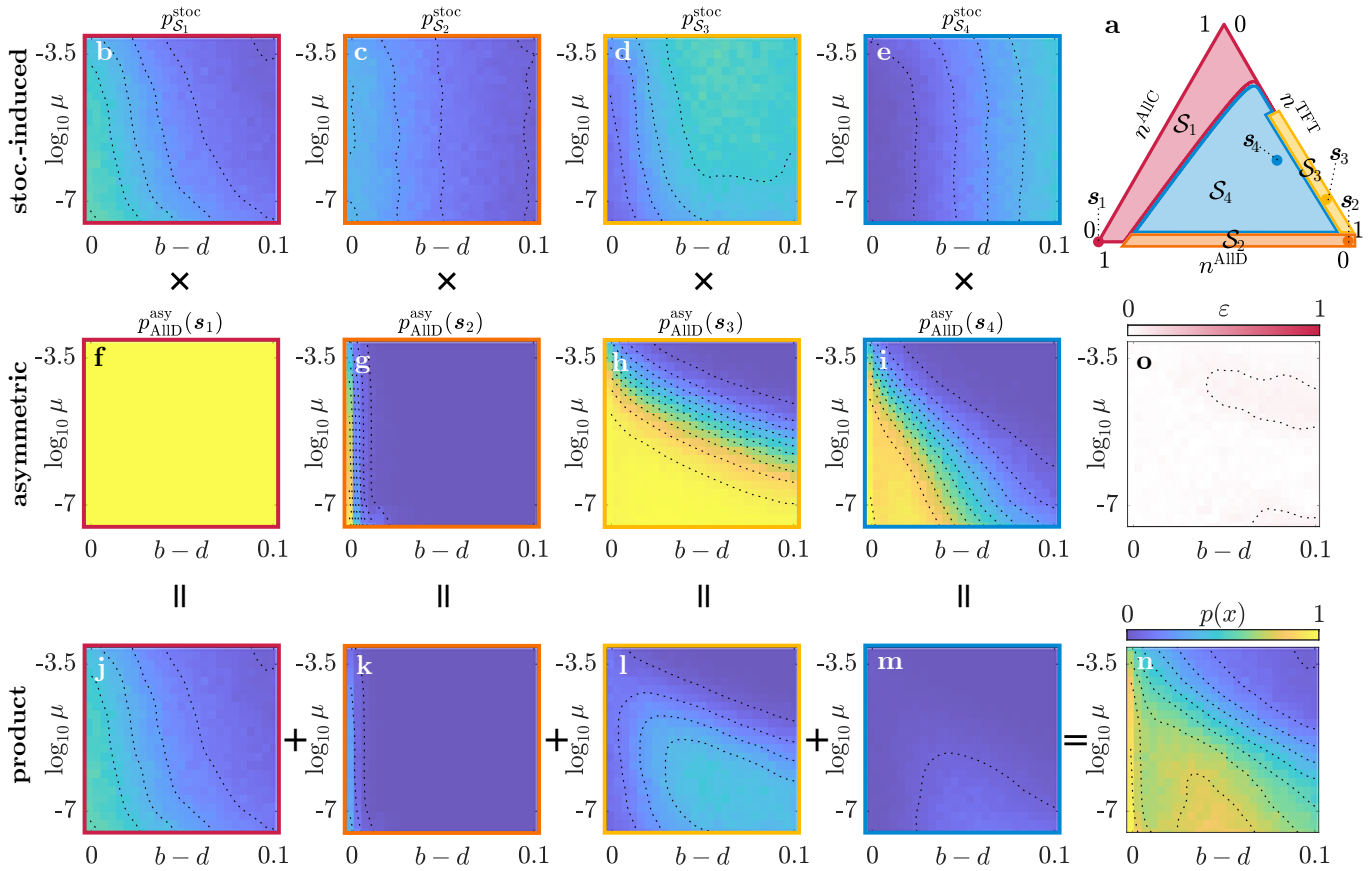


FIG. 7. The dependence of fixation probabilities on rates of growth and mutation is reproduced by a decomposition based on effective phases and equivalence classes. Fixation probabilities can be decomposed in terms of conditional probabilities that are based on equivalence classes of the asymmetric phase (panel **a**). Within each class \mathcal{S}_i , $i = 1, \dots, 3$, the outcome of the asymmetric phase does not rely on the specific demographic mix at the onset of the phase. Moreover, the remaining states \mathcal{S}_4 contribute very little to the overall fixation probabilities (panels **e**, **i** & **m**, and main text). This permits the approximation in Eq. (6), graphically represented by panels **b-n**. The resulting reconstruction is in good agreement with the full stochastic simulations (panels **n** & **o** and Fig. 3a). The same decomposition is shown in Supplementary Information, Sec. 11, for additional initial conditions. The initial population size is $N_0 = 128$, whilst the times t_{asy} and t_{lock} are derived from the critical population sizes identified in Fig. 4.

where $p_{\mathcal{S}_i}^{\text{stoc}} = \Pr\{\mathbf{n}_{t_{\text{asy}}} \in \mathcal{S}_i | N_0\}$ (again, with avoidance of extinction assumed implicitly), and states \mathbf{s}_i can be chosen arbitrarily from \mathcal{S}_i . This approximation allows us to verify our heuristic understanding of how the rates of growth and mutation impact fixation probabilities by controlling the duration and stochastic character of dynamical phases and therefore the likelihood of (stochastic) behaviours, such as crossing the separatrix, that are crucial in dictating long-term outcomes. It also dramatically reduces the computational time needed to calculate p_{AllID} for a range of different initial conditions, since we only need to recalculate the four likelihoods that the stochastically-induced regime finishes in each of the equivalence classes, respectively. The conditional probability $p_{\text{AllID}}^{\text{asy}}$, by contrast, does not depend on the specific initial conditions (but rather in which equivalence class the system is at time t_{asy}).

Growth and mutation

Fixing the mixed-strategy fixed point as the initial demographic mix, we can use Eq. (6) to deconstruct the dependence of p_{AllID} on the rates of growth and mutation (Fig. 7).

In the stochastically-induced regime, fluctuation gradients ‘drive’ trajectories from the centre of the mixed-strategy basin towards the simplex edges and then the corners. As a result, the growth-dependent (average) duration of the regime, $t^{\text{asy}} = \log(10^{2.7}/N_0)/(b-d)$, dictates in which class the trajectories are likely to start the asymmetric phase: rapid growth rates are required to confine trajectories that end in the mixed-strategy basin to the \mathcal{S}_4 region, whilst intermediate and slow growth rates suffice for the \mathcal{S}_3 and \mathcal{S}_2 regions, respectively (Figs. 7c-e). Although the average time spent in the

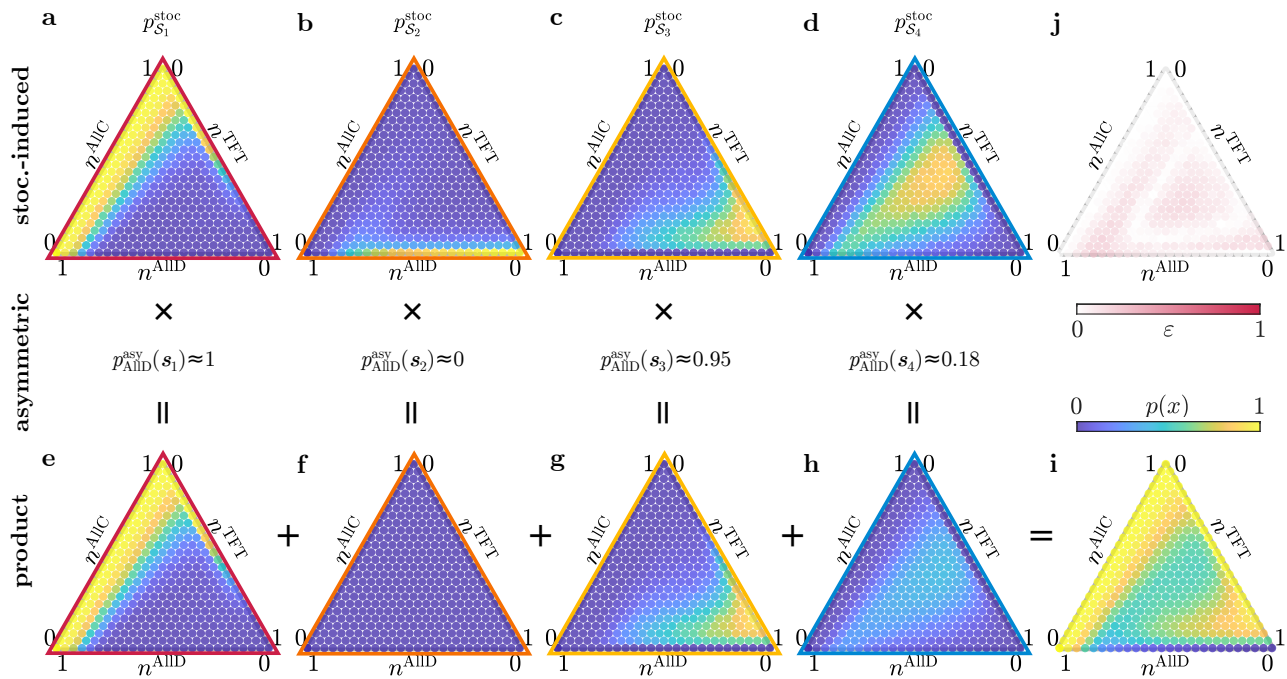


FIG. 8. The sensitivity to initial conditions can be reproduced by a decomposition based on effective phases and equivalence classes. For certain rates—*i.e.*, those below the red line, Fig. 3a—fixation probabilities can rely critically on initial conditions. That is, initial demographic mixes that are ‘close’ (and not necessarily near the separatrix) can give rise to dramatically different long-term fixation probabilities. This is particularly striking at low levels of mutation ($\log \mu = -6.5$) and comparatively high initial population size ($N_0 = 250$) (*cf.* Fig. 3f). The reason for this is twofold. First, larger initial populations reduce the effect of the stochastically-induced phase, which is to expel populations to the simplex boundaries and the AllD and TFT corners (panels c & d). Second, the duration of the asymmetric phase is significant and cannot be neglected (middle row). This implies that trajectories finishing the stochastically-induced phase in classes \mathcal{S}_3 or \mathcal{S}_4 can still cross the separatrix and fixate on the AllD fixed point (panels g & h). Whilst this provides a heuristic understanding, the contribution to p_{AllD} from the \mathcal{S}_4 region violates one of the assumptions of (6), and therefore the error increases with N_0 for the stated values of μ and $b - d$ (panel j and Supplementary Information, Figs. 19 & 20).

two basins by such trajectories during the stochastically-induced regime is independent of μ (Fig. 4a), the likelihood that the system is in \mathcal{S}_1 at time t_{asy} actually increases with decreasing μ (Fig. 7b). The reason is that decreasing μ changes the shape of the separatrix, therefore reducing the size of the \mathcal{S}_3 region (Supplementary Information, Sec. 10).

The asymmetric regime, by contrast, has a duration that is both growth- and μ -dependent: $\tau_\mu = t_{\text{lock}} - t_{\text{asy}} = -\log(\mu 10^{2.7}) / (b - d)$. Here, the probability of crossing the separatrix hinges, principally, on the likelihood of avoiding the TFT corner and its associated large confinement times (Figs. 7g-i). For example, τ_μ^{asy} must be extremely long in order to permit crossings from the \mathcal{S}_2 region, since populations starting the asymmetric regime from this region encounter the TFT corner with almost certainty. Crossings from the \mathcal{S}_4 region, however, occur at more modest τ_μ^{asy} —achieved by either low growth and high μ , or modest growth and low μ —reflecting the possibility that trajectories might avoid the TFT corner. Those from the \mathcal{S}_3 region can happen at the smallest τ_μ^{asy} , since there is a high likelihood that trajectories will avoid the TFT corner (recall the anti-clockwise dynam-

ics) and the ratio of the magnitude of stochastic effects to the magnitude of the deterministic flow is large in \mathcal{S}_3 (see Fig. 4f).

Combining these conditional probabilities using Eq. (6) satisfactorily reproduces the overall statistics of demographic outcomes (*cf.* Fig. 7n and Fig. 3a): the difference between our approximation and the full simulations have a mean value of 0.027, when averaged over growth and mutation rates, and a maximum value of 0.099 (Fig. 7o). This also confirms our assertion that the trajectories that start the asymmetric phase from the \mathcal{S}_4 region do not impact long-term outcomes. The reason is that this only happens with significant likelihood when growth rates are high, which simultaneously ensures that such trajectories never cross the separatrix (Figs. 7e, i & m). The same decomposition is shown for initial demographic mixes other than the mixed-strategy fixed point in Supplementary, Sec. 11.

Initial conditions

The approximation in Eq. (6) further provides insight into the founder-like dependence of fixation probabilities on initial conditions (Fig. 8 and Supplementary Information, Figs. 16-18). At high values of μ , the asymmetric phase has negligible duration and behaviour is trivial (Figs. 3b & c, Supplementary Information, Figs. 16 & 17). At low values of μ , however, the asymmetric regime cannot be ignored and has a significant bearing on fixation probabilities.

In this case, if the initial population size is small (Fig. 3e, Supplementary Information, Fig. 18), then the stochastically-induced phase is sufficiently long as to expel trajectories to the AllD or TFT corners from the AllD or mixed-strategy basins, respectively. Only the former trajectories impact fixation probabilities, however, since for all but the slowest growth rates, trajectories stuck in the TFT corner have a residence time longer than the duration of the asymmetric regime.

