The inherent fragility of collective proliferative control

Michael G. Caldwell[1] and Arthur D. Lander[1,2,3,*]

[1]Center for Complex Biological Systems, [2]Department of Developmental and Cell Biology, and [3]Department of Biomedical Engineering, University of California, Irvine, CA 92697-2300

*corresponding author; adlander@uci.edu

ABSTRACT

Tissues achieve and maintain their sizes through active feedback, whereby cells collectively regulate proliferation and differentiation so as to facilitate homeostasis and the ability to respond to disturbances. One of the best understood feedback mechanisms—*renewal control*—achieves remarkable feats of robustness in determining and maintaining desired sizes. Yet in a variety of biologically relevant situations, we show that stochastic effects should cause rare but catastrophic failures of renewal control. We define the circumstances under which this occurs and raise the possibility such events account for important non-genetic steps in the development of cancer. We further suggest that the spontaneous stochastic reversal of these events could explain cases of cancer normalization or dormancy following treatment. Indeed, we show that the kinetics of post-treatment recurrence for many cancers are often better fit by a model of stochastic re-emergence due to loss of collective proliferative control, than by deterministic models of cancer relapse.

INTRODUCTION

Animal species depend for survival on the ability to build and maintain tissues that are stable (relatively unchanging in time) and robust (near in size and morphology to a desired set-point). Among other things, achieving these goals requires a mechanism for proliferative homeostasis, a process typically carried out in the context of cell lineages—i.e., the production of differentiated or "terminal" cells by stem cells or stem cell-derived intermediate progenitors. Although it is widely understood that homeostasis must, at some level, amount to a process of matching the production of terminal cells to their turnover, surprisingly little is known about how this is accomplished across the multitude of tissue types.

In recent years, mathematical models have helped elucidate fundamental mechanisms underlying stem cell behavior and tissue growth[1-7]. Here we exploit such models to identify tradeoffs that tissues may encounter in attempting to achieve homeostasis. In particular, we investigate the possibility that control mechanisms that evolved to ensure tight homeostasis are also a source of fragility, so that, on vary rare occasions, one may expect well-controlled tissues to spontaneously start growing without bound.

To the extent such a fragility exists—we will argue that it is likely unavoidable—understanding how tissues minimize it can shed light on the constraints under which normal tissues operate. Perhaps even more importantly, understanding how such a fragility arises can suggest new ways to think about cancer—a situation defined by the absence of controlled growth. Cancer initiation is often portrayed of as the result of genetic changes that "drive" growth or "inhibit" cell death, but if homeostasis is controlled through feedback (and we argue here that it must be), one would not generally expect quantitative adjustments to such drives to produce a total collapse of control. In contrast, if total collapse can occur on its own at random, then an alternative way to think about oncogenes and tumor suppressor genes is as factors that make such rare transitions more likely.

Here we use simulation and mathematical analysis to show that a relative common form of homeostatic feedback control necessarily creates opportunities for rare, stochastic escape, particularly when the spatial constraints of cells and tissues are taken into consideration. At the heart of this behavior lies a stochastic transition that is collective, rather than cell-autonomous.

Although we do not currently know whether the development of some cancers depends upon the existence of this transition, we argue that its inherent reversibility could explain how some treated cancers become dormant, re-emerging stochastically at a much later time. Using clinical data on the timing of cancer recurrence in patients treated for a variety of different cancers, we show that recurrence kinetics are, indeed, often better fit by a model based on stochastic re-emergence than by traditional, deterministically motivated models.

RESULTS

Producing or maintaining reliable tissue size requires feedback control

Maintaining self-renewing tissues, such as blood, skin and intestinal epithelia, at fixed sizes requires stem cells to replace differentiated, or "terminal", cells at a rate matching cell turnover. It is widely recognized that, to replace terminal cells at a constant rate, exactly half the offspring of stem cell divisions must remain stem cells—any more and a tissue will grow in size indefinitely; any fewer and it will shrink and disappear.

Less appreciated is the fact that, even if every stem cell did renew with probability exactly equal to one half, homeostasis still could not be effectively maintained. The reason is that, for most types of stem cells, division outcomes are probabilistic, meaning that sometimes cell division produces two stem cells (symmetric renewal), sometimes two differentiated cells (symmetric extinction), and sometimes one of each (asymmetric division). Even if renewal exactly balances extinction, fluctuations induced by the probabilistic nature of these outcomes will induce, in any finite-sized population, a variance in population size that grows indefinitely over time. This is illustrated in Figure 1A, which simulates scenarios that all start with the same number of stem and terminal cells, and in which each stem cell offspring differentiates independently and probabilistically, with

probability of exactly one half. Although the average cell number across simulations remains roughly constant, the variance grows without bound (also see Appendix).

Of course, this behavior can be avoided if all stem cell divisions are asymmetric, but the fact that few stem cell systems undergo strictly asymmetric division implies that variance must get stabilized by some other means. That in turn implies that stem cells cannot behave independently of each other, a conclusion also supported by direct observation of fluctuations in systems sustained by small numbers of stem cellse.g. [8,9]. Specifically, stem cells must adjust their renewal behavior in response to changes in tissue size, or some proxy for tissue size, such as terminal cell number, degree of cell crowding, net rate of growth, availability of a limiting nutrient, etc. In other words, *stem cells must be feedback-controlled*.

One of the simplest and most efficient ways to achieve such control is for stem cell renewal probability to fall (i.e., for differentiation probability to rise) with the number of terminal cells, a feedback strategy that has been termed "renewal control"[5]. Renewal control has been documented in several tissues including muscle, sensory epithelia, and blood, wherein the concentrations of secreted factors (often members of the TGFβ superfamily) serve as a proxy for differentiated cell number[2,4,10-12]. Renewal control not only suppresses fluctuations (Fig. 1B), it displays a property known as "perfect robustness" or "robust perfect adaptation"[13], whereby the steady state number of terminal cells becomes independent of all parameters that lie outside of the feedback loop itself—for example the rate of the cell cycle, the rate at which terminal cells die, or the initial number of cells that the system begins with[4,14-16]. Robust perfect adaptation arises because renewal control is an example of integral negative feedback, a strategy widely used in both biology and engineering to maintain systems at

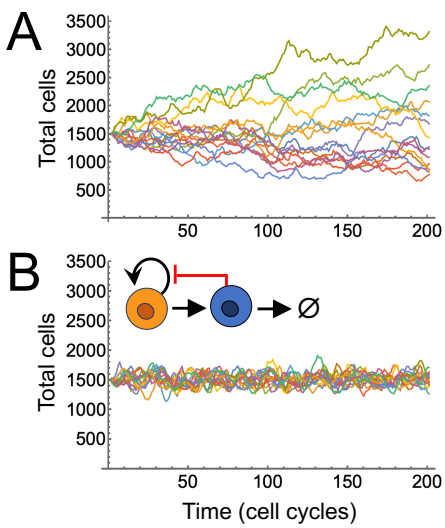

Figure 1. Monte Carlo simulations of stem cell systems **A.** Starting from an initial state of 500 stem and 1000 terminal cells, stem cells divide every cell cycle, each differentiating with probability ½, and terminal cells die at a rate of 0.5 per cell cycle. Each trace is an independent simulation. **B.** Stem cells renew probabilistically, but the probability of renewal is a declining function of the number of terminal cells, with feedback strength adjusted to produce an average terminal cell number of 1000. Initial conditions and terminal cell turnover are as in panel A.

fixed set-points in the face of unpredictable disturbances [17].

Renewal control provides a simple explanation for the extraordinary capacity of self-renewing tissues to grow and regenerate to genetically specified sizes without fine-tuning of initial conditions or parameters. Renewal control can also explain the robust development of non-renewing tissues—e.g. brain, retina, cartilage, etc.—which can be modeled as instances of renewal control in which terminal cells do not turn over: As terminal cells accumulate, net stem cell renewal eventually falls below 50%, leading to stem cell exhaustion, and a static "final state" (as opposed to a steady state) where only non-dividing cells remain. As long as such a final state is much larger than the initial state, it can be shown that the final state is nearly perfectly robust to initial conditions, as well as to other parameters lying outside the feedback loop[5,14].

Figure 2 illustrates the flexibility and robustness of renewal control in both steady- and final state contexts. Panels A and E diagram feedback circuits in which terminal cells (blue) that either do (A-D) or don't (E-H) turn over inhibit the renewal of stem cells (orange). B and F show the solutions to (deterministic) differential equations, for a single set of parameters, in which cells are represented as continuous concentrations, and feedback follows a declining Hill function (i.e., a function that, with enough terminal cells, would drive stem cell renewal to zero); notice how steady and final states are reached, respectively, in these two cases. In panels C and G, solutions to the same equations are presented using a phase diagram, with stem and total cell numbers on the two axes, and time represented by arrows on streamlines. From these panels one can see how solutions that start from different initial conditions evolve to converge on a single fixed point (panel C) or small region on the ordinate axis (panel G). Panels D and H show the outcomes of stochastic simulations similar to those in Fig. 1—cases in which cells were treated as discrete units undergoing probabilistic divisions—and the results of 500 such simulations are plotted using phase diagrams (as in C and G). Although the stochastic and deterministic trajectories are attracted to similar locations on the phase diagram, it is clear that stochasticity allows for oscillations about the deterministic steady state, a phenomenon also apparent in Fig. 1B.