By contrast, if the initial population size is large (Fig. 3f and Fig. 8), then the duration of the stochastically-induced phase is not long enough to expel trajectories to the boundaries and/or corners, resulting in a non-zero likelihood of starting the asymmetric regime from either the \mathcal{S}_3 or \mathcal{S}_4 regions. For the former, there is a high probability of crossing to the AllD basin during the asymmetric regime, since the deterministic flows direct demographic trajectories towards the separatrix. For the latter, this probability is much lower, since trajectories are more likely to be entrained to the limit cycle. Nevertheless, the small contribution that results from the \mathcal{S}_4 region breaks one of the assumptions on which Eq. (6) is based, which also explains why the error in our decomposition increases with N_0 for certain values of μ and $b - d$ (the mean value of the error in Fig. 8j is 0.079, while the maximum value is 0.285). Of note, the values shown in Fig. 8 are the worst case of those we have simulated (Supplementary Information, Sec. 12).

III. DISCUSSION

Using a growing variant of an iconic model of evolutionary game theory, we have demonstrated the existence of non-trivial fluctuation-mediated effects in growing populations, whereby the rates of growth and mutation critically determine how fixation probabilities are conditioned on initial population size and demographic mix. In particular, we show that, for low mutation rates and intermediate growth rates, behaviour defies traditional monotonic expectations on the likelihood of fixations. This has ramifications for the understanding of population bottlenecks and their long-term impact. The implication being that the population growth commonly associated with post-bottleneck recovery can, in fact, be as important as, if not more than, the effects of the bottleneck on population size and demographic mix.

Our findings apply to populations that are well-mixed, and as such pertain to systems that combine short-range interactions with a mechanism for mixing that is fast on the timescales of the population dynamics, or those that otherwise have effectively long-range interactions, either explicitly or via the mutual interaction with public resource (although some of these assumptions have recently been brought into question [29]).

In this context, we follow several pioneering works that have characterised various effects of growth in well-mixed populations [10–14]. Although they do not explicitly examine the role of initial conditions, the behaviours appearing in two of these works [13, 14] are related to the those reported here. As are several studies concerning growing systems of binary ‘spins’ [30–32].

The shared mechanism in all these cases is the presence of $O(1/\sqrt{N})$ intrinsic fluctuations that decrease in time as the system grows. This has two ramifications: first, ergodicity is broken, implying that a population fixates to its deterministic attractors in long times, and; second, there is a decreasing scale by which fluctuations can couple dynamically to the underlying geometry of those attractors (and their basins). At a high level, it is the combination of these effects that gives rise to our headline behaviour: fixation probabilities that depend non-trivially on initial conditions (or otherwise stochastic events at early times), even as $t \rightarrow \infty$.

Such fluctuation-mediated effects are further characterised by a dependence on the rates of growth and mutation, which we explain by putting forward the notion of *effective* non-equilibrium phase transitions, and showing that these delimit distinct demographic regimes in our model. The rate of growth controls the overall rate at which intrinsic fluctuations decrease, and therefore also the relative duration of such demographic regimes. Mutation, by contrast, has two related effects. First, it changes the nature of the underlying deterministic attractors (the size of the mixed-strategy limit cycle, in our case) and hence the structure to which decreasing fluctuations couple. Second, it also changes the state-dependence of the fluctuations. Whilst the former impacts the character of the initial stochastically-induced phase, it is the combination of both of these effects that sets the $\sim 1/\mu$ dependence of the critical transition between the intermediate asymmetric and the final locked-in phases.

The latter behaviour is an example of a potentially interesting and unexpected antagonistic relationship between two sources of stochasticity: mutation and the intrinsic effects of finite-sized populations. In our model and other studies, mutation promotes heterogeneity, whilst the intrinsic fluctuations that arise from frequency-dependent birth drive the system towards homogeneity. This is particularly important in the context of population growth, since the effects of mutation do not depend on the population size, whereas the stochastically-induced forcing due to intrinsic noise decreases with increasing population size. As a result, we

speculate the other growing systems may also exhibit mutation-dependent critical transitions, where the effects of intrinsic fluctuations are balanced by those of mutation.

More generally, for growing well-mixed populations with all but the simplest of deterministic attractors—i.e., fixed points, limit cycles and stable manifolds etc.—the implication is that growth may be synonymous with effective non-equilibrium phase transitions. Understanding whether and how these fit into the existing literature is an open question. Of potential interest is the ongoing challenge to classify non-equilibrium phase transitions by their universality classes [18]. Work has been undertaken to describe single absorbing state transitions (of the directed-percolation type) [17, 33, 34] and also symmetric absorbing state transitions [35], but there appears to be very little literature on asymmetric absorbing state transitions.

Theoretical considerations aside, we put forward that our ideas may be examined within the context of directed evolution [36, 37]. Here, another, albeit direct, interplay between growth and mutation has already been demonstrated: mutations occurring at the genetic loci associated with growth-control promote so-called genetic instabilities [38]. Our results also appear relevant to state-of-the-art *in silico* representations of directed evolution, where the role of intrinsic fluctuations during growth has so far been overlooked [39].

A further setting that may prove relevant is that of RNA viruses, which need to survive many population bottlenecks involved in host-to-host, as well as intra-host (e.g., plaque-to-plaque), transmission [3, 4, 40]. Due to their high mutation rate, these viruses typically consist of a whole spectra of nucleotide sequences—a so-called quasispecies [40]. After each transmission, a small number of viral particles need to be capable of restoring either the original mutant spectra, or a different one that is adapted to a new host or environment. The success of RNA viruses in overcoming this challenge is hypothesised to be related to their high mutation rate, which is estimated to be very close for some viruses to the theoretical error threshold above which viral identity cannot be maintained [40]. We speculate that a previously unappreciated basis for such high mutation rates could be the

stochastic effects demonstrated in this study, which drive small replicating populations towards homogeneity. In particular, during the initial stages of post-bottleneck recovery, a viral population will become increasingly dominated by a single sequence chosen mostly by chance. In this context, high mutation rates are seen more as a ‘correction’ for the action of finite-size effects, rather than for the compositional restriction due to the bottleneck itself. Answering this question and more is left for future work.

Nevertheless, bridging the gap between the abstract setting of the present work and the aforementioned applications will undoubtedly involve significant work. The extent to which this will be possible remains an open question, and may hinge on features that are not included in our model, such as spatial structure and/or other physical constraints [41–44]. So-called “patch” models of interacting locally well-mixed subpopulations is one avenue that may prove promising. Exploring how the ideas set out here translate across a wider class of systems is therefore an important avenue of future research, and we welcome further work in the area.

AUTHOR CONTRIBUTIONS

EC and JP performed simulations and analysis, under the guidance of ST and RM. EC and RM wrote the manuscript, with help from JP and ST. The project initially arose from discussions between JP, ST and RM.

ACKNOWLEDGMENTS

The authors would like to thank V. Guttal for helpful discussions at the project’s outset. RGM and EC acknowledge EMBL Australia for funding. We further acknowledge support from the Simons Foundation (Grant No. 287975 to ST), the Max Planck Society through a Max-Planck-Partner-Group at NCBS-TIFR (ST) and the Department of Atomic Energy, Government of India, under projects RTI4001 and RTI4006. This research includes computations using the computational cluster Katana supported by Research Technology Services at UNSW Sydney.

-
- [1] Mayr, E. Change of genetic environment and evolution (1954).
 - [2] Nei, M., Maruyama, T. & Chakraborty, R. The Bottleneck Effect and Genetic Variability in Populations. *Evolution* **29**, 1 (1975). URL <https://www.jstor.org/stable/2407137?origin=crossref>.
 - [3] Domingo, E., Sheldon, J. & Perales, C. Viral quasispecies evolution. *Microbiology and Molecular Biology Reviews* **76**, 159–216 (2012).
 - [4] Weaver, S. C., Forrester, N. L., Liu, J. & Vasilakis, N. Population bottlenecks and founder effects: implications for mosquito-borne arboviral emergence. *Nature Reviews Microbiology* **19**, 184–195 (2021).
 - [5] Keller, L. F. *et al.* Immigration and the ephemerality of a natural population bottleneck: evidence from molecular markers. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **268**, 1387–1394 (2001). URL <https://royalsocietypublishing.org/doi/10.1098/rspb.2001.1607>.
 - [6] Keller, L. F., Arcese, P., Smith, J. N. M., Hochachka, W. M. & Stearns, S. C. Selection against inbred song sparrows during a natural population bottleneck. *Nature*

- ture **372**, 356–357 (1994). URL <http://www.nature.com/articles/372356a0>.
- [7] Doherty, P. J. *et al.* High mortality during settlement is a population bottleneck for a tropical surgeonfish. *Ecology* **85**, 2422–2428 (2004). URL <http://doi.wiley.com/10.1890/04-0366>.
- [8] Hoelzel, A. R., Fleischer, R. C., Campagna, C., Le Boeuf, B. J. & Alvord, G. Impact of a population bottleneck on symmetry and genetic diversity in the northern elephant seal. *Journal of Evolutionary Biology* **15**, 567–575 (2002). URL <http://doi.wiley.com/10.1046/j.1420-9101.2002.00419.x>.
- [9] Manica, A., Amos, W., Balloux, F. & Hanihara, T. The effect of ancient population bottlenecks on human phenotypic variation. *Nature* **448**, 346–348 (2007). URL <http://www.nature.com/articles/nature05951>.
- [10] Melbinger, A., Cremer, J. & Frey, E. Evolutionary Game Theory in Growing Populations. *Physical Review Letters* **105**, 178101 (2010). URL <https://link.aps.org/doi/10.1103/PhysRevLett.105.178101>.
- [11] Cremer, J., Melbinger, A. & Frey, E. Evolutionary and population dynamics: A coupled approach. *Physical Review E* **84**, 051921 (2011). URL <https://link.aps.org/doi/10.1103/PhysRevE.84.051921>.
- [12] Huang, W., Hauert, C. & Traulsen, A. Stochastic game dynamics under demographic fluctuations. *Proceedings of the National Academy of Sciences* **112**, 9064–9069 (2015). URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1418745112>.
- [13] Constable, G. W. A., Rogers, T., McKane, A. J. & Tarnita, C. E. Demographic noise can reverse the direction of deterministic selection. *Proceedings of the National Academy of Sciences* **113**, E4745–E4754 (2016). URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1603693113>.
- [14] Ashcroft, P., Smith, C. E., Garrod, M. & Galla, T. Effects of population growth on the success of invading mutants. *Journal of Theoretical Biology* **420**, 232–240 (2017). URL <https://linkinghub.elsevier.com/retrieve/pii/S0022519317301261>.
- [15] Imhof, L. A., Fudenberg, D. & Nowak, M. A. Evolutionary cycles of cooperation and defection. *Proceedings of the National Academy of Sciences* **102**, 10797–10800 (2005). URL <http://www.pnas.org/cgi/doi/10.1073/pnas.0502589102>.
- [16] Bladon, A. J., Galla, T. & McKane, A. J. Evolutionary dynamics, intrinsic noise, and cycles of cooperation. *Physical Review E* **81**, 066122 (2010). URL <https://link.aps.org/doi/10.1103/PhysRevE.81.066122>.
- [17] Hinrichsen, H. Non-equilibrium critical phenomena and phase transitions into absorbing states. *Advances in Physics* **49**, 815–958 (2000). URL <http://www.tandfonline.com/doi/abs/10.1080/00018730050198152>.
- [18] Ódor, G. Universality classes in nonequilibrium lattice systems. *Reviews of Modern Physics* **76**, 663–724 (2004). URL <https://link.aps.org/doi/10.1103/RevModPhys.76.663>.
- [19] Broadbent, S. R. & Hammersley, J. M. Percolation processes. *Mathematical Proceedings of the Cambridge Philosophical Society* **53**, 629–641 (1957). URL <https://www.cambridge.org/core/product/identifier/S0305004100032680/type/journal-article>.
- [20] Corté, L., Chaikin, P. M., Gollub, J. P. & Pine, D. J. Random organization in periodically driven systems. *Nature Physics* **4**, 420–424 (2008). URL <http://www.nature.com/articles/nphys891>.
- [21] Tu, Y., Grinstein, G. & Muñoz, M. A. Systems with Multiplicative Noise: Critical Behavior from KPZ Equation and Numerics. *Physical Review Letters* **78**, 274–277 (1997). URL <https://link.aps.org/doi/10.1103/PhysRevLett.78.274>.
- [22] Cardy, J. L. & Grassberger, P. Epidemic models and percolation. *Journal of Physics A: Mathematical and General* **18**, L267–L271 (1985). URL <https://iopscience.iop.org/article/10.1088/0305-4470/18/6/001>.
- [23] Wilken, S., Guerra, R. E., Levine, D. & Chaikin, P. M. Random Close Packing as a Dynamical Phase Transition. *Physical Review Letters* **127**, 038002 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.038002>.
- [24] Smith, D. & Betbeder-Matibet, L. Katana. Tech. Rep. (2010).
- [25] Touchette, H. Introduction to dynamical large deviations of markov processes. *Physica A: Statistical Mechanics and its Applications* **504**, 5–19 (2018).
- [26] van Kampen, N. G. *Stochastic processes in physics and chemistry* (Elsevier, Amsterdam, 1992).
- [27] Horsthemke, W. & Lefever, R. *Noise-Induced Transitions*, vol. 15 of *Springer Series in Synergetics* (Springer Berlin Heidelberg, 2006). URL <http://link.springer.com/10.1007/3-540-36852-3>.
- [28] Jhavar, J. *et al.* Noise-induced schooling of fish. *Nature Physics* (2020). URL <http://www.nature.com/articles/s41567-020-0787-y>.
- [29] Herrerías-Azcué, F., Pérez-Muñuzuri, V. & Galla, T. Stirring does not make populations well mixed. *Scientific Reports* **8**, 4068 (2018). URL <http://www.nature.com/articles/s41598-018-22062-w>.
- [30] Morris, R. G. & Rogers, T. Growth-induced breaking and unbreaking of ergodicity in fully-connected spin systems. *Journal of Physics A: Mathematical and Theoretical* **47**, 342003 (2014). URL <http://stacks.iop.org/1751-8121/47/i=34/a=342003?key=crossref.80b743c6ffee60b17655974228dcf8f9>.
- [31] Klymko, K., Garrahan, J. P. & Whitelam, S. Similarity of ensembles of trajectories of reversible and irreversible growth processes. *Physical Review E* **96**, 042126 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevE.96.042126>.
- [32] Jack, R. L. Large deviations in models of growing clusters with symmetry-breaking transitions. *Physical Review E* **100**, 012140 (2019). URL <https://link.aps.org/doi/10.1103/PhysRevE.100.012140>.
- [33] Janssen, H.-K. On the nonequilibrium phase transition in reaction-diffusion systems with an absorbing stationary state. *Zeitschrift für Physik B Condensed Matter* **42**, 151–154 (1981).
- [34] Grassberger, P. On Phase Transitions in Schlögl’s Second Model. 262–262 (1981). URL http://link.springer.com/10.1007/978-3-642-81778-6_{_}49.
- [35] Al Hammal, O., Chaté, H., Dornic, I. & Muñoz, M. A. Langevin Description of Critical Phenomena with Two Symmetric Absorbing States. *Physical Review Letters* **94**, 230601 (2005). URL <https://link.aps.org/doi/10.1103/PhysRevLett.94.230601>.

- [36] Romero, P. A. & Arnold, F. H. Exploring protein fitness landscapes by directed evolution. *Nature reviews Molecular cell biology* **10**, 866–876 (2009).
- [37] Tracewell, C. A. & Arnold, F. H. Directed enzyme evolution: climbing fitness peaks one amino acid at a time. *Current opinion in chemical biology* **13**, 3–9 (2009).
- [38] Coelho, M. C., Pinto, R. M. & Murray, A. W. Heterozygous mutations cause genetic instability in a yeast model of cancer evolution. *Nature* **566**, 275–278 (2019). URL <http://www.nature.com/articles/s41586-019-0887-y>.
- [39] Chang, C.-Y. *et al.* Engineering complex communities by directed evolution. *Nature ecology & evolution* **5**, 1011–1023 (2021).
- [40] Domingo, E., García-Crespo, C. & Perales, C. Historical perspective on the discovery of the quasispecies concept. *Annual Review of Virology* **8**, 51–72 (2021).
- [41] Lieberman, E., Hauert, C. & Nowak, M. A. Evolutionary dynamics on graphs. *Nature* **433**, 312–316 (2005).
- [42] Allesina, S. & Levine, J. M. A competitive network theory of species diversity. *Proceedings of the National Academy of Sciences* **108**, 5638–5642 (2011).
- [43] Hindersin, L. & Traulsen, A. Most undirected random graphs are amplifiers of selection for birth-death dynamics, but suppressors of selection for death-birth dynamics. *PLoS computational biology* **11**, e1004437 (2015).
- [44] Marrec, L., Lamberti, I. & Bitbol, A.-F. Toward a universal model for spatially structured populations. *Physical Review Letters* **127**, 218102 (2021).
- [45] Gillespie, D. T. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry* **81**, 2340–2361 (1977).
- [46] Gillespie, D. T. The chemical Langevin equation. *The Journal of Chemical Physics* **113**, 297–306 (2000). URL <http://aip.scitation.org/doi/10.1063/1.481811>.
- [47] Gardiner, C. W. *Handbook of stochastic methods* (Springer, 2003), 3rd edn.
-

SUPPLEMENTARY INFORMATION

1. Gillespie-Itô simulations

The Gillespie algorithm [45, 46] allows the exact simulation of the stochastic dynamics of the growing IPD when N is small, *i.e.*, when such dynamics cannot be approximated by simply integrating the SDEs (see Secs. 3 & 4). The downside of this algorithm is that its computational time scales linearly with N , becoming impractical as the population grows. We therefore adopt a hybrid approach: when N is smaller than a chosen threshold the system's dynamics are simulated with the Gillespie algorithm, and when the threshold is exceeded they are simulated by numerically integrating the SDEs (Euler-Maruyama). Since the latter method results in values of N^i along the real line, rounding is required if stochastic fluctuations trigger a switch back to the Gillespie algorithm.

Our choice for the algorithm switch threshold is $N = \mu^{-1}$. This is motivated by the existence of the locked-in phase, which begins when $N \approx \mu^{-1}$ (see Fig. 4). During this phase, the risk that the SDE approximation would lead to an 'accidental' crossing of the separatrix is extremely small. A minimum threshold of $N = 10,000$ is used for simulations with large μ .

All results were obtained using the High Performance Computing facility Katana [24]. The results in Fig. 3a (and Figs. 13-15, panel **p**), were obtained with 10,000 repetitions of the hybrid Gillespie-Itô simulations for each combination of b and μ (fixing $d = 1$). Each repetition was carried out until $N > 10^{10}$. The results in Fig. 3b-g were similarly obtained with 1000 repetitions for each initial condition. The results in Fig. 7b-i (and Figs. 13-15, panels **b-i**) were obtained with 1000 repetitions for each combination of μ and b . 1000 repetitions were also used to obtain the results in Figs. 8a-d and Sec. 12.

The fixed-size dynamics were also simulated using the hybrid Gillespie-Itô approach. The Gillespie algorithm was modified by setting $d = 0$ and 'killing' a randomly selected player at every birth. The Itô part of the algorithm involved replacing the noise correlation matrix (see details in Sec. 7). The results in Fig. 4 were obtained by running the fixed-size Gillespie-Itô algorithm 1000 times for 20,000 seconds for each combination of μ and N .

2. The deterministic limit

In the $N \rightarrow \infty$ limit, the behaviour of the growing Iterated Prisoner Dilemma described in the main text is equivalent to that of the following continuous, deterministic system:

$$\frac{dn^i}{dt} = \frac{b\mu}{\varphi} \sum_{j \neq i} (f_j n^j - f_i n^i) + \left(\frac{bf_i}{\varphi} - d \right) n^i, \quad (7)$$

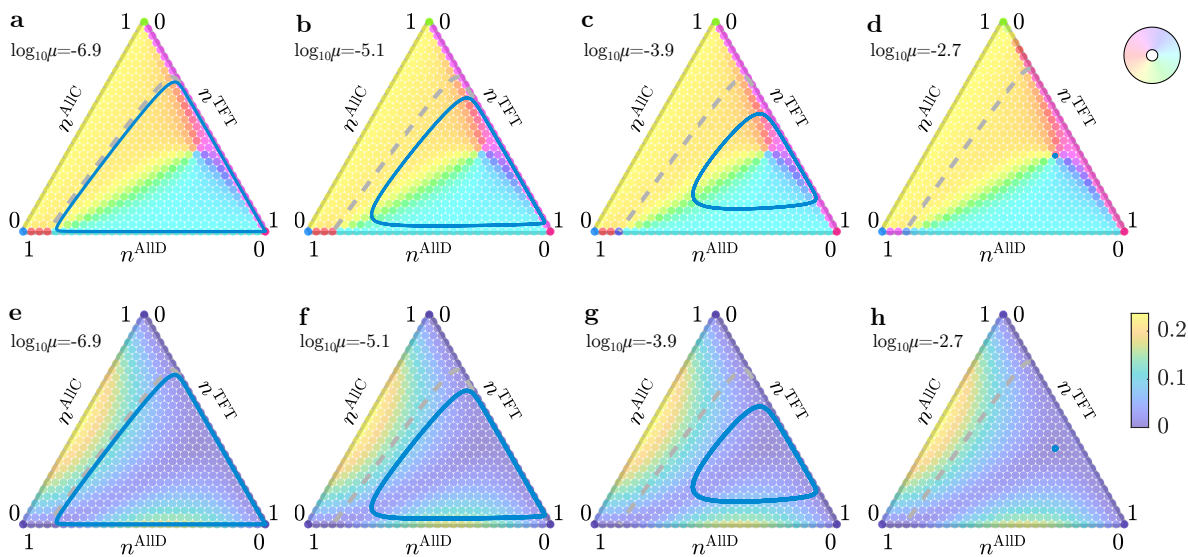


FIG. 9. Population dynamics in the deterministic limit of the growing IPD at different values of μ . The top row (panels a-d) show the direction of the deterministic flow, while the bottom row (panels e-h) show the magnitude (speed). For all panels, the growth rate is $b - d = 0.05$.

where $n^i = N^i/N$. These equations can be easily derived from Fig. 1.

As shown in Fig. 9 (see also [15]), the mutation rate has a clear effect on the dynamics of the continuous system. For $10^{-7.5} \leq \mu \leq 10^{-2.5}$, the system exhibits two stable attractors: a fixed point characterised by a population of almost entirely ALLD players and a stable mixed-strategy limit cycle. The most notable effect of varying μ within this range is the resizing of the limit cycle: the smaller the mutation rate, the larger the limit cycle (panels **a-c** and **e-g**). For large values of μ (approximately $10^{-3.1} \leq \mu \leq 10^{-2.5}$), the limit cycle collapses into a point, although the separatrix still exists (panels **d** and **h**). Varying μ also slightly changes the shape of the separatrix, especially in the proximity of the ALLC-TFT edge of the simplex.

3. System-Size Expansion

Following [26, 47], we derive a set of coupled SDEs that approximate the dynamics of the underlying protocol when $N \gg 1$. In terms of step-operators, the Master-equation has the form:

$$\begin{aligned} \frac{dP(\vec{N}; t)}{dt} &= \sum_i (\mathbb{E}_{N^i}^{-1} - 1) \frac{bf_i N^i}{\varphi} P(\vec{N}; t) \\ &+ \sum_i (\mathbb{E}_{N^i}^{+1} - 1) dN^i P(\vec{N}; t) \\ &+ \sum_i \sum_{j \neq i} (\mathbb{E}_{N^i}^{+1} \mathbb{E}_{N^j}^{-1} - 1) \frac{b\mu f_i N^i}{\varphi} P(\vec{N}; t). \end{aligned} \quad (8)$$

Expanding the step-operators in the usual fashion, and retaining only the leading and next-to-leading order terms, gives

$$\mathbb{E}_{N^i}^r = 1 + r \frac{\partial}{\partial N^i} + \frac{r^2}{2} \frac{\partial^2}{\partial (N^i)^2} + O(1), \quad (9)$$

where $r = \pm 1$. Substituting Eq. (9) into Eq. (8) results in an equation of the Fokker-Planck type [47]:

$$\begin{aligned} \frac{\partial P(\vec{N}; t)}{\partial t} &= - \sum_i \frac{\partial}{\partial N^i} \left[A_i(\vec{N}) P(\vec{N}; t) \right] \\ &+ \frac{1}{2} \sum_i \sum_j \frac{\partial^2}{\partial N^i \partial N^j} \left[B_{ij}(\vec{N}) P(\vec{N}; t) \right]. \end{aligned} \quad (10)$$

After some manipulation, it can be shown that

$$A_i = \frac{b(1-2\mu)f_i N^i}{\varphi} + \sum_{j \neq i} \frac{b\mu f_j N^j}{\varphi} - dN^i, \quad (11)$$

and

$$B_{ij} = \begin{cases} \frac{b(1+2\mu)f_i N^i}{\varphi} + \sum_{k \neq i} \frac{b\mu f_k N^k}{\varphi} + dN^i & \text{if } i = j \\ -\frac{b\mu}{\varphi} (f_i N^i + f_j N^j) & \text{if } i \neq j \end{cases}. \quad (12)$$

Equation (10) implies an SDE for the variables N^i :

$$\frac{dN^i}{dt} = A_i + \xi_i. \quad (13)$$

The deterministic part of Eq. (13) is equivalent to Eq. (7). The noise sources that appear in Eq. (13) have zero mean (*i.e.*, $\langle \xi_i \rangle = 0$) and are correlated according to:

$$\langle \xi_i(t) \xi_j(t') \rangle = B_{ij} \delta(t - t'). \quad (14)$$

Summing over index i in Eq. (13), we have

$$\frac{dN}{dt} = (b-d)N + \sqrt{(b+d)N} \xi, \quad (15)$$

where ξ is a single source of zero-mean Gaussian white noise.