Stability and instability in the feedback control of growth

Any biologically useful strategy for controlling tissue size must produce a stable steady state (or, when needed, a final state). The ideal strategy should be globally stable, i.e., stable regardless of parameter choices, since that would ensure that disturbances to parameters (e.g., mutations that affect gene expression) could not easily lead to uncontrolled growth. The equations that govern the behaviors depicted in Fig. 2A-H are, in fact, globally stable. However, as with most simple models, the equations rely on making assumptions. For example, representing feedback with a declining Hill function is arbitrary. As it turns out, global stability would still be guaranteed if feedback were represented by any monotonically declining function that preserves the potential to drive the renewal probability from above ½ to below ½. A more serious issue is the assumption that the behaviors of stem cells (or the probabilities associated with those behaviors) depend only on time, and not space, i.e., stem cells all receive the same feedback regardless of location. In reality—in solid tissues at least—stem cells at different locations would likely be exposed

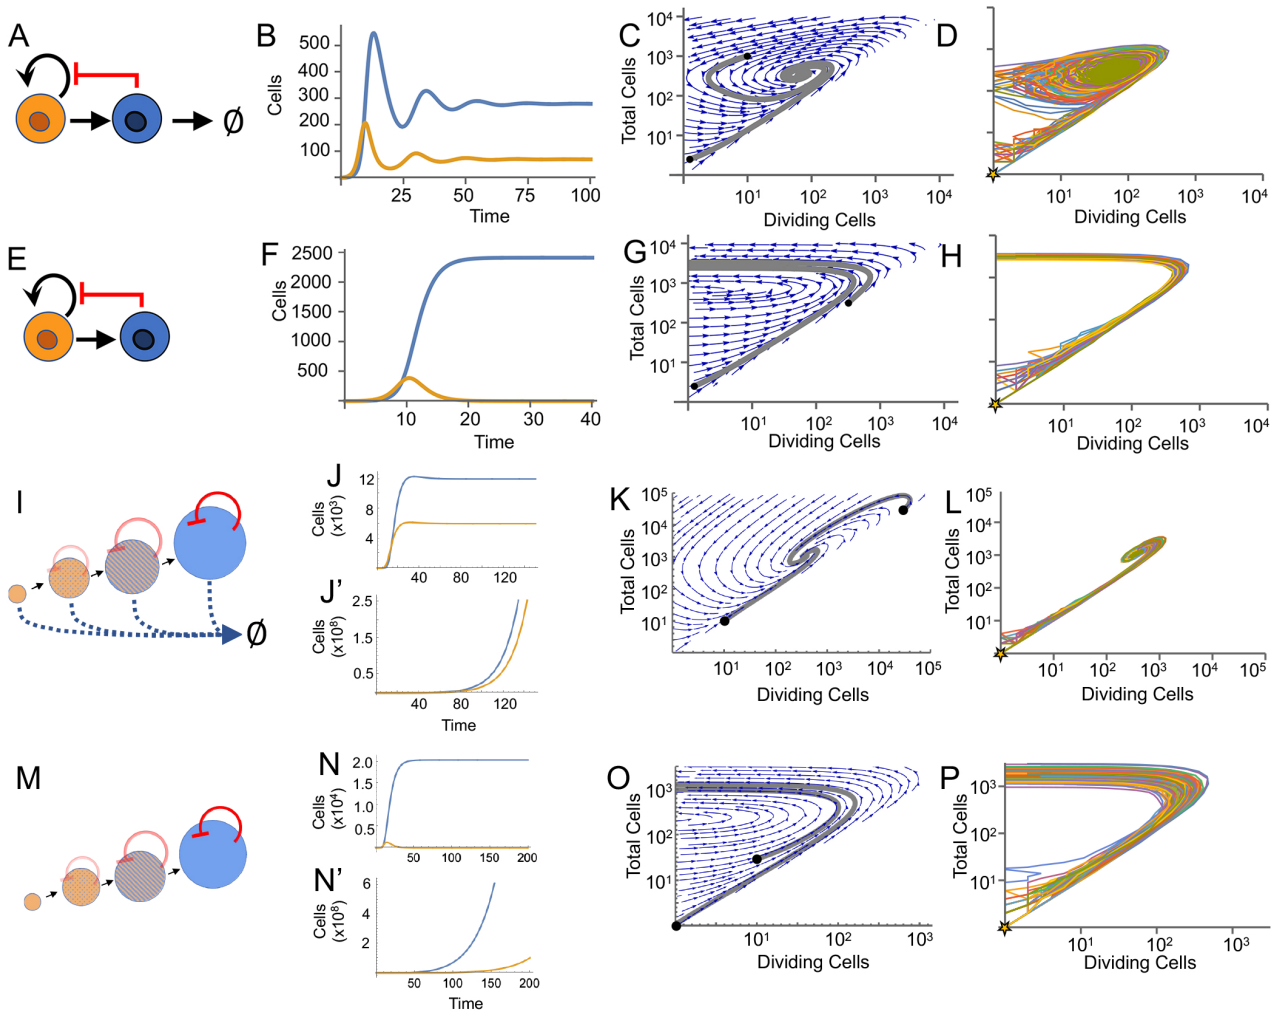

Figure 2. Behaviors enabled by renewal control. A-D. Feedback inhibition of stem cell renewal by long-lived terminal cells produces a robust steady state. Ordinary differential equations that model the feedback circuit in A were solved for a single set of parameters in B-C. In A-B, orange denotes stem cells and blue terminal cells. Panel C plots solutions in the phase plane. In D, 500 Monte Carlo simulations were performed in which differentiation of stem cell offspring was probabilistic. **E-H.** Feedback inhibition of stem cell renewal by terminal cells that do not turn over produces a robust final state. Panels were generated as in A-D. **I-P.** Simplified models of renewal feedback in solid tissues, taking into account declining sensitivity to feedback as tissues grow beyond the characteristic decay length of a diffusing feedback signal. I-L represent the steady-state case; M-P the final-state case (no terminal cell turnover). Alternative panels J-J' and N-N' display stable (J, N) and unstable (J'-N') behaviors for different parameter sets. Phase plane results are shown in K and O and stochastic simulations in L and P. For parameter values, see Table 1.

to different levels of feedback signals. For example, if such signals are carried by diffusible molecules, concentration gradients should arise.

The stability behavior of such systems is better modeled using partial differential equations, but their analytical treatment is complicated[18]. To build intuition more easily, we considered an intermediate formulation that still uses ordinary differential equations (ODE) and accounts for spatial effects in a simplified manner. Specifically, we modeled a solid tissue as a growing disc in two-dimensions, the area of which is determined by the number of total (stem and terminal) cells. We assume stem and terminal cells are uniformly mixed ("well stirred") within the disc (we relax this assumption later). At every location, we consider that a negative feedback signal is produced in proportion to the local concentration of terminal cells, and that the carrier of this signal diffuses away, and is taken up (or destroyed) at a constant rate. Assuming diffusion and uptake are fast relative to tissue growth, we may convert the spatial pattern of production of the feedback signal into a steady-state concentration gradient across the disc. As described elsewhere[19], that shape may be calculated exactly, and has the expected property of being highest in the center of the disc and lowest at the edge. Accordingly, the amount of feedback "felt" by stem cells will differ by location. To obtain an expression for the aggregate behavior of the entire stem cell pool, one integrates the renewal probability across the disc, deriving an average renewal probability. The time-evolution of the total numbers of stem and terminal cells from any initial condition may then be formulated as an ODE problem.

Figure 2I-P explores the behavior of such systems for cases in which terminal cells do (Fig. 2I-L) and don't (Fig. 2M-P) turn over. In both cases they are no longer globally stable (see

Appendix for proof). When feedback is sufficiently strong, steady or final states are reached (Fig. 2J and 2N), but with weaker feedback, unbounded growth occurs (Fig. 2J' and 2N'). This is because, as such systems expand, they experience "declining gain feedback", where the total feedback produced by the system does not rise fast enough to keep pace with system growth. This happens because gradients created by diffusing substances produced within a spatial domain always approach a saturating shape as domain size becomes large relative to the intrinsic decay length set by the parameters of diffusivity and decay[19].

On the other hand, such systems are still "locally stable", meaning that, if one chooses parameters appropriately, they behave like the non-spatial models, exhibiting robust perfect adaptation of steady states and near-perfect adaptation of final states (Fig. 2K, O, respectively). Importantly, stochastic trajectories in these regimes closely mirror deterministic ones.