4. Multiplicative delta-correlated noise

Eqn. (13) can be re-written in terms of delta-correlated noise and multiplicative pre-factors. To do this we must choose a matrix \mathbf{b} that satisfies $\mathbf{B} = \mathbf{b}^T \cdot \mathbf{b}$. For \mathbf{b} to be square, this requires a Cholesky decomposition, and ensures no more independent noise sources than there are variables in the system. However, a more natural approach is to decompose \mathbf{B} according to the rules set out in Gillespie's CLE approach [46]. This results in SDEs of the following form:

$$\frac{dN^i}{dt} = A_i + \sum_{\alpha} b_{\alpha i} \eta_{\alpha}, \quad (16)$$

where the index α runs from $1 \dots 12$ and correspond to the reactions in Fig. 1. The twelve independent noise sources each have mean zero, and are delta-correlated— *i.e.*, $\langle \eta_{\alpha} \rangle = 0$ and $\langle \eta_{\alpha}(t) \eta_{\beta}(t') \rangle = \delta_{\alpha\beta} \delta(t - t')$. The matrix \mathbf{b}^T is given by

$$b_{\alpha i} = \begin{pmatrix} \sqrt{\frac{bf_1 N^1}{\varphi}} & 0 & 0 \\ 0 & \sqrt{\frac{bf_2 N^2}{\varphi}} & 0 \\ 0 & 0 & \sqrt{\frac{bf_3 N^3}{\varphi}} \\ -\sqrt{d N^1} & 0 & 0 \\ 0 & -\sqrt{d N^2} & 0 \\ 0 & 0 & -\sqrt{d N^3} \\ -\sqrt{\frac{b\mu f_1 N^1}{\varphi}} & \sqrt{\frac{b\mu f_1 N^1}{\varphi}} & 0 \\ \sqrt{\frac{b\mu f_2 N^2}{\varphi}} & -\sqrt{\frac{b\mu f_2 N^2}{\varphi}} & 0 \\ 0 & -\sqrt{\frac{b\mu f_2 N^2}{\varphi}} & \sqrt{\frac{b\mu f_2 N^2}{\varphi}} \\ 0 & \sqrt{\frac{b\mu f_3 N^3}{\varphi}} & -\sqrt{\frac{b\mu f_3 N^3}{\varphi}} \\ -\sqrt{\frac{b\mu f_1 N^1}{\varphi}} & 0 & \sqrt{\frac{b\mu f_1 N^1}{\varphi}} \\ \sqrt{\frac{b\mu f_3 N^3}{\varphi}} & 0 & -\sqrt{\frac{b\mu f_3 N^3}{\varphi}} \end{pmatrix}. \quad (17)$$

5. Projected dynamics

We wish to project the dynamics onto the unit simplex— *i.e.*, in terms of variables $n^i = N^i / \sum_i N^i$. For this, we require the multivariate form of Itô's lemma [47]. For finite N , we have

$$\frac{dn^i}{dt} = \sum_k A_k \partial_k n^i + \frac{1}{2} \sum_{k,j} B_{kj} \partial_k \partial_j n^i + \sum_j \sum_{\alpha} b_{j\alpha} (\partial_j n^i) \eta_{\alpha}, \quad (18)$$

where the shorthand $\partial_i = \partial / \partial N^i$ has been used. First, we deal with the deterministic parts. Using

$$\frac{\partial n^i}{\partial N^j} = \begin{cases} \frac{1-n^i}{N} & \text{if } i = j \\ -\frac{n^i}{N} & \text{if } i \neq j \end{cases}, \quad (19)$$

alongside (11), it can be shown that

$$\sum_k A_k \partial_k n^i = \frac{b\mu}{\varphi} \left(\sum_{j \neq i} f_j n^j - 2f_i n^i \right) + b \left(\frac{f_i}{\varphi} - 1 \right) n^i. \quad (20)$$

Similarly, using

$$\frac{\partial^2 n^i}{\partial N^j \partial N^k} = \begin{cases} \frac{2(n^i-1)}{N^2} & \text{if } i = j = k \\ \frac{2n^i-1}{N^2} & \text{if } j = i \neq k \text{ or } k = i \neq j \\ \frac{2n^i}{N^2} & \text{if } j \neq i \text{ and } k \neq i \end{cases}, \quad (21)$$

gives

$$\frac{1}{2} \sum_{k,j} B_{kj} \partial_k \partial_j n^i = -b \left(\frac{f_i}{\varphi} - 1 \right) \frac{n^i}{N}. \quad (22)$$

Notice that, when summed over i , both (20) and (22) are zero by virtue of the fact that $\sum_i f_i n^i / \varphi = 1$. Also, trivially, (22) goes to zero in the deterministic $N \rightarrow \infty$ limit. For the stochastic part of (18), define a new matrix $b_{\alpha i}^\dagger = \sum_j b_{\alpha j} (\partial_j n^i)$ such that (18) can be re-cast in terms of correlated noise sources— *i.e.*,

$$\frac{dn^i}{dt} = \frac{b\mu}{\varphi} \left(\sum_{j \neq i} f_j n^j - 2f_i n^i \right) + b \left(\frac{f_i}{\varphi} - 1 \right) n^i \left(1 - \frac{1}{N} \right) + \zeta_i, \quad (23)$$

with $\langle \zeta_i \rangle = 0$ and $\langle \zeta_i(t) \zeta_j(t') \rangle = B_{ij}^\dagger \delta(t - t')$, where

$$B_{ij}^\dagger = \sum_{\alpha} b_{\alpha i}^\dagger b_{\alpha j}^\dagger = \begin{cases} \frac{b\mu}{N} \left[n^i + \frac{f_i}{\varphi} (1 - 2n^i) \right] + \frac{b\mu}{N} \left(1 + \frac{f_i n^i}{\varphi} \right) + \frac{dn^i}{N} (1 - n^i) & \text{if } i = j \\ -\frac{b\mu}{N} \frac{n^i}{\varphi} \left(\frac{f_i}{\varphi} + \frac{f_j}{\varphi} - 1 \right) - \frac{b\mu}{N\varphi} (f_i n^i + f_j n^j) - \frac{dn^i n^j}{N} & \text{if } i \neq j \end{cases}, \quad (24)$$

To understand the impact of these correlations, (24) can be computationally decomposed into an eigenbasis for different values of n^i . This always results in one zero-eigenvalue eigenvector pointing perpendicular to the simplex. The remaining in-simplex eigenvalues reveal that populations towards the centre of the simplex experience large uncorrelated fluctuations whilst, closer to the simplex boundary, correlations suppress fluctuations in the direction normal to the boundary (Fig. 4b-c, orange crosses). Moreover, the magnitude of along-boundary fluctuations decrease as a corner is approached.

6. Boundary effects

Equation (23) can be evaluated at the simplex edges, where we are particularly interested in the both deterministic drift and fluctuations in the direction of the bulk, which is captured by the dynamics of the strategy who's concentration is zero along a given edge.

- AllD-AllC edge: setting $n^3 = 0$ and $n^2 = 1 - n^1$, gives $dn^3/dt = b\mu + \zeta_3$, where

$$\langle \zeta_3 \zeta_j \rangle = \begin{pmatrix} \frac{b\mu n^1 [n^1(R-S)+S]}{N\{P(n^1-1)^2+n^1[n^1(R-S-T)+S+T]\}} \\ \frac{b\mu(n^1-1)[(n^1-1)P-n^1T]}{N\{P(n^1-1)^2+n^1[n^1(R-S-T)+S+T]\}} \\ \frac{b\mu}{N} \end{pmatrix}. \quad (25)$$

- AllC-TFT edge: setting $n^2 = 0$ and $n^1 = 1 - n^3$, gives $dn^2/dt = b\mu + \zeta_2$, where

$$\langle \zeta_2 \zeta_j \rangle = \begin{pmatrix} \frac{b\mu m(n^3-1)R}{N(mR-cn^3)} \\ \frac{b\mu}{N} \\ \frac{b\mu n^3(mR-c)}{N(cn^3-mR)} \end{pmatrix}. \quad (26)$$

- TFT-AllD edge: setting $n^1 = 0$ and $n^3 = 1 - n^2$, gives $dn^1/dt = b\mu + \zeta_1$, where

$$\langle \zeta_1 \zeta_j \rangle = \begin{pmatrix} \frac{b\mu}{N} \\ \frac{b\mu n^2 [P(m+n^2-1)+T(n^2-1)]}{N\{c(n^2-1)-(n^2)^2[(m-2)P-mR+S+T]+n^2[2(m-1)P-2mR+S+T]+mR\}} \\ \frac{b\mu(n^2-1)\{-c+n^2[(m-1)P-mR+S]+mR\}}{N\{c(n^2-1)-(n^2)^2[(m-2)P-mR+S+T]+n^2[2(m-1)P-2mR+S+T]+mR\}} \end{pmatrix}. \quad (27)$$

In all three of the above cases, the deterministic repulsion from the edge is $O(\mu)$. The fluctuations (positive and negative) in the same direction are $O(\sqrt{\mu/N})$. The implication, in the context of the convex hull analysis of the main manuscript, is that, very close to the AllD-TFT edge, fluctuations can overcome the deterministic forces only if $N \sim 1/\mu$.

7. Langevin equation for the fixed-size model

The dynamics of the system were simulated using the hybrid Gillespie-Itô approach (see Sec. 1). A fixed-size version of the system was also used in our analysis (see Figs. 4 & 5): The Gillespie algorithm was modified by setting $d = 0$ and ‘killing’ a randomly selected player at every birth, while the Itô part of the algorithm involved replacing the matrix \mathbf{b}^T in Eq. (17) with

$$b_{\alpha i} = \begin{pmatrix} (1-n^1)\sqrt{\frac{bf_1N^1}{\varphi}} & -n^1\sqrt{\frac{bf_2N^2}{\varphi}} & -n^1\sqrt{\frac{bf_3N^3}{\varphi}} \\ -n^2\sqrt{\frac{bf_1N^1}{\varphi}} & (1-n^2)\sqrt{\frac{bf_2N^2}{\varphi}} & -n^2\sqrt{\frac{bf_3N^3}{\varphi}} \\ -n^3\sqrt{\frac{bf_1N^1}{\varphi}} & -n^3\sqrt{\frac{bf_2N^2}{\varphi}} & (1-n^3)\sqrt{\frac{bf_3N^3}{\varphi}} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ -\sqrt{\frac{b\mu f_1N^1}{\varphi}} & \sqrt{\frac{b\mu f_1N^1}{\varphi}} & 0 \\ \sqrt{\frac{b\mu f_2N^2}{\varphi}} & -\sqrt{\frac{b\mu f_2N^2}{\varphi}} & 0 \\ 0 & -\sqrt{\frac{b\mu f_2N^2}{\varphi}} & \sqrt{\frac{b\mu f_2N^2}{\varphi}} \\ 0 & \sqrt{\frac{b\mu f_3N^3}{\varphi}} & -\sqrt{\frac{b\mu f_3N^3}{\varphi}} \\ -\sqrt{\frac{b\mu f_1N^1}{\varphi}} & 0 & \sqrt{\frac{b\mu f_1N^1}{\varphi}} \\ \sqrt{\frac{b\mu f_3N^3}{\varphi}} & 0 & -\sqrt{\frac{b\mu f_3N^3}{\varphi}} \end{pmatrix}. \quad (28)$$

8. Escape time from corners in the stochastically-induced regime

Here we show the statistics of the time t^\dagger that is needed by the fixed-size system to cross the separatrix for the first time (Fig. 10). We compare the cases of the systems starting from the AllD corner (panels **a** and **c**) and the TFT corner (panels **b** and **d**). In both cases, t^\dagger is proportional to μ : in the order of thousands of seconds when $\log_{10} \mu = -3.5$ (panels **a-b**) and in the order of millions of seconds when $\log_{10} \mu = -6.5$ (panels **c-d**). However, the influence of the population size N_0 on t^\dagger is different in the two cases. On the one hand, the crossing time starting from the AllD corner increases very rapidly when N_0 approaches the onset of the asymmetric regime at around $N_0 = 10^{2.7}$ (panels **a** and **c**), indeed showing that crossing the separatrix from the AllD basin to the mixed-strategy basin becomes extremely unlikely in the asymmetric regime. On the other hand, the crossing time starting from the TFT corner decreases with N_0 .

9. Equivalence classes decomposition

Consider a trajectory of the growing IPD that starts from an arbitrary state in the stochastically-induced phase (i.e., $N_0 < 10^{2.7}$) at time $t = 0$, enters the asymmetric phase at time $t = t_{\text{asy}}$ and the locked-in phase at time $t = t_{\text{lock}}$. Due to the nature of the locked-in phase, the evolutionary outcome at $t \rightarrow \infty$ is already known at time t_{lock} , and thus Eq. (3) can be simplified:

$$p_{\text{AllD}} = \Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{N}_0\}, \quad (29)$$

where avoidance of extinction is implicitly assumed and AllD-b indicates the AllD basin of the simplex.

Moreover, since the growing IPD is a Markovian process, p_{AllD} can be decomposed as follows:

$$p_{\text{AllD}} = \sum_{\mathbf{s} \in U} \Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\} \Pr\{\mathbf{n}_{t_{\text{asy}}} = \mathbf{s} \mid \mathbf{N}_0\}, \quad (30)$$

where U is the set of all possible states of the system at time t_{asy} .

Eq. (30) is impractical, since it requires the numerical estimation of the probabilities $\Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\}$ and $\Pr\{\mathbf{n}_{t_{\text{asy}}} = \mathbf{s} \mid \mathbf{N}_0\}$ for all \mathbf{s} in the very large set, U . However, the approximation given in Eq. (6) can be made by inspecting $\Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\}$ in Fig. 11a, which was numerically estimated for a subset of all points \mathbf{s} (for illustration purposes we only show the case of $b - d = 0.05$ and $\log_{10} \mu = -6.5$).