Stochastic instability can arise when cell types do not thoroughly mix

Of course, solid tissues do not grow as perfect discs or spheres, nor do cells normally intermingle perfectly. To investigate the consequences of a more realistic accounting of spatial details, we turned to agent-based modeling. In such models each cell and its behavior are represented explicitly, so stochastic effects are automatically captured. Fig. 3A-B shows the rules of an agent-based model in which cells divide on a grid, pushing their neighbors aside as they do; terminal cells arise but do not turn over; and the probability that stem cell offspring renew (as opposed to differentiate) decreases according to a diffusion gradient created by the release of a feedback signal from each differentiated cell (see methods for further details).

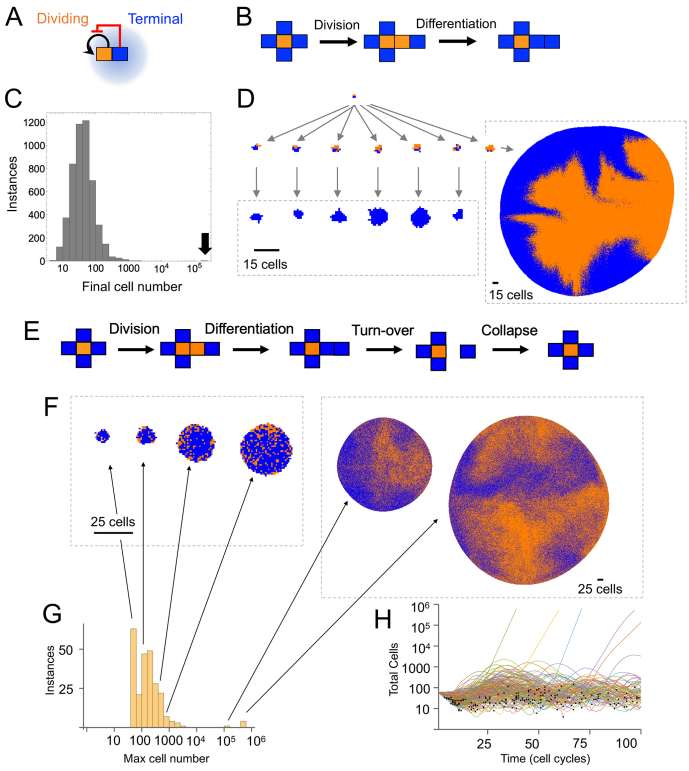

Such a model thus seeks to capture the "final state" scenario of Fig. 2M-P.

Simulations were initialized with a square of 4 cells, two dividing cells atop two differentiated ones, and run until either all divisions stopped, or the number of stem cells exceeded 70,000. Fig. 3C summarizes the results of 5,000 simulations, and shows that the vast majority ran out of dividing cells, stopping with a median size of 39 cells. They formed a roughly log-normal distribution of final sizes, with 99% of cases stopping at fewer than 420 cells. These simulations qualitatively mirrored the behavior in Fig. 2P, although the final size for these parameter choices was considerably smaller.

Three cases, however, continued to grow on what appeared to be an exponential trajectory. Fig. 3D contrasts the behaviors of 6 typical cases with one of the three that failed to stop growing, displaying the arrangements of cells after 4 cell cycles and, below that, at the end of the simulation.

We also modeled the scenario in which terminal cells continuously turn over so that, for sufficiently strong feedback, steady state behavior might be expected, as in Fig. 2I-L. This required modifying the agent-based simulation to allow cells to move towards each other so as to collapse empty spaces left by cell loss (Fig. 3E). Parameters were found for which simulations reached and oscillated about apparent steady states for long periods of time (e.g., 100 cell cycles), and 250 new simulations were initialized from those approximate steady state conditions. Most cases continued to oscillate about steady state, with a maximum total cell number in range of a few hundred (Fig. 3F-H). In a few cases, however, simulations switched at seemingly random times to what appeared to be unbounded exponential growth (Fig. 3H), resembling that in Fig. 3D.

Figure 3. Spontaneous stochastic instability. Agent-based models were used to simulate cells growing and differentiating in a two-dimensional space, influenced by renewal feedback, either without (A-D) or with (E-H) turnover of terminal cells. **A.** Cells were initialized on a fixed grid, and diffusion gradients from each differentiated cell calculated and summed to produce a feedback field, from which probabilities of renewal *vs.* differentiation at each position were calculated. **B.** Starting from initial conditions of two dividing cells juxtaposed with two differentiated cells, and neglecting death or removal of differentiated cells, simulations were run until either no dividing cells remained, or the number of dividing cells exceeded a pre-set maximum. **C.** Distribution of final sizes produced by 5000 such simulations. The arrow indicates rare simulations that were stopped because dividing cell numbers exceeded 70,000. **D**. Images from six representative simulations that had ceased growing, and one of three that had not. The initial condition for all simulations is shown at top, the results after 4 rounds of cell division in the middle, and those at the end of simulation at bottom. The image at right (note smaller scale bar) reached 92,998 dividing and 70,768 differentiated cells after 25 cell divisions. **E.** Starting from initial conditions close to the steady state for the chosen parameters (see Methods), simulations were run until either no dividing cells remained, the number of dividing cells exceeded 500,000, or 100 cell cycles had elapsed. **F.** Six representative simulations at their largest sizes. The four on the left continually oscillated about a steady state while the two on the right appeared to be increase approximately exponentially, and without bound. **G.** Distribution of the largest sizes (total cell number) observed among 250 simulations. The arrows show the bins into which the images in panel F fall (note logarithmic axis). **H.** Cell numbers over time for the 250 simulations summarized in panel G. Cases in which all stem cells went extinct are indicated by black dots at the time of complete stem cell disappearance. At least five cases appear to display unbounded growth (note logarithmic axis).

Such spontaneous escape from growth control was never seen in simulations of any of the models in Figure 2. Examination of the patterns of cell arrangement in cases of escape showed that dividing and differentiated cells were no longer thoroughly mixed but had spontaneously broken up into more homogeneous patches (Fig. 3D, F). Suspecting that insufficient cell mixing might be responsible for driving rare escape from growth control, we modified the ODE models of Fig. 2 to eliminate cell mixing, by mandating that stem and terminal cells always and instantaneously sort away from each other (this allows the model to still be posed in terms of ODEs). Fig. 4A-F shows the results when stem cells were made to sort to the outside of growing discs; in Fig. 4G-L, stem cells all moved to the interior.

When the equations describing these scenarios were analyzed using phase diagrams (Fig. 4B,E,H,K), bimodality was noticed: For a single set of parameters, the same system could be attracted either to a stable state or undergo unbounded growth, depending on the initial conditions (depicted as black dots on the phase diagrams). In general, conditions that grew without bound were those that started with an already large number of stem cells.

Whereas initial conditions fully determined the outcomes of

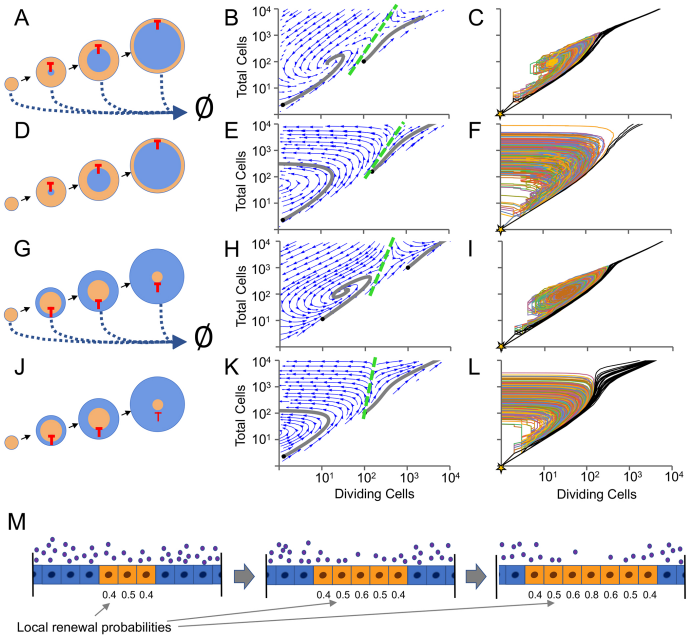

Figure 4. An absence of cell mixing predisposes to stochastic instability. **A,D,G,J**. Each row represents a system in which dividing cells (orange) grow as discs in space, and generate differentiated cells (blue), under the influence of diffusible negative feedback, modeled as in Fig. 2 panels I and M. In A and D, dividing cells instantaneously sort to the outsides of discs; in G and J they sort to the insides. In A and G, differentiated cells turn over; in D and J they do not. **B,E,H,K**. Phase portraits calculated for each of the four systems; note the existence of two basins of attraction in each case (separatrices are shown as dashed green lines). **C,F,I,L.** 500 stochastic trajectories of each system, starting from identical initial conditions (a single dividing cell), plotted in phase space. Trajectories that spontaneously escape from growth control are shown in black. **M.** Diagram explaining the source of cryptic positive feedback. If regions containing stem cells randomly grow large enough, cells in their centers will experience reduced negative feedback, favoring further growth of such regions. See Table 1 for parameter values.

deterministic solutions, stochastic simulations starting from conditions expected to lead to stability occasionally ended in unbounded growth (such trajectories are shown in black in Fig. 4C,F,I and L). This is just the sort of behavior that was seen in Fig. 3. By plotting stochastic trajectories onto the models' phase diagrams, it could be seen that escape from control arises when random fluctuations allow the system to cross, on occasion, the separatrix that divides the two possible modes of behavior.