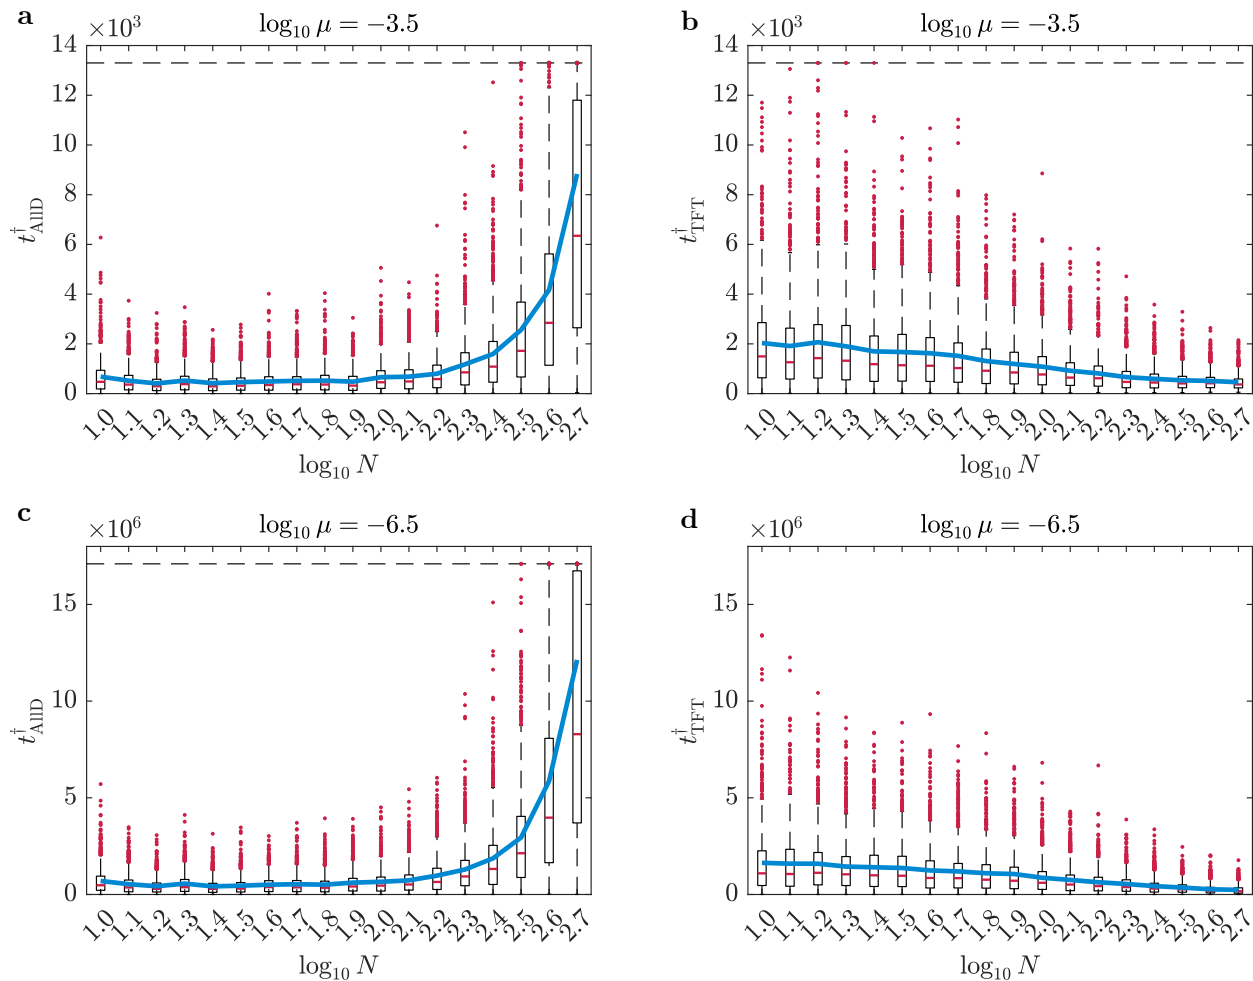


FIG. 10. Statistics of separatrix crossing time in the stochastically induced regime. The figure shows the statistics of the time taken by the system to cross the separatrix for the first time using the fixed-size model with $\mathbf{n}_0 = \mathbf{s}_1$, i.e. the AllD corner (panels **a** for $\log_{10} \mu = -3.5$ and **c** for $\log_{10} \mu = -6.5$) and with $\mathbf{n}_0 = \mathbf{s}_2$, i.e. the TFT corner (panels **b** for $\log_{10} \mu = -3.5$ and **d** for $\log_{10} \mu = -6.5$). The statistics, for each of the four cases, are obtained from 10,000 repetitions. On each box, the central red mark indicates the median, and the bottom and top edges indicate the 25th and 75th percentiles, respectively, the whiskers extend to the most extreme data points not considered outliers, and the red dots are the outliers (a data value is considered an outlier if it is greater than $Q_3 + 1.5(Q_3 - Q_1)$ or less than $Q_1 - 1.5(Q_3 - Q_1)$, where Q_1 and Q_3 are the 25th and 75th percentiles, respectively). For an easier visualisation of the statistics, all data values exceeding a maximum threshold are collapsed into the horizontal dashed lines on the top part of each plot. The blue lines represent the average crossing time over all repetitions.

A first area of the simplex, \mathcal{S}_1 (outlined in red), corresponding to the AllD basin, can be immediately identified as it is characterised by a homogenous probability of AllD outcomes of approximately 1. If the system is anywhere within \mathcal{S}_1 at time t_{asy} , then it is expected never to move to the mixed-strategy basin, since the crossing of the separatrix in this direction is extremely unlikely. Any point \mathbf{s}_1 would be good candidate for representing the entire area \mathcal{S}_1 , however, we chose the AllD corner (red dot in Fig. 11b) as the system is in the proximity of such point for the vast majority of the time spent in \mathcal{S}_1 .

A second area can similarly be identified: \mathcal{S}_2 (orange), corresponding to the mixed-strategy part of the AllD-TFT edge (including the TFT corner) and characterised by a homogenous probability of AllD outcomes of approximately 0. If the system is in this area at time t_{asy} , then it is very unlikely to cross the separatrix because of the long time (inversely proportional to μ) spent in the TFT corner while growing (i.e., as the fluctuations become smaller and smaller). Again, any point \mathbf{s}_2 can represent the area \mathcal{S}_2 , but we chose the TFT corner (orange dot in Fig. 11b).

A third area, \mathcal{S}_3 (yellow), corresponding to the mixed-strategy part of the AllC-TFT edge with exclusion of the TFT corner, also stands out: here we see the probability of AllD outcomes quickly increases with the fraction of AllC players. Due to the anti-clockwise dynamics, in \mathcal{S}_3 the system directed towards the point where the separatrix

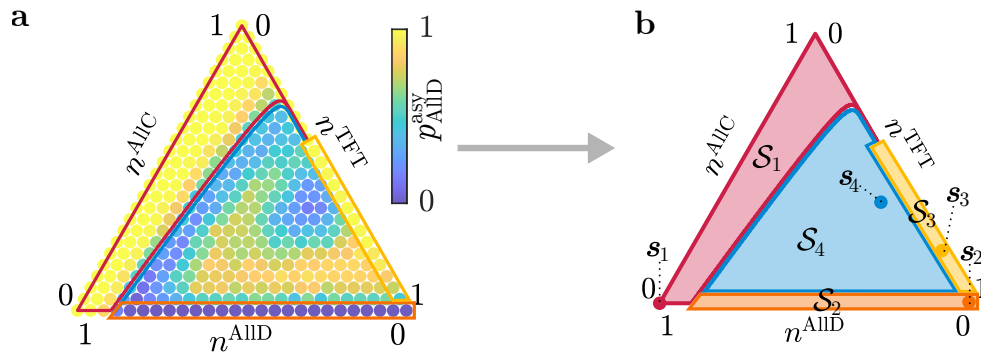


FIG. 11. **Equivalence classes of the asymmetric regime.** $p_{\text{AllD}}^{\text{asy}} = \Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\}$ was numerically estimated via Gillespi-Itô simulations for a subset of all points \mathbf{s} (panel a). The results motivates the decomposition of the simplex into four equivalence classes, \mathcal{S}_i , as well as the choice of their representative point \mathbf{s}_i (panel b). The parameters used for this illustration are $b - d = 0.05$ and $\log_{10} \mu = -6.5$.

meets the AllC-TFT edge, which is where a move from the mixed-strategy basin to the AllD basin is most likely (see Fig. 4f). Since the probability of AllD outcomes is less homogeneous, an approximation must be made for \mathbf{s}_3 . The dynamics along the AllC-TFT edges are slow, and they get even slower close to the corners. We observe that a good proxy for the average position of the system over time during a climb of the AllC-TFT edge from the TFT corner to the separatrix is around $n^{\text{AllC}} = 0.25$ and $n^{\text{TFT}} = 0.75$ (yellow dot in Fig. 11b).

Within the remainder of the simplex, \mathcal{S}_4 (blue), $\Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\}$ can be very heterogeneous, however, as discussed in the main text, this contributes very little to the outcome statistics for small values of N_0 . We chose $\mathbf{s}_4 = \text{MS-fp}$ to represent this area (blue dot in Fig. 11b).

Fig. 11a was obtained with $b - d = 0.05$ and $\log_{10} \mu = -6.5$. Different values of these parameters produce different values of $\Pr\{\mathbf{n}_{t_{\text{lock}}} \in \text{AllD-b} \mid \mathbf{n}_{t_{\text{asy}}} = \mathbf{s}\}$, however, the four areas can always be identified. This decomposition is validated by the successful reconstruction (see Figs. 7 & 8 as well as Figs 13-20) of the all statistics of the evolutionary outcomes obtained via full simulations (Fig. 3). The conditional probability decomposition is always very accurate except for the case of low mutation rate (e.g., $\log_{10} \mu < -6$), growth rates around $b - d \approx 0.05$ and large initial populations (e.g., $N_0 > 200$), for which it begins to produce less accurate results (Fig. 8 panel o).

10. Statistics of the stochastically-induced regime in detail

In this section we describe the effect of the mutation rate on the statistics of the stochastically-induced regime outcomes more in detail. Fig. 12a shows the same probabilities $p_{\mathcal{S}_i}^{\text{stoc}}$ shown in Figs. 7b-e, but for a single value of $b - d = 0.026$ (we remind that $\mathbf{n}_0 = \text{MS-fp}$ and $N_0 = 128$). We can see that all probabilities depend on μ for high mutation rates (i.e., $\log_{10} \mu > -4$), but only $p_{\mathcal{S}_1}^{\text{stoc}}$ and $p_{\mathcal{S}_3}^{\text{stoc}}$ depend on μ for lower mutation rates.

For a more detailed analysis, we decompose $p_{\mathcal{S}_1}^{\text{stoc}}$ into $p_{\mathcal{S}_1^{\text{AllD-c}}}^{\text{stoc}}$, the probability of $\mathbf{n}_{t_{\text{asy}}}$ being the AllD corner, $p_{\mathcal{S}_1^{\text{AllC-c}}}^{\text{stoc}}$, the probability of $\mathbf{n}_{t_{\text{asy}}}$ being the AllC corner, $p_{\mathcal{S}_1^{\text{AllC-TFT}}}^{\text{stoc}}$, the probability of $\mathbf{n}_{t_{\text{asy}}}$ being in the AllD basin side of the AllC-TFT edge excluding the AllC corner and $p_{\mathcal{S}_1^{\text{Other}}}^{\text{stoc}}$, the probability of $\mathbf{n}_{t_{\text{asy}}}$ being anywhere else within the AllD basin. We can now see that $p_{\mathcal{S}_1^{\text{AllC-TFT}}}^{\text{stoc}}$ decreases with μ , in contrast with $p_{\mathcal{S}_3}^{\text{stoc}}$, which instead increases with μ . This is explained by the shape of the separatrix changing with μ : the point where the separatrix meets the AllC-TFT edge moves towards the TFT corner as μ decreases (see panels c & d).

At the same time, we also see that the probability of the system being in one of the two corners at time t_{asy} , $p_{\mathcal{S}_1^{\text{AllD-c}}}^{\text{stoc}}$ and $p_{\mathcal{S}_1^{\text{AllC-c}}}^{\text{stoc}}$, are independent on the mutation rate for $\log_{10} \mu < -5$. However, for higher mutation rates it becomes more and more likely for the system to be in the AllD corner rather than in the AllC corner (see also panels c & d). Finally, for mutation rates $\log_{10} \mu > -5$ the probability of the system being away from the boundaries of the simplex $p_{\mathcal{S}_1^{\text{Other}}}^{\text{stoc}}$ becomes higher (the same of course can be observed for $p_{\mathcal{S}_3}^{\text{stoc}}$).