The cartoon in Fig. 4M provides a rationale for such observations. Here stem (orange) and terminal (blue) cells are envisioned growing in one dimension, with terminal cells sorting automatically to the outside of the domain. The gradient of a diffusible feedback factor produced by terminal cells is depicted as blue dots. If, during any cell cycle, the size of the stem cell domain increases, then the average level of feedback within it will fall. The fall will be particularly large when the size of the stem cell domain exceeds the characteristic decay length of the diffusible factor. Since less feedback means more renewal, the chance that the stem cell domain will grow even larger in the next cell cycle goes up. That will lower feedback even further, further elevating renewal, until unbounded growth ultimately ensues. In effect, spatial inhomogeneity generates a cryptic positive feedback loop.

In Fig. 4, spatial inhomogeneity was enforced by the model itself (which assumes perfect cell sorting), but it is almost certainly the same phenomenon that accounts for stochastic escape in agent-based modeling (Fig. 3), where cell sorting was not mandated. This is because fluctuations in the outcomes and directionality of stem cell divisions should lead, at random, to the creation of small islands disproportionately populated with stem cells and others disproportionately populated with terminal cells. Should any one of the stem cell islands get large enough (compared with the characteristic decay length of the feedback signal), and should the average self-renewal probability within such a domain exceed ½, then effectively the same cryptic positive feedback loop as in Fig. 4 would exist.

Positive feedback on self-renewal is sufficient to generate stochastic instability

Dynamical systems with more than one attractor state are often generated by circuits that incorporate positive feedback, sometimes combined with negative. Indeed, the behaviors of renewal control systems that utilize both positive and negative feedback on self-renewal have previously been analyzed[14]. Under appropriate conditions, such systems can be bi-stable—i.e. admitting of two stable steady states (or two possible final states).

Here we show that, under other conditions (different relative strengths of the two types of feedback), such systems can be bi-modal, producing either stability or unbounded growth, depending on the initial conditions: In Figure 5, cell dynamics were modeled using ODEs (space was not explicitly considered), and positive feedback on stem cell renewal was directly introduced, with the source either being stem cells themselves (Fig. 5A-F), or terminal cells (Fig. 5G-I). For simplicity, positive feedback was modeled as inhibition of negative feedback, although other formulations are possible[14].

The results indicate that, for appropriate parameter values, mixed negative-positive feedback systems can also be bimodal, exhibiting phase diagrams much like those in Fig. 4. Moreover,

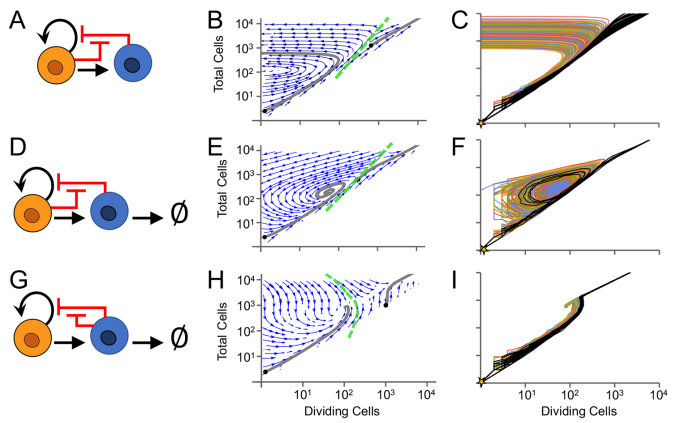

Figure 5. Positive feedback allows for stochastic escape from growth control. Each row presents the analysis of one of the mixed positive-negative feedback models in **A, D, and G.** Phase portraits are shown in **B, E, and H,** and stochastic trajectories (500 instances) are plotted in the same phase space in **C, F and I**. Trajectories that escape growth control are shown in black. Initial conditions consisted of a single dividing cell. See Table 1 for parameter values.

stochastic simulations that started from initial conditions expected to achieve stability occasionally displayed complete loss of control, much as was also observed in Figures 3 and 4. Examination of the trajectories that escaped control suggests that, as in Figure 4, they did so because stochastic fluctuations drove them across the separatrix that divides the phase diagram into regimes of controlled and uncontrolled growth.

In summary, positive feedback on dividing cell renewal, when present, creates a scenario in which stochastic effects can produce rare escape from otherwise well-controlled growth. Such positive feedback may be an explicit process—e.g. a consequence of cell signaling—but it need not be, as effects equivalent to positive feedback tend naturally arise spontaneously, solely due to random fluctuations in spatial homogeneity.

Reversibility of collective transitions

When thinking about rare, catastrophic events of unbounded cell proliferation, cancer naturally comes to mind. Although cancers have long been recognized as resulting from stochastic events, the traditional "multi-hit model" of carcinogenesis equates such events with the mutation of oncogenes or tumor suppressor genes[20,21]. Increasingly, it has been recognized that changes of a non-genetic nature, such as cells randomly flipping between epigenetic states, could also account for steps in cancer progression, and possibly also cancer initiation[22,23]. In the cancer literature, non-genetic state changes are nearly always described in cell-autonomous terms, i.e., as something that happens to a single cell, much like mutation[22]; indeed use of the term "epimutation" to describe such events has become common[24,25].

Collective stochastic transitions—i.e., "tipping point" phenomena that arise out of fluctuations in collective behavior—are relatively common in organismal biology[26], and are sometimes studied in developmental biology[27], but have received little attention by cancer biologists (for an exception see [28]). A unique feature of collective transitions is their propensity for reversibility in response to changes in collective circumstances. For example, if a tumor had escaped growth control through a collective transition, e.g., as modeled in Fig. 3-5, one might expect to be able to return those cells to homeostasis just by altering their numbers or spatial arrangements. In essence, if stochastic fluctuations drove crossing of a separatrix from controlled to unbounded growth, then an appropriate change of circumstances should be able to drive cells back across that same separatrix.

Simulations support this intuition. In Figure 6 the underlying model was, for convenience of analysis, the mixed positive-negative feedback model of Fig. 5G. Panel A shows 100 trajectories that all initiated from near the deterministic steady state for this system, at 4,000 total cells (700 dividing and 3300 terminal). Notice how two trajectories randomly "escape" control and start to grow without bound (similar to the black trajectories in Fig. 5G). To mimic cancer treatment, we started from the point at which one of the escaped trajectories had reached 15,000 cells, and then removed a fixed number of those cells (both dividing and terminal), in the process driving the system to the other side of the separatrix shown in Fig. 5H.

Simulation was then re-initiated from initial conditions in which cell removal had driven the system to a point where total number was either above (panel B), somewhat below (panel C) or far below (panel D) the predicted (deterministic) steady state (see figure legend). In all cases, a large proportion of the ensuing trajectories returned to tightly controlled fluctuations, indicating resumption of homeostasis. At long times, all three conditions behaved similarly, with occasional trajectories re-escaping from control, however their short-term behaviors differed. In B, a large proportion of trajectories returned quickly to unbounded growth, whereas that did not occur in C, Interestingly, panel D exhibited an early burst of escaping trajectories like those in panel B. The early escapes in B undoubtedly reflect of the fact that the initial condition was sufficiently close to the separatrix that random re-crossing was relatively likely. The early escapes in D, however, reflect a different phenomenon: the inherently oscillatory nature of renewal control[5], which causes perturbations that go too far below a steady state to elicit compensatory overshooting. Such overshoots enable a large proportion of trajectories to make close approaches to the separatrix, again allowing for a high probability of stochastic re-crossing.

The abrupt removal of a large fraction of the cells from an escaping simulation might be seen as a model of tumor treatment by surgical removal or ablation. As many cancers are treated non-surgically, e.g., with various dosing schedules of radiation, chemotherapy, or immunotherapy, we asked whether those kinds of treatments might also be expected to "normalize" the behaviors of growth-escaped cells. The results are summarized in Fig. 6E-J, in which the strength or duration of treatment is varied on the abscissa, and the ordinate shows the fraction of trajectories that either resumed homeostasis ("normalization"), fluctuated to zero ("extinction") or exhibited early resumption of uncontrolled growth ("relapse").

As can be seen in Fig. 6E-G, maximal normalization often occurs at relatively low treatment strengths, giving way to a high level of relapse at treatment strengths that begin to produce a large amount of extinction. Interestingly, whereas minimizing relapse with short term treatments requires fine-tuning treatment strength, it appears that low-intensity treatment over a long

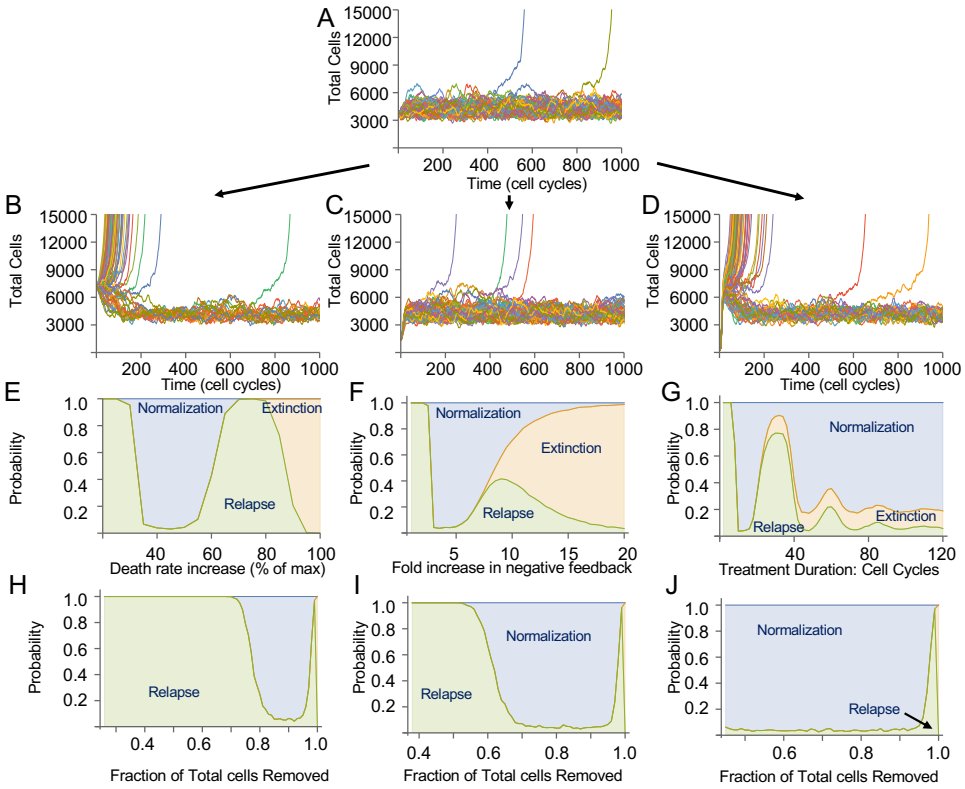

Figure 6. **Transient treatments can reverse stochastic escape from growth control. A.** 100 stochastic trajectories of the model in Fig. 5G initialized near steady state showing rare (~1-2%) escape from growth over course of 1,000 cell cycles. **(B-D)** Simulations initialized above (B), somewhat below (C), and well below (D) steady state. Note the high level of early "relapse" in B and D. (**E-G)** Simulations were initialized from a high level (15,000 cells, of which 50% were stem cells) and then subjected to various "treatments", i.e., temporary or persistent changes to parameters, to determine the fraction of trajectories that normalized (returned to steady state), relapsed, or went entirely extinct, over the course of 1,000 additional cell cycles. In (E) treatment consisted of increasing the probability of death of dividing and terminal cells, by the indicated amount, for 5 cell cycles. In (F) treatment consisted of increasing the strength of negative feedback by indicated factor for 5 cell cycles. In (G) negative feedback was increased by a factor of two for the indicated number of cell cycles. **H-J.** Simulations were initialized as in E-G, and treatment consisted of a one-time removal of the indicated fraction of total cells, but with the removed cells having different proportions of stem vs. terminal cells. In (H) the proportion was 2:1; in (I) it was 4:1 and in (J) it was 9:1.

duration often does equally well. Although these results come from the exploration of a simplified, "toy" model, it is intriguing that they predict a phenomenon that has been observed in the clinic, which is that long-duration treatment of cancers with low dose chemotherapeutic agents (so-called "metronomic therapy") sometimes achieves unexpectedly good results[29-32]. Results in Fig. 6H-J further show that the success of ablative therapies can depend strongly on which type of cell is primarily being ablated—dividing cells or their non-dividing (or slowly dividing) offspring.

Are there signs of stochastic instability in clinical data?

The analyses above suggest that, for cancers in which the final step in malignant transformation is a collective, stochastic transition, tumor treatment might sometimes succeed not because cells capable of initiating cancer are entirely eliminated, but because remaining malignant cells return, at least temporarily, to homeostatic growth. In fact, there are some striking examples of cancer "normalization" in animal models[33-35], as well as rare human tumor types that spontaneously normalize without treatment[36]. There is also growing recognition that very late recurrences of tumors following treatments intended to be curative reflect "dormant" residual disease, either at the primary tumor site or elsewhere. Although mechanisms underlying dormancy are unknown[37], it likely involves at least some degree of restoration of normal cell behavior to cells that were once cancer.

It is improbable that normalization could ever occur by chance reversion of mutations or epimutations that act in a cell-autonomous fashion. Even if back-mutation (or back-epimutation) were not rare at the level of the individual cell, it would be necessary for multiple cells all to undergo such an event at the same time to restore homeostasis—which would be extremely unlikely. For cancer cells that have disseminated, it has been proposed that normalization may occur because of deficiencies in the cellular environment, but exactly how that might happen, and how escape would subsequently occur, is not well understood.

With collective transitions, both the establishment and escape from dormancy are far easier to explain. As Fig. 6 illustrates, with systems that have escaped growth control collectively, it should be possible to re-establish a controlled steady state solely by sufficiently decreasing the number of cells or adjusting, even transiently, the parameters of feedback.

Does this ever happen when cancers are treated? One way to investigate this is to examine the kinetics with which tumors