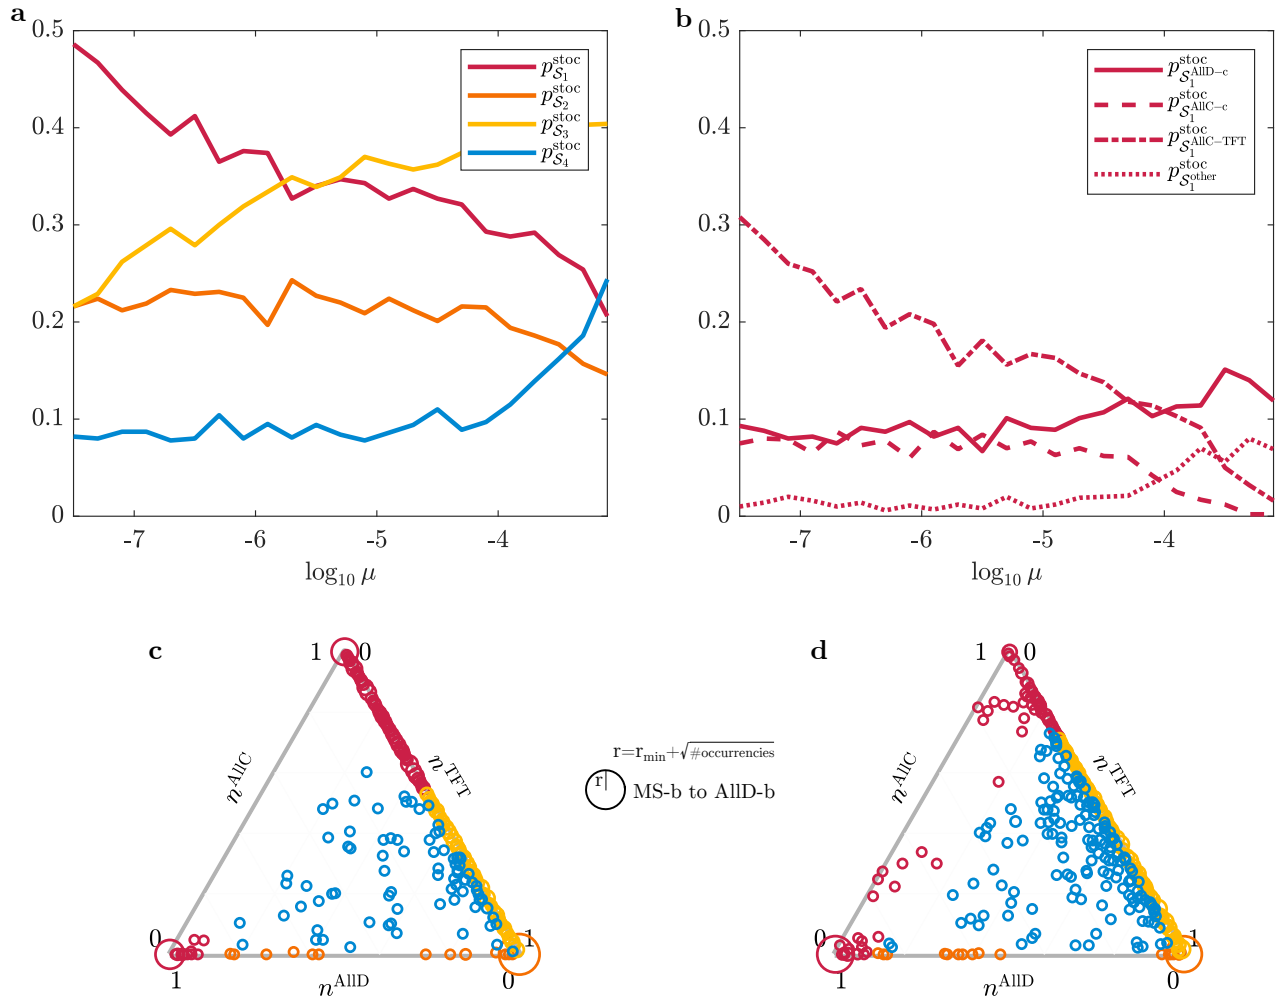


FIG. 12. **Statistics of the stochastically induced regime $p_{S_i}^{\text{stoc}}$ in detail.** Panel **a** shows the probabilities $p_{S_i}^{\text{stoc}}$ over μ with $b - d = 0.026$, $\mathbf{n}_0 = \text{MS-fp}$ and $N_0 = 128$. Panel **b** decomposes $p_{S_1}^{\text{stoc}}$ into cases of the AllD corner, the AllC corner, the AllD basin side of the AllC-TFT edge (excluding the AllC corner) and the rest of the AllD basin. Panels **c** and **d** show the location of $\mathbf{n}_{t_{\text{asy}}}$ within the simplex for the same 1000 stochastic trajectories used for estimating the probabilities in **a** and **b**, for the cases of $\log_{10} \mu = -6.5$ and $\log_{10} \mu = -3.5$ respectively. The size of the circles is proportional to the occurrences of $\mathbf{n}_{t_{\text{asy}}}$ in the location and the colours reflect the different equivalence classes S_i (see panel **a**).

11. Outcome statistics over growth/mutation rates: more initial mixes

Here we show that our decomposition of the outcome statistics based on equivalence classes holds not only for $\mathbf{n}_0 = \text{MS-fp}$ but in general. We consider other three cases that are very different: $\mathbf{n}_0 = \text{centre}$, i.e., $n^{\text{AllC}} = n^{\text{AllD}} = n^{\text{TFT}} = 1/3$ (Fig. 13); $\mathbf{n}_0 = \text{TFT-c}$, i.e., $n^{\text{AllC}} = n^{\text{AllD}} = 0$, $n^{\text{TFT}} = 1$ (Fig. 14); and $\mathbf{n}_0 = \text{AllD-c}$, i.e., $n^{\text{AllC}} = n^{\text{TFT}} = 0$, $n^{\text{AllD}} = 1$ (Fig. 15). In all figures, comparing the results obtained via decomposition (panels **n**) against those obtained via full simulation (panels **p**) yields a low error (panels **o**), demonstrates the validity of our method. The average error over all the considered values of growth and mutation is approximately 0.029 for the $\mathbf{n}_0 = \text{centre}$, 0.003 for $\mathbf{n}_0 = \text{TFT-c}$, and 0.001 for $\mathbf{n}_0 = \text{AllD-c}$.

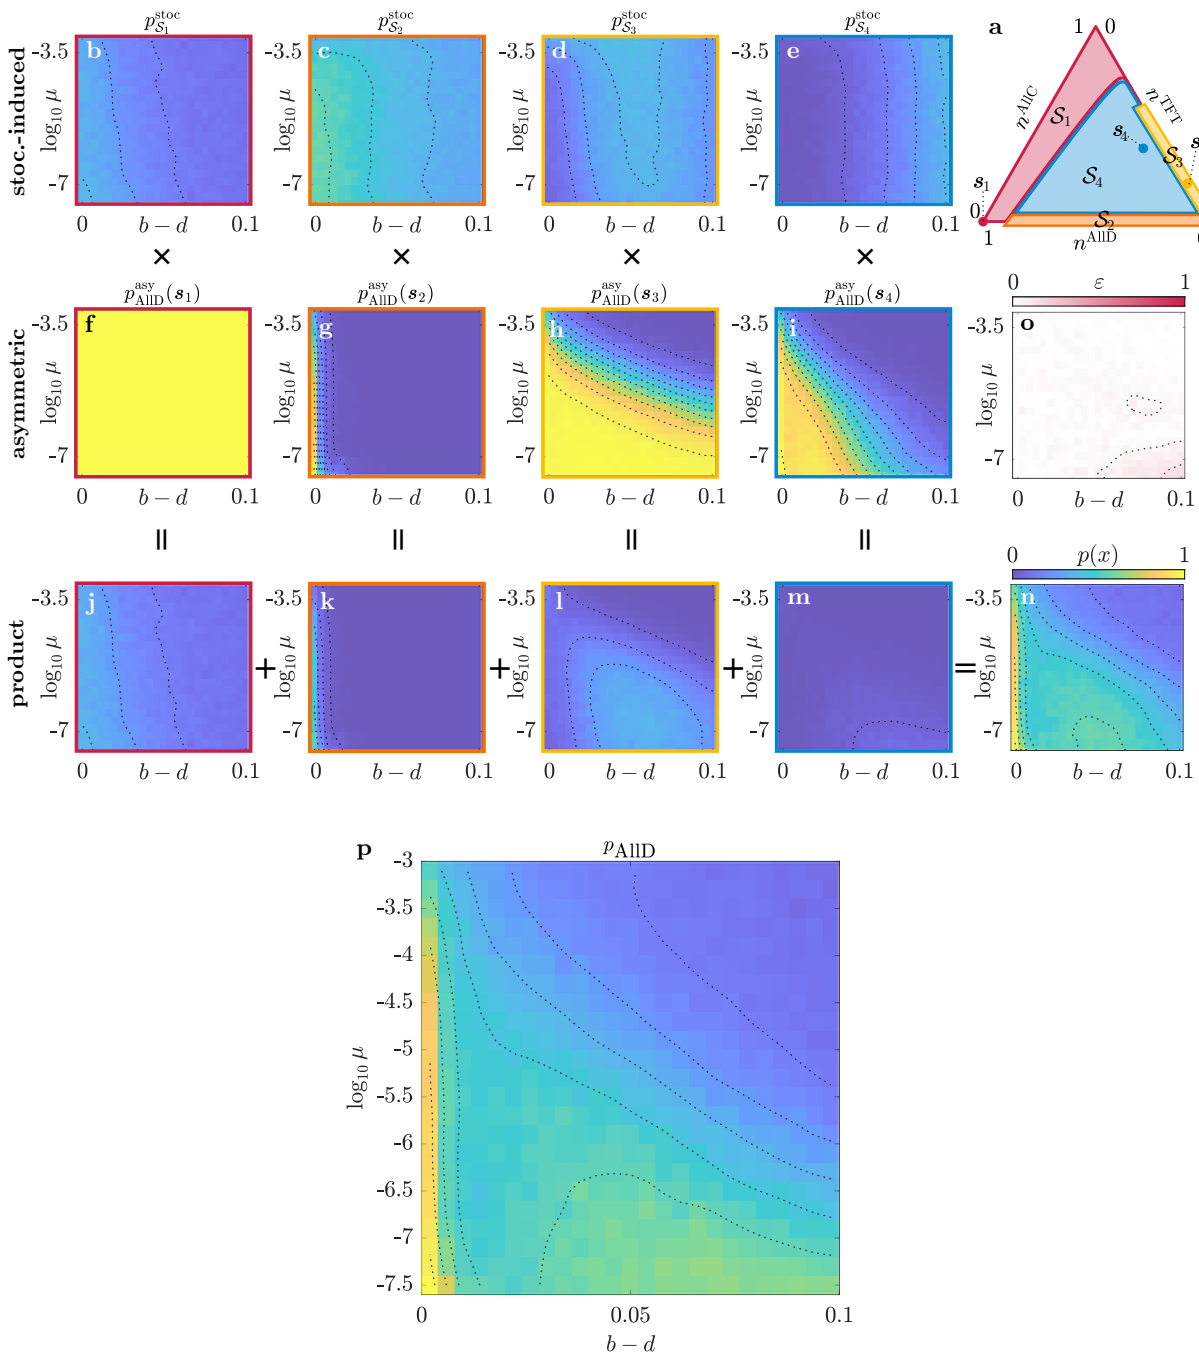


FIG. 13. **Decomposition of outcome statistics based on equivalence classes ($n_0 = \text{centre}$).** The statistics of evolutionary outcomes can be decomposed in terms of conditional probabilities that are based on four equivalence classes of states (panels a-i). The landscape of $p_{AIID}(\text{centre})$ for different rates of growth ($b-d$) and mutation (μ) can be reconstructed with good agreement (panels j-n and panel o). The initial population size is $N_0 = 128$. t_{asy} and t_{lock} are derived from the critical population sizes identified in Fig. 4.