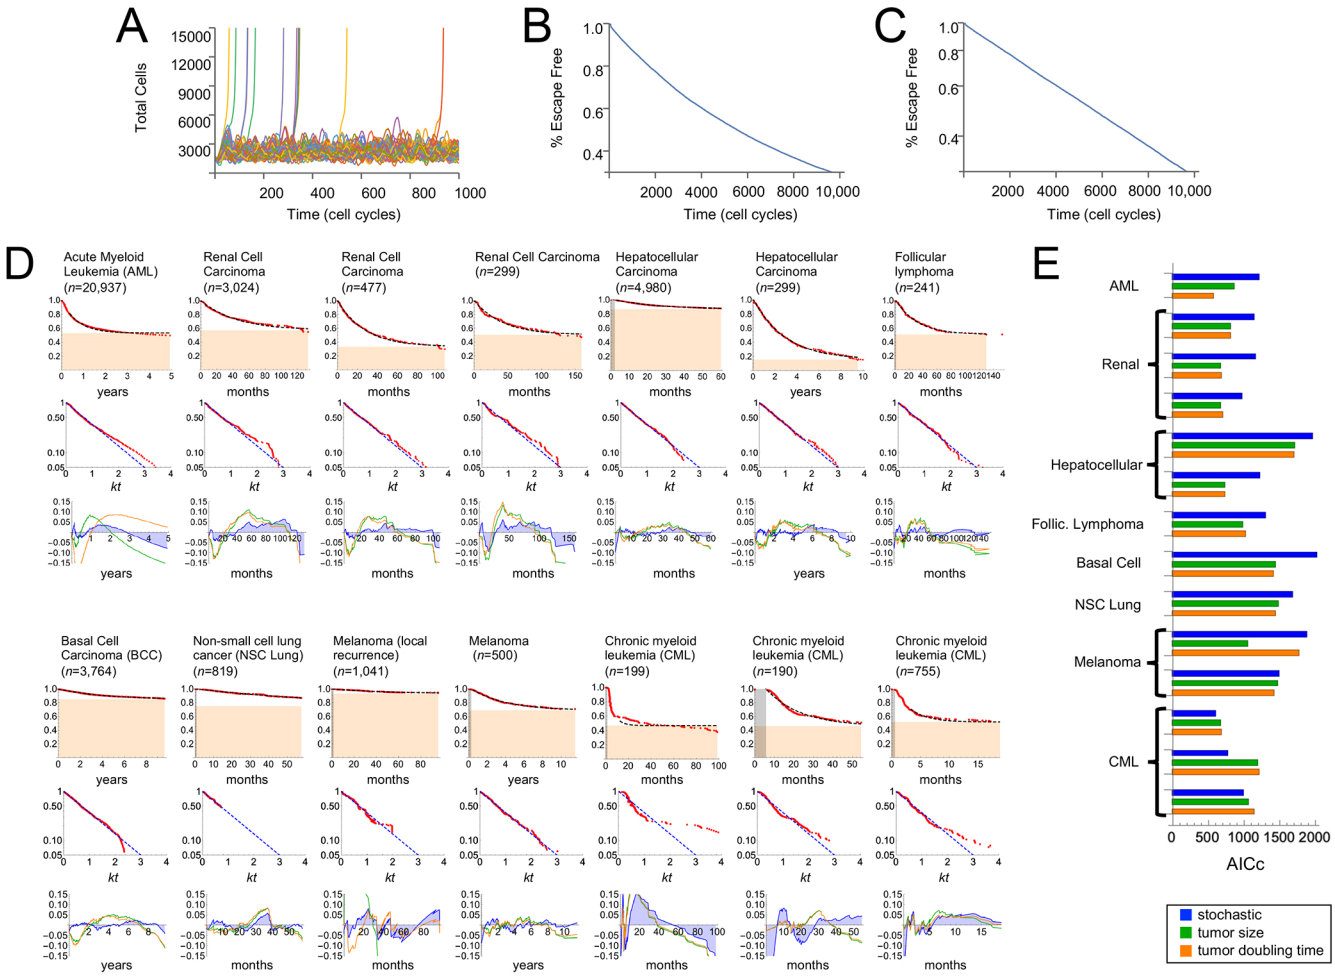

Figure 7. Single-hit kinetics in cancer recurrence. A. 100 stochastic trajectories of the model in Fig. 5G initialized near steady state run long enough to observe ~12% escape. **B-C.** Escape-free survival (defined as never exceeding 100,000 total cells) as a function of time, for 25,000 simulations from the same initial conditions as those in A. Axes are linear in B and logarithmic in C. **D-E.** Reported kinetics of cancer recurrence after treatment with curative intent. In D, for each cancer type, the top graph reproduces a published Kaplan-Meier curve (red symbols) and blue curves represent the fit to a "single-hit" stochastic model (the fitted fraction of patients "cured" and thus not available to relapse is shown as orange fill). In each middle graph, the same data were transformed to remove the influence of the cured fraction and any detection lag, so that single-hit kinetics would be expected to follow the dashed blue line of slope 1 (the abscissa shows time scaled to the fitted growth rate constant, *k*). Each lower graph plots the fit residuals for the fits in the upper graphs (the differences, over time, between the fit and the data points), not only for the single-hit stochastic model (blue curve with blue filling) but also models based on deterministic regrowth from a distribution of tumor sizes (green) or following a distribution of tumor doubling times (orange). Note that the deterministic models tend to overfit early times and underfit late ones. In E, the corrected Akaike Information Criterion (AICc)—a measure of goodness of fit—is plotted for each of the three models for each cancer type. References for these 14 studies are, in the order shown above, AML[49], renal[50], renal[51], renal[52], hepatocellular[53], hepatocellular[54], follicular lymphoma[55], basal cell carcinoma[56], non-small cell lung carcinoma[57], melanoma[58], melanoma[59], CML[46], CML[60] and CML[61]. See also Table S2.

recur after treatments that were intended to be curative. If, following treatment (or following metastatic dissemination), tumor cells undergo normalization through reversal of a collective transition, then the kinetics of recurrence should match those of a random process. This is illustrated by Fig. 7A-C, in which we simulated what happens to stochastic trajectories produced by the model in Fig. 6 after simulated treatment (removal of a large number of cells). With sufficiently long follow-up, trajectories that had been restored to homeostasis re-escaped at times that exactly follow an exponential distribution (Fig. 7B-C). Exponential waiting times are indicative of a "single-hit" stochastic process, as they are a reflection of a probability rate (also known as a hazard rate) that is constant in time.

We wondered whether similar exponential waiting times might be observed in the recurrence of actual cancers following treatment. If so, it would strongly suggest recurrence is a "single-hit" stochastic process, consistent with a model in which tumor cells have been driven back over a separatrix across which they must once again fluctuate to escape.

We therefore gathered published clinical data from a broad range of cancers treated with curative intent (i.e., with the expectation of long-term survival of a substantial fraction of patients), focusing primarily on those in which treatment was brief and follow-up was treatment-free (i.e., we rejected cases in which recurrence could be explained by the acquisition of resistance to ongoing therapy). In such cases, the traditional explanation for recurrence is that a local or distant tumor residuum remains after treatment and continues growing until a

detectable size is again reached. Even though individual residua may not all start from the same size or grow at the same rate, such an explanation is still deterministic—knowing the size and growth properties of an individual tumor residuum, one will know when that tumor would reappear. The probability rate associated with recurrence would thus not be constant in time; rather it should be peaked around the rate associated with the average tumor residuum size.

In principle, it should be possible to use clinical data on cancer recurrence to distinguish whether tumors reemerge deterministically or stochastically. In practice, however, we don't usually know the details of how tumor residua are distributed with respect to size and growth rate, the size at which recurring tumors become recognizable, nor the fraction of patients, if any, that are cured by treatment (i.e., have no residua, and thus no probability of recurrence). Allowing these parameters to vary freely enables both stochastic and deterministic models to fit most clinical data to some degree. Still, we may investigate whether one type of model consistently does a better job of fitting across multiple data sets.

Such an analysis is shown in Figure 7D-E using data from 14 different clinical studies. These include tumors treated by surgical resection, ablation and bone marrow transplantation, or targeted pharmacological therapy. Recurrences included both local, locoregional and metastatic. For each cancer, the upper and middle panels in Fig. 7D compare published Kaplan-Meier curves (red) with fits to a single-hit stochastic model (blue) that allows for a detection lag (gray) and a "cured" patient fraction (orange). The middle panels show how, for most of these cancers, recurrence in at-risk patients is consistent with exponential kinetics (constant probability rate). The one clear exception is chronic myeloid leukemia (CML), in which recurrence is notably biphasic.

The lower panels show the fit residuals (measures of the difference between data and fit at each time point) for the stochastic model and compare them with two deterministic models with the same number of parameters, in which tumors either start from a distribution of sizes or follow a distribution of doubling times (see Methods and Appendix). The residuals of the stochastic model tended, on average, to be smaller than those of either deterministic model. More importantly, the residuals of the deterministic models tended to be negative at early times and positive at intermediate times, as expected for models with a peaked probability rate distribution. In addition, for the deterministic models, the fit parameters describing the distributions of initial residuum sizes and doubling times were often at the extremes of their allowable ranges (Supplemental Table 3).

Overall goodness of fit for the three model types was quantified in Fig. 7E using the Akaike information criterion (AICc). For all but CML, the stochastic models displayed the best AICc values. Data from an additional six studies are presented in Supplemental Figure 1. In some of those cases treatments were protracted, or we were unable to determine treatment durations, or follow-up times were very short compared to the size of the apparent at-risk population. For these reasons, confidence that such data should fit any of the proposed models was lower; nevertheless, the stochastic model performed at least as well, if not better than, other models in the majority of these cases.

DISCUSSION

The results presented here come both from exploring theoretical models and fitting such models to clinical observations. The goal of exploring models was to build intuitions about how stochastic fluctuations and feedback control of cell proliferation might interact in unexpected ways. What was shown was that feedback circuitry that normally guarantees stability and robustness over a large portion of parameter space can, in the presence of positive feedback, also be the source of rare, stochastic instability.

The kind of positive feedback required to do this can be either explicit or cryptic. Cryptic positive feedback was shown to arise in any spatial system in which the spread of a feedback signal can in some circumstances not keep up with domain growth—at the very least this means any system in which feedback depends on diffusion, or any process that decays similarly over space (e.g., mechanical tension[38]). It arises particularly easily when stem cells and their differentiated offspring explicitly sort away from each other (Fig. 4) which, as it happens, is a common situation in many, if not most, epithelia, wherein stem and progenitor cells reside in a distinct compartment, usually next to the basement membrane. However, as we show by agent-based modeling (Fig. 3), such a situation even arises when stem and differentiated cells simply remain near where they are born and are only moved about passively (Fig. 3).

In addition to the possibility of cryptic positive feedback due to spatial limitations, there are many tissues in which positive feedback may be explicit, through the combined actions of feedback factors that have opposing effects on self-renewal[14]. For example, tissues that use TGFβ-family molecules to feedback negatively on stem cell renewal commonly co-express MAP kinase-stimulating ligands that feedback in the opposite direction. Thus, beyond the cryptic feedback that arises due to spatial effects, there are strong reasons to expect the potential for stochastic instability to exist broadly in renewal-controlled tissues.

Stochastic instability may thus be viewed as a generic, structural shortcoming of renewal control circuitry. The idea that stabilizing, homeostatic feedback can be a source of instability may seem surprising, as control processes typically dampen, not enhance, stochastic effects (as, for example, in Fig. 1B). Such behavior nicely exemplifies what Doyle and colleagues termed the "robust-yet-fragile" nature of complex control systems—where the price of good control is typically a predisposition to rare, catastrophic failures[39].

Even if self-regulating tissues can never avoid the possibility of spontaneous escape from control, one would expect them to be under evolutionary pressure to minimize the risk, which they could presumably achieve through appropriate choices of parameters (e.g., stronger negative feedback, longer decay lengths of feedback gradients). Thus, under normal circumstances, the probability of spontaneous escape from growth control during an organism's lifetime might be kept close to zero. Oncogenic mutations, however, might change this. Indeed, thinking along these lines enables one to conceptualize oncogenes less as instructive signals that tell cells to grow, than as factors that adjust parameters so that the risk of spontaneous loss of control becomes significant. Such a view could help explain why the expression of oncogenes (or loss of expression

of tumor suppressor genes) is so often associated with little or no phenotypic effect in vivo[40,41], and why transgenic animal models with widespread oncogene expression commonly result in the production of only small numbers of tumors, e.g. [42,43]. Whether collective, stochastic transitions to unbounded growth indeed drive cancer initiation or progression remains to be determined empirically, but several of the modeling results presented here support that view. As shown in Fig. 6, the unique tendency of collective transitions to reverse themselves following removal of cells or appropriate changes to parameters provides both a model for cancer dormancy and a justification for the ability of metronomic chemotherapy to perform comparably to high-dose, short term treatments[29-32].

Yet by no means do these results imply all cancers should be able to revert after treatment. Even if collective transitions do play a role in carcinogenesis, subsequent cell autonomous ones (e.g., mutations and epimutations) that are effectively irreversible could block any return to normalcy. Still, the idea of reversion provides a conceptual framework for understanding cancer dormancy that does not require postulating the existence of any novel, cancer-specific mechanisms.

Clearly, the fact that cancer recurrence kinetics often fit a model in which re-emergence involves a single stochastic event (Fig. 7) does not prove the event is the one described here. Other phenomena, such as escape from immune surveillance, or induction of neovascularization, could produce similar kinetics if they occurred truly at random, but appropriate random mechanisms (e.g., mechanisms with a low, constant probability rate, such as mutation and epimutation) have not so far been found as necessary triggers of such processes. At the very least, the idea of stochastic escape from collective control provides a mechanistically plausible alternative.

The one cancer type that was not well fit by the single-hit stochastic model, CML, displayed recurrence that was rapid at early times, slowing distinctly afterward (Fig. 7D, CML data). Similar bursts of early recurrence have been reported following treatment of other cancers (e.g., breast cancer), and have been suggested to reflect systemic growth-stimulatory effects triggered by surgery[44,45]. In the case of the CML, however, there is no surgical intervention; recurrence is measured following the abrupt withdrawal of a long-term suppressive pharmacotherapy, so the explanation must lie elsewhere.

Interestingly, excessive early recurrence is precisely what was seen in the simulations of Figure 6D, in which simulated tumors were reduced to a very small size by a treatment that was abruptly withdrawn. As mentioned above, the explanation for such behavior is the inherently oscillatory nature of renewal control[5], which creates greater opportunity for stochastic escape when cell numbers start too far below their stable steady state values. In fact, in one of the CML studies included in Fig. 7, *BCR-ABL* transcripts were tracked throughout the study and frequently showed marked oscillations prior to the return of leukemia [46]. Although such similarities do not establish definitively that CML relapse is due to collective stochastic escape, the plausibility of that explanation stands in contrast to the presumed alternative: that treated patient populations are intrinsically heterogeneous. Determining which view is correct has obvious implications for patient management.

The mathematical models explored here were kept deliberately simple, and the shapes of the distributions underlying deterministic models were chosen somewhat arbitrarily. Other models may well fit recurrence data equally well, especially models with greater flexibility due to a greater number of parameters. Indeed, others have generated empirical models of cancer recurrence requiring an arbitrary number of slow, sequential, stochastic events, and these also provide a good fit to clinical data[47]. Still, given that such transitions must be relatively rare to account for broad ranges of recurrence times, it is not easy to identify many possible cellular mechanisms for producing them. What we can say here is that we have identified a very simple mechanism, one that emerges spontaneously, and unexpectedly, out of the need of tissues to control their own sizes.

METHODS

Deterministic models of growth control (non-spatial)

Renewal control—in which differentiated cells dose dependently feedback upon the renewal probability of their progenitors—was adapted from[4], which introduces the following equations:

$$c0'(t) = c0(t)(2p - 1) \quad (1)$$

$$c1'(t) = 2c0(t)(1 - p) - d * c1\,(t) \quad (2)$$

In which *c0* is the number of dividing progenitor cells, *c1* is the number of terminally differentiated cells, *d* is the turnover rate constant for differentiated cells, and *p* is the probability that a daughter cell created upon division remains dividing. If *p* is a constant, it must equal exactly ½ or the system cannot achieve a non-zero steady state. If *p* is subject to negative feedback from *c1*, i.e. is a declining function of *c1* that has a value above ½ when *c1*=0, then, for *d*>0, the system reaches a stable steady state in which dividing and differentiated (terminal) cells co-exist; if *d*=0, the system reaches a final state entirely comprised of terminal cells[14].

For both ODE and non-spatial stochastic models, the form of the function that describes *p* differed depending on the type of feedback used into the model.

Pure negative feedback from terminal cells was represented using:

$$p = \frac{p_0}{\gamma * c1^{n1} + 1}$$

Mixed negative feedback from terminal cells and positive feedback from dividing cells was represented using:

$$p = \frac{p_0}{\frac{\gamma * c1^{n1}}{\phi * c0^{n2} + 1} + 1}$$

And positive and negative feedback from terminal cells was represented using:

$$p = \frac{p_0}{\frac{\gamma * c1^{n1}}{\phi * c1^{n2} + 1} + 1}$$

In these formulae, p_0 is the renewal probability in the absence of feedback, γ and ϕ are constants representing the strength of negative and positive feedback respectively, and *n1* and *n2* are the gains of negative and positive feedback respectively. Parameter values and initial conditions for all figures can be found in Table S1.

Growing disk models

To model for the effect of space on the distribution of feedback signaling without having to solve partial differential equations, we constructed an idealized scenario in which a collection of cells was represented as a disk, the area of which was

proportional to the number of cells. We then evolved discs according to equations 1-2, except the formula for *p* depended on the expected spatial distribution of a feedback signal produced by terminal cells uniformly distributed throughout the disk. Specifically, we modeled negative feedback as being carried by a factor that is produced in proportion to the number of terminal cells in the disk, that diffuses and linearly decays, reaching a steady state diffusion gradient rapidly, relative to cell growth.

In such a situation it has been shown[19] that the steady state shape of the diffusion gradient produced is of the form

$$c(x) = v(1 - \frac{r}{\lambda} I_0\left(\frac{x}{\lambda}\right) K_1\left(\frac{r}{\lambda}\right)) \quad (3)$$

Where *c* is the concentration of the diffusing factor, *v* its rate of production, *x* the radial distance from the center of the disk, *r* the radius of the disk, λ a constant representing the decay length of the diffusing factor, and I_0 and K_1 being modified Bessel functions of the first and second kinds, respectively. To compute the renewal probability *p* for disk as whole we integrate to obtain the average of

$$p = \frac{p_0}{1 + c(x)/k}$$

across the disk, where p_0 is the renewal probability in the absence of feedback and *k* is a constant. In this way we obtain a single value that can be used in equations 1-2.

Under the assumption that dividing and differentiated cells were uniformly mixed within the disc, *v* was taken to be proportional to the fraction of total cells that was differentiated. In some simulations, disks were divided into an inner disk and outer annulus with dividing and differentiated cells assumed to sort immediately into one or the other compartment. In such cases it was necessary to use the formula for the diffusion gradient outside of a uniformly producing disk:

$$c(x) = v(\frac{r}{\lambda} K_0\left(\frac{x}{\lambda}\right) I_1\left(\frac{r}{\lambda}\right)) \quad (4)$$

If the dividing cells occupied an annulus outside a disk of differentiated cells, Eq. 4 was used directly; if dividing cells occupied a disk inside an annulus of differentiated cells we calculated equation 3 for both the inner and outer radii and subtracted one result from the other. For parameter values see Table S1.

Stochastic simulations of renewal control

Monte Carlo simulations were used to model the effects of fluctuations in cell number due to the probabilistic nature of cell differentiation and death. At fixed time steps all dividing cells were duplicated, and the fates of the resulting cells were probabilistically determined. Additionally, the survival of each differentiated cell was probabilistically determined, according to a constant probability of *d* of disappearing per time step. The order in which events were executed was:

1) The probability *p* of renewal was computed using the current number of dividing and differentiated cells.
2) The number of differentiated cells that die was chosen probabilistically, given *d*, and removed from the system.
3) All dividing cells were replaced by two daughter cells, and they were assigned fates so that, for any given value of *p*, the expected number of dividing cells would be *p* times the number of daughter cells. There are multiple ways in which this can be implemented, however, which differ based on the degree of correlation or anticorrelation to expect between sister cells. One approach is to have each daughter cell behave independently, differentiating with probability 1-*p*. This was done for the stochastic simulations in Figures 1-3. However, some stem cell populations may be inclined toward asymmetric divisions (sisters choosing opposite fates) as opposed to symmetric one (sisters choosing the same fate). Because asymmetric divisions balance opposing fluctuations, they tend to reduce stochastic effects. To capture bias toward asymmetry, we note that, for any given value of *p*, the maximum proportion of divisions that can ever be asymmetric is $1 - 2\left|p - \frac{1}{2}\right|$; we thus define a parameter α such that sister cells have a probability of $\alpha(1 - 2\left|p - \frac{1}{2}\right|)$ of producing one dividing and one differentiated cell. In Figures 4-5 we use α=1 (maximal asymmetric bias), except for panel 4F where we use α=0.5.

Stochastic simulations were run until either no dividing cells remained; total cell numbers reached a predetermined limit (generally 250,000 cells); or a predetermined number of time steps had elapsed. For parameter values see Table S1.

Agent Based Modeling

Agent-based models were used to explore stochastic spatial effects. Conditions with the potential to reach a static final state (with no dividing cells remaining) were obtained by simulating renewal control in the absence of terminal cell death. Conditions with the potential to reach a dynamic steady state were obtained by simulating renewal control in the presence of terminal cell death.

In the first case, simulations were carried out on an 800 x 800 square grid and initiated with two adjacent dividing cells juxtaposed with two adjacent terminal cells in the center. At each time step, cells could divide, and rules were used to determine both the fate of each offspring (renewal vs. differentiation) and how neighboring cells would adjust their positions to the extent necessary to create an open space for a newly generated cell. At each time step all dividing cells divided simultaneously, but the movements of other cells to accommodate new offspring were calculated sequentially.

The fate of every newborn daughter cell was assigned probabilistically, determined by the strength of local negative feedback field at its position. Field strength was determined by summing the diffusion fields produced by every terminal cell. Following eq. 4, each which was assumed to be of the form

$$\frac{1}{2\lambda} I_1(\frac{1}{2\lambda}) K_0(\frac{r}{\lambda})$$

Where *r* is the radial distance from the terminal cell, λ is a constant representing the decay length of the diffusing factor (in units of cell diameters) and I_1 and K_0 are modified Bessel functions of the first and second kinds, respectively. Renewal probability at each location was then calculated as

$$\frac{p_0}{1 + (f/k)^2}$$

where f is the field strength at that location, γ is a constant and p_0 the renewal probability in the absence of feedback, which was usually set to 1. After the fates of new cells were assigned, space was made available by identifying the nearest unoccupied grid point to the location of each cell that divided (if more than one was equidistant one was chosen at random), and then "shoving" occupied cells toward that space. If that space was at the same horizontal or vertical level as the dividing cell, shoving consisted of simply inserting an empty cell either above or below, or to the left or the right, of the dividing cell, chosen at random. If the space differed in both its horizontal and vertical coordinates, both a row and a column (the first of which being chosen at random) were shoved to make space for the new cell while also filling the unoccupied point. This process was carried out sequentially (asynchronously) to allocate every newly generated cell to its own position. Typically, simulations were run until no dividing cells remained, or dividing cell numbers grew too numerous to simulate in reasonable time frames (about 70,000 dividing cells).

To simulate situations in which terminally differentiated cells die, it was necessary to apply a constant probability of disappearance to every terminal cell and implement a rule to close up the spaces created by terminal cell death. These simulations were carried out on an 1100 x 1000 square grid, initialized from approximately steady state conditions (to identify steady state conditions, we initialized 100 simulations with a 7 x 7 grid of cells whose starting identities were chosen at random, and identified a typical configuration characteristic of systems that remained close to a common state).

At each time step, a random subset of differentiated cells was selected to die based on a death parameter, and subsequently cells surrounding the unoccupied space collapses inward to fill the void. Specifically, a cell furthest cell from the void was identified, and cells were shoved from that location in the manner described above, such that the void became filled. Additionally, cells divided and created additional space as described above for final state modeling, and differentiated following the same rules except that in this case the feedback function was slightly less steep, namely:

$$p = \frac{p_0}{1 + (f/k)}$$

where f is the local strength of the feedback field, k is a constant, and $p_0 = 1$.

For any given set parameters, a group of simulations were run until most fluctuated about a relative constant level of dividing and terminal cells, and this apparent steady state was used subsequently as an initial condition for the analyses shown. Simulations were then run until either all dividing cells had extinguished themselves; the number of dividing cells exceeded 500,000; or 100 cell cycles had elapsed, whichever came first.

Final state simulations were performed using *Mathematica* software; steady state simulations were hand-coded in Julia. All code will be made freely available upon request.

Simulating tumor treatment

A non-spatial stochastic model of renewal control, which includes positive feedback but can be run as an ODE model, facilitating analysis, was used to explore how tumor treatment might influence stochastic escape from growth control. In this case the *p*-value formula was

$$p = \frac{p_0}{\frac{\gamma * c1^{\mathrm{n1}}}{\phi * c1^{\mathrm{n2}} + 1} + 1}$$

which also corresponds to the scenario modeled in Fig. 5G.

It was determined empirically that the steady state total cell number, for the parameters chosen (Table S1), was ~4200 and that while most simulations fluctuated above and below that number, a threshold of 15,000 total cells was sufficient to distinguish simulations that had truly escaped control, from ones that were merely fluctuating at random.

To model treatment of tumors that had collectively escaped growth control, we initialized simulations at 15,000 total cells and then, to simulate therapy, changed parameters: In the first case, we increased the rate of turnover of both dividing and differentiated cells, according to the formula

$$\hat{d} = d + (1 - d) * \sigma$$

Where d is the previous probability of death per cell per timestep (before treatment) and $0 \leq \sigma \leq 1$ is a parameter representing the strength of treatment.

In a second case, treatment was implemented as an increase in negative feedback by multiplying the feedback strength γ by a positive constant. In a third case the feedback strength γ was increased by a factor of two and treatment duration was varied

In the remainder of cases, treatment consisted not in a change of parameters, but of initial conditions, corresponding to the removal of a fixed proportion of dividing cells, differentiated cells or both. The same parameters used in the previous treatment panels were used here. Treatment was begun at 15,000 total cells of which half were dividing and half differentiated. Treatment was implemented as a change in initial conditions, which were adjusted to be $15{,}000\alpha(1-\sigma)$ and $15{,}000(1-\alpha)\,(1-\sigma)$, for dividing and differentiated cells, respectively, where $\sigma\ (0 \leq \sigma \leq 1)$ quantifies overall treatment strength and $\alpha\ (0 \leq \alpha \leq 1)$ quantifies any bias of treatment toward selectively removing one cell type or the other. The values of σ and α were chosen such that the new initial conditions would consist of dividing and differentiated cells both between 0 and 7,500. Parameters values can be found in Supplemental Table 1.

Fitting cancer recurrence data

Cancer recurrence data were obtained from published studies (see Supplemental Table 2), in most cases by extracting data points directly from published Kaplan Meier curves (or their equivalent) and resampling in proportion to the estimated number of patients at each time point. Disease-free survival (as a fraction of total study subjects) was fit to each of three models.

<u>Model 1</u> ("stochastic emergence from dormancy") assumes each residuum has a constant probability per time of resuming growth, expressed as *ln*2 over a rate constant *k*. Once growth has started, a delay of ψ must elapse before recurrence is detected.

<u>Model 2</u> ("variable tumor residuum size") assumes the sizes of tumor residua or undetected metastases are distributed according to probability density function $Prob[x] = (1-x)^{-1+\frac{1}{\mu}}$, where the mean size μ is scaled relative to the size at which recurrence becomes detectable (i.e. μ=0.25 implies that the average size is one fourth of the size threshold for detection; for

details see Appendix). Tumors are assumed to grow with rate constant k (i.e., doubling time $ln2/k$). We constrain μ to be less than 0.33, as higher values would imply that the *average* residuum reaches detectable size in less than 1.6 doubling times.

Model 3 ("variable tumor growth rates") assumes a common tumor residuum size and tumor doubling times drawn from a log-normal distribution. The parameter χ characterizes the breadth of that distribution and is related to the coefficient of variation of the distribution, c, by $\chi = \frac{\sqrt{ln[1+c^2]}}{\sqrt{2}}$. Based on clinical data we constrain to be less than 0.59 [48]. The parameter ϕ represents the time required for a tumor residuum that starts from the average size to reach detectable size.

The parameters ρ, k and ψ (Model 1), ρ, k and μ (Model 2), or ρ, χ and ϕ (Model 3) were obtained by nonlinear fitting using the NonlinearModelFit function of *Mathematica* software (values are shown in supplemental Table 3).

Note that we chose to fit survival curve data directly, rather than first calculate hazard rates from data, because such calculations depend on assumptions about at-risk population sizes, and invariably create artifacts (especially at very early and late times) due to the need for data smoothing.

Note also that the values obtained for ρ (the "cured fraction") were similar for all three models (see supplemental Table 3), implying that differences in quality of fit largely reflect the remaining two parameters. Note also that, for model 1, values of ψ were, in most cases, very small and thus had little influence on model fit. For model 2, fitted values of μ were, in many classes, at the maximum allowable value (0.33), implying that a better fit would have required a large number of residua to be very close to detectable size immediately after treatment, which seems implausible. Similarly, for model 3, for all cases except Chronic Myeloid Leukemia (CML) fitted values of χ were all at the maximum of 0.59, implying that a better fit would have required a greater diversity of tumor growth rates than has been observed clinically [48].

For more details on the derivation and limitations of the models, see Appendix.

Supplementary Material

Figure S1. Fitting additional cancer recurrence data sets to stochastic and deterministic models. Data were analyzed as in Figure 7D-M.

Supplemental Table 1. Simulation parameters

Supplemental Table 2. Cancer data sets used in Fig. 7 and Fig. S1

Supplemental Table 3. Parameters obtained by data fitting (relates to Fig. 