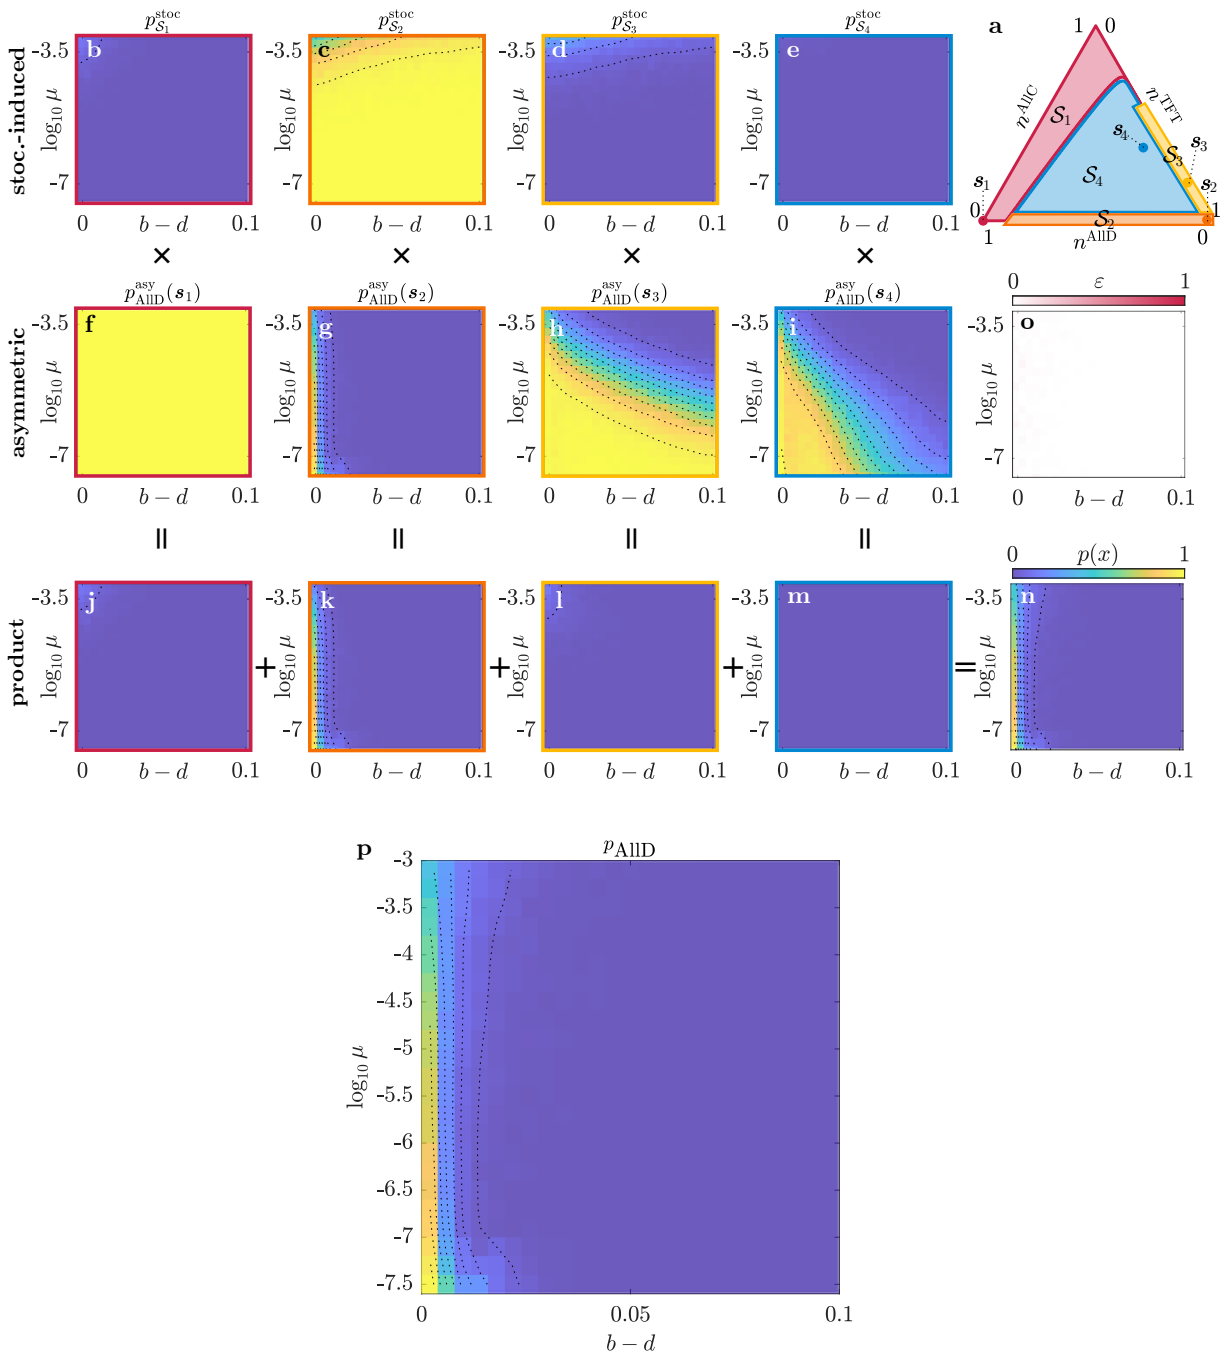


FIG. 14. **Decomposition of outcome statistics based on equivalence classes ($n_0 = \text{TFT-c}$).** The statistics of evolutionary outcomes can be decomposed in terms of conditional probabilities that are based on four equivalence classes of states (panels a-i). The landscape of $p_{\text{AllD}}(\text{TFT})$ for different rates of growth ($b-d$) and mutation (μ) can be reconstructed with good agreement (panels j-n and panel o). The initial population size is $N_0 = 128$. t_{asy} and t_{lock} are derived from the critical population sizes identified in Fig. 4.

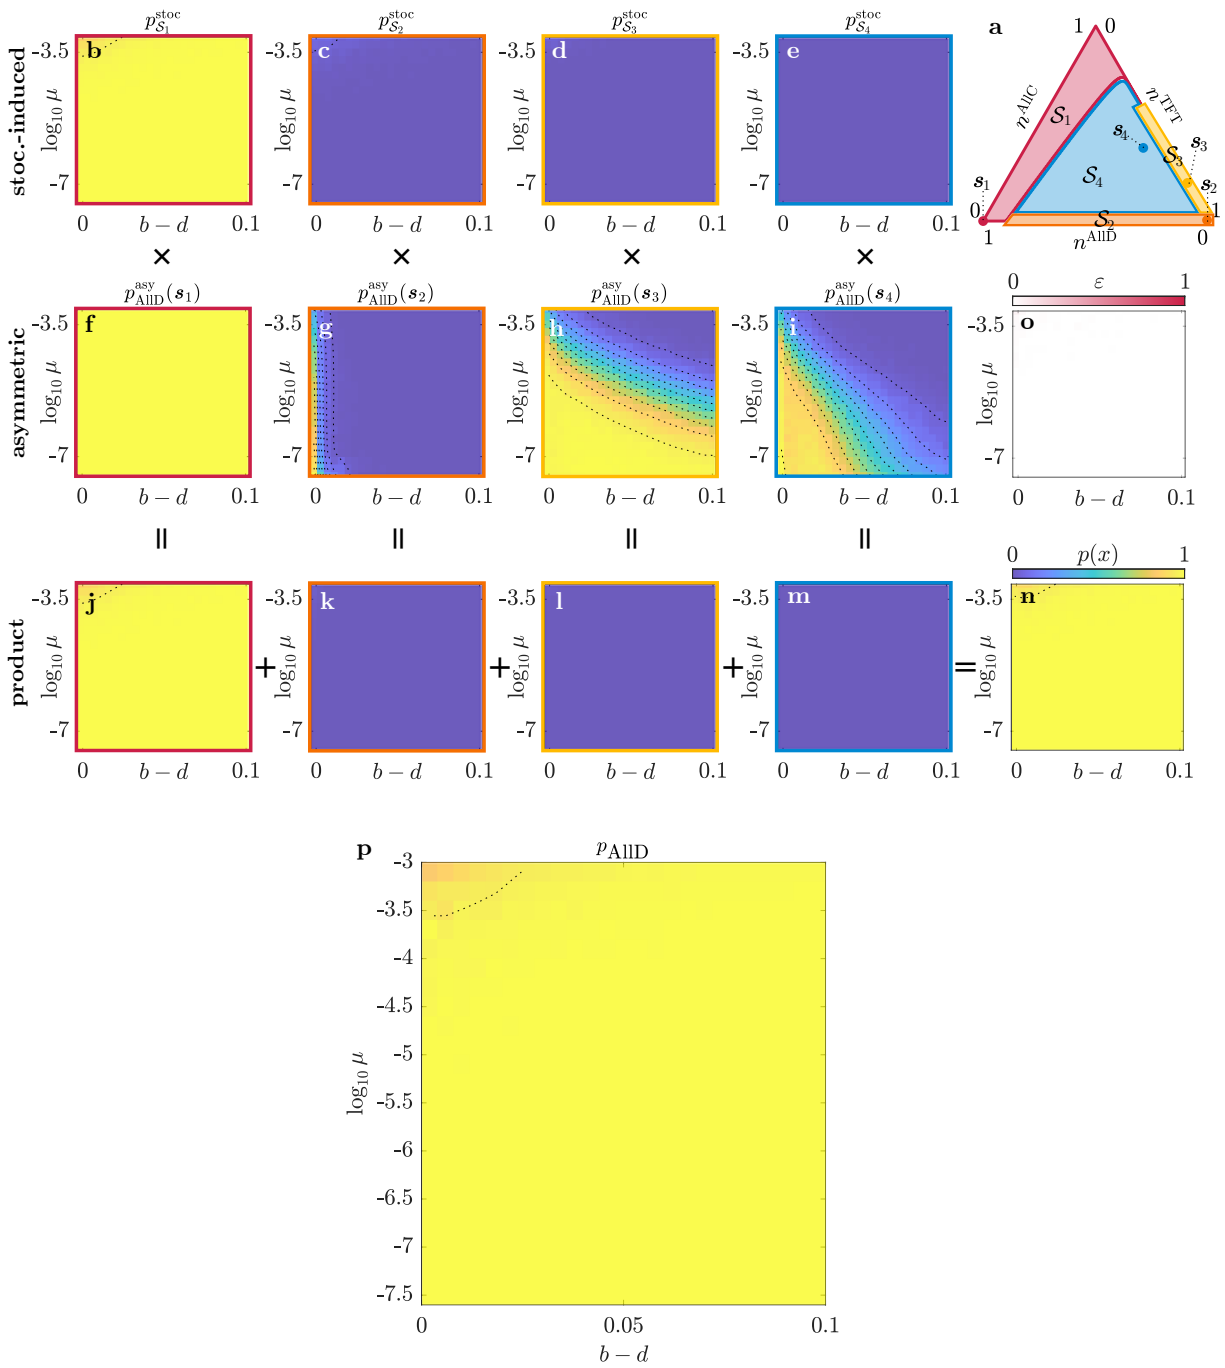


FIG. 15. **Decomposition of outcome statistics based on equivalence classes ($n_0 = \text{AID-c}$).** The statistics of evolutionary outcomes can be decomposed in terms of conditional probabilities that are based on four equivalence classes of states (panels a-i). The landscape of $p_{\text{AID}}(\text{AID-c})$ for different rates of growth ($b-d$) and mutation (μ) can be reconstructed with good agreement (panels j-n and panel o). The initial population size is $N_0 = 128$. t_{asy} and t_{lock} are derived from the critical population sizes identified in Fig. 4.

12. Outcome statistics over initial population's size and mix

In this section we illustrate how the outcome statistics in Fig. 3 panels **b**, **c**, **d**, **e** and **g** can be reconstructed using our probability decomposition method.

At high mutation rate (Figs. 3**b-c**), the asymmetric phase is very short. This is reflected in the probability decomposition in Figs. 16 and 17, where the probabilities in panels **f-h** are all zeros, i.e., the statistics of the outcome are given exclusively by the fixation of the system into one of the two basins during the stochastically induced phase (panels **a**). The accuracy of the probability decomposition is very high for both small (Fig. 16) and large (Fig. 17) initial populations, as illustrated in panels **j**.

At low mutation rate the asymmetric regime is very long, however, such regime has little effect on the outcome statistics for small initial populations. This is explained by the probability decomposition in Fig. 18. For some starting points N_0 the system can be in the equivalence class \mathcal{S}_2 at the end of the stochastically-induced regime (panel **b**), however, from this area of the simplex it is extremely unlikely for the system to cross the separatrix towards the AllD basin, even if the asymmetric regime is long (panel **f**). Moreover, it is very unlikely for the system to be in \mathcal{S}_4 at the end of the stochastically induced regime, since with a small initial population such regime is longer and fluctuations are likely to drive the systems towards the edges and corners (panels **d**). Thus, even if the system can cross the separatrix from \mathcal{S}_4 during the asymmetric regime, the combined probability in panel **h** is approximately zero for every initial condition. Finally, we can see that for some initial state s the system has a small chance of being in \mathcal{S}_3 at the end of the stochastically-induced regime (panel **c**). From this area the system is very likely to cross the separatrix during the asymmetric regime, leading to the small outcome probability contribution in panel **g**.

Since the probabilities in panel **h** are very low, the accuracy of the probability decomposition is very high (panel **j**) also in this case. The accuracy gets worse only for middle values of growth rate (around $b - d = 0.05$), very small mutation rates (e.g., $\log_{10} \mu < -6$) and large initial populations (e.g., $N_0 > 200$). This case is reported in Fig. 8, which shows that the probability of AllD outcomes is underestimated for some initial starts around the mixed-strategy unstable fixed point.

Fig. 3 panels **d** and **g** can be reconstructed in a similar way, as shown in Figs. 19 and 20.

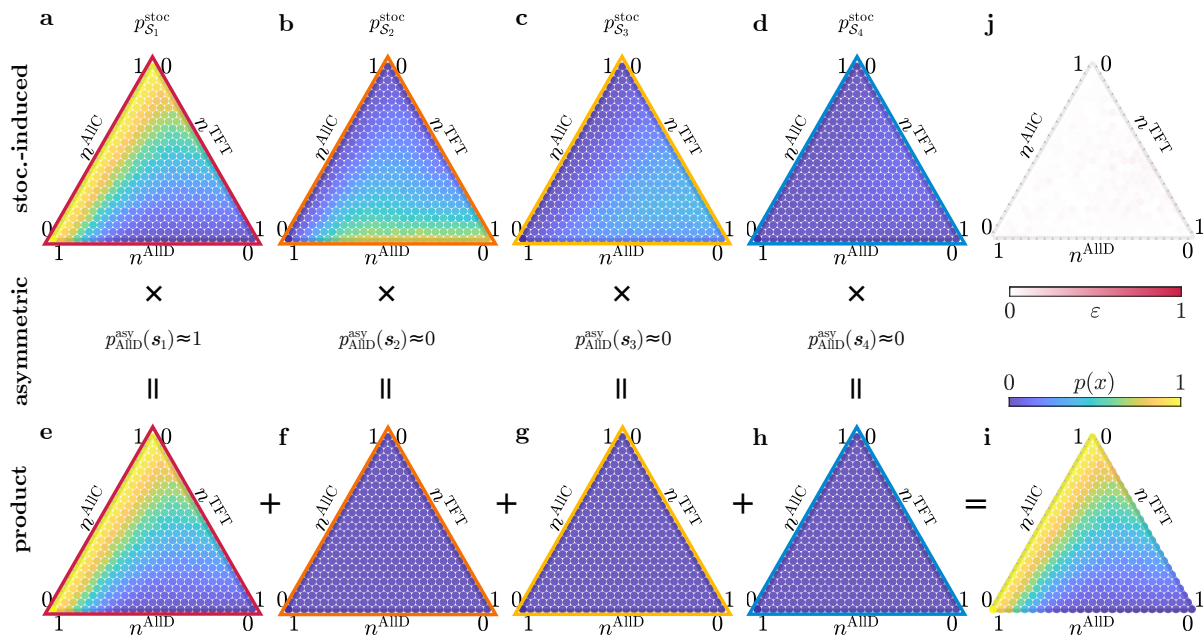


FIG. 16. **Outcome statistics decomposition over n_0 : $\log_{10} \mu = -3.5$, $N_0 = 50$ and $b - d = 0.05$.** Panels **a-d** show the probability of the system being in the four equivalence classes \mathcal{S}_i at the end of the stochastically-induced regime. The probabilities of AllD outcomes given that the system is in the representative points s_i at time beginning of the asymmetric regime are reported in the second row. Panels **e-i** show how the probabilities are combined to approximate the outcome probability and panel **j** shows accuracy of such approximation (cf. Fig. 3**b**).

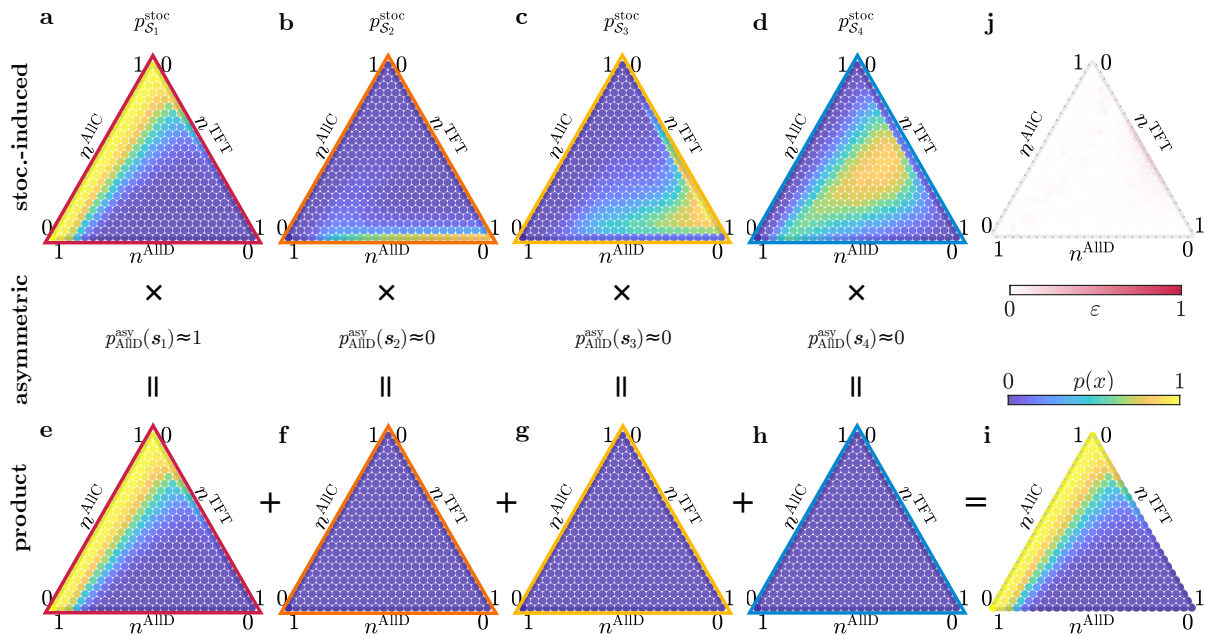


FIG. 17. **Outcome statistics decomposition over n_0 : $\log_{10}\mu = -3.5$, $N_0 = 250$ and $b - d = 0.05$.** Panels a-d show the probability of the system being in the four equivalence classes S_i at the end of the stochastically-induced regime. The probabilities of AllD outcomes given that the system is in the representative points s_i at time beginning of the asymmetric regime are reported in the second row. Panels e-i show how the probabilities are combined to approximate the outcome probability and panel j shows accuracy of such approximation (cf. Fig. 3c).

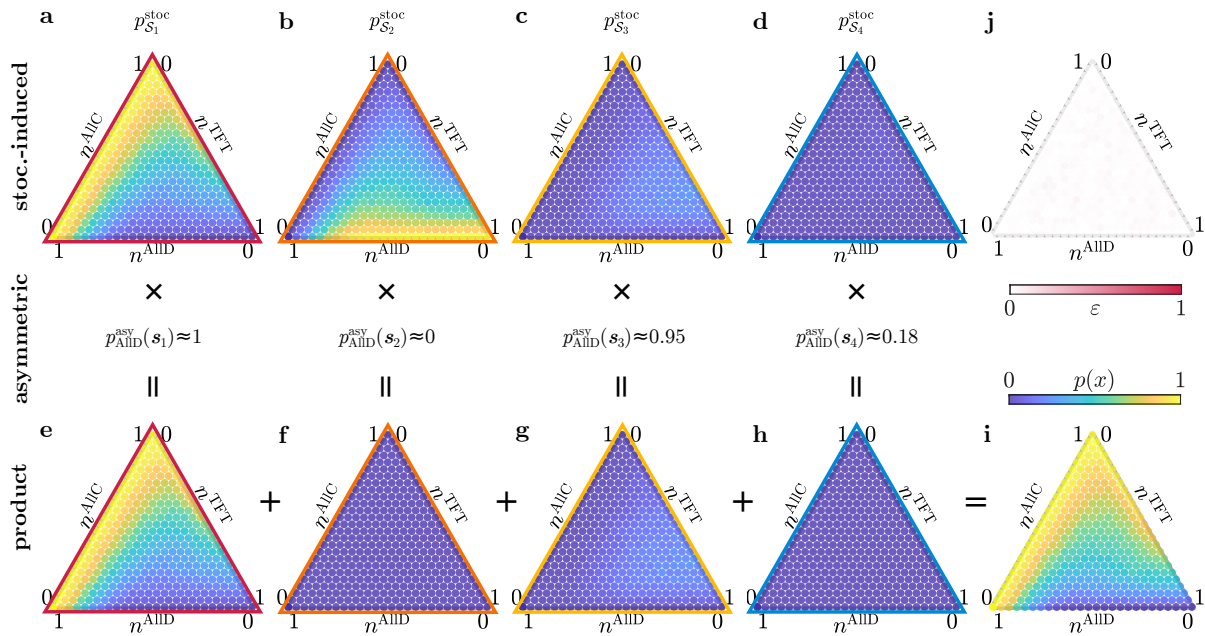


FIG. 18. **Outcome statistics decomposition over n_0 : $\log_{10}\mu = -6.5$, $N_0 = 50$ and $b - d = 0.05$.** Panels a-d show the probability of the system being in the four equivalence classes S_i at the end of the stochastically-induced regime. The probabilities of AllD outcomes given that the system is in the representative points s_i at time beginning of the asymmetric regime are reported in the second row. Panels e-i show how the probabilities are combined to approximate the outcome probability and panel j shows accuracy of such approximation (cf. Fig. 3e).

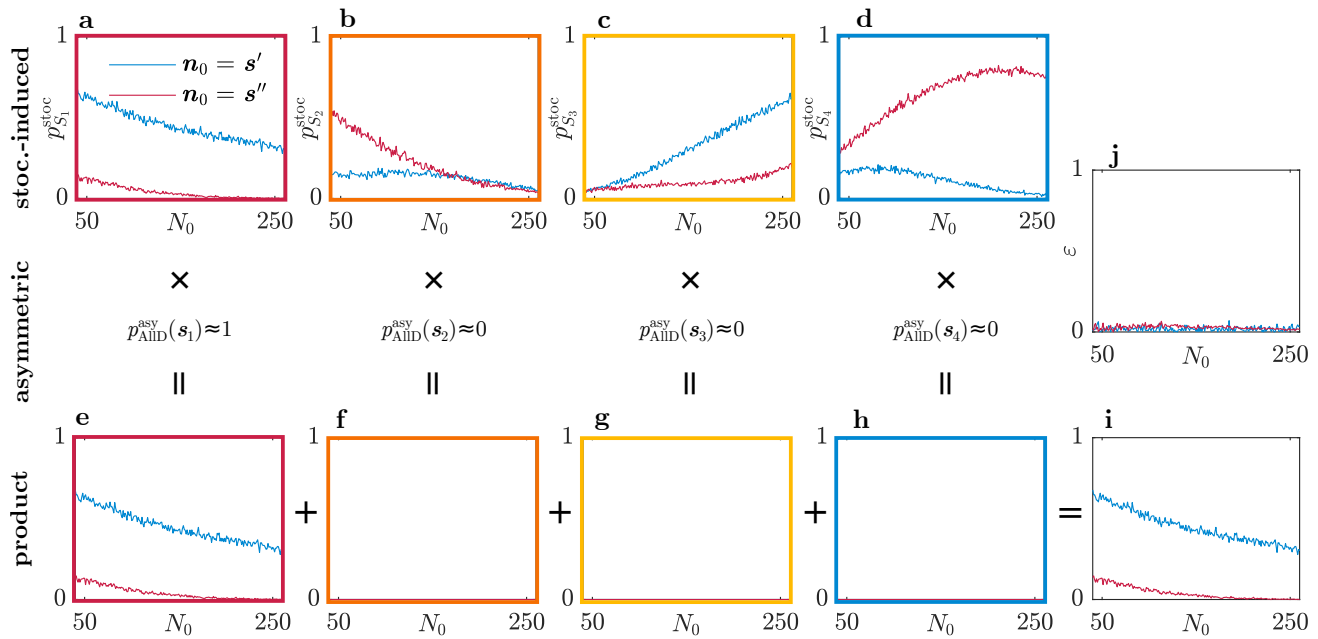


FIG. 19. **Outcome statistics decomposition over N_0 : $\log_{10}\mu = -3.5$ and $b - d = 0.05$.** The results are shown for the two initial states s' and s'' in Fig. 3. Panels a-d show the probability of the system being in the four equivalence classes S_i at the end of the stochastically-induced regime. The probabilities of AllD outcomes given that the system is in the representative points s_i at time beginning of the asymmetric regime are reported in the second row. Panels e-i show how the probabilities are combined to approximate the outcome probability and panel j shows accuracy of such approximation (cf. Fig. 3d).

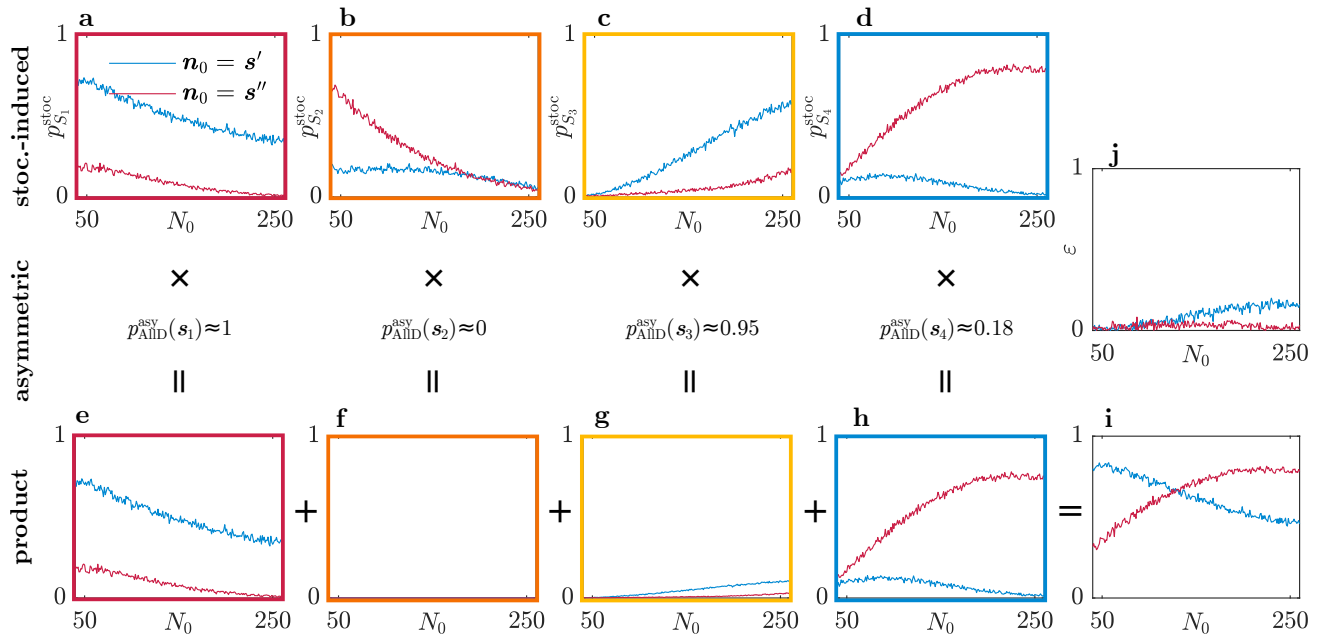


FIG. 20. **Outcome statistics decomposition over N_0 : $\log_{10}\mu = -6.5$ and $b - d = 0.05$.** The results are shown for the two initial states s' and s'' in Fig. 3. Panels a-d show the probability of the system being in the four equivalence classes S_i at the end of the stochastically-induced regime. The probabilities of AllD outcomes given that the system is in the representative points s_i at time beginning of the asymmetric regime are reported in the second row. Panels e-i show how the probabilities are combined to approximate the outcome probability and panel j shows accuracy of such approximation (cf. Fig. 3g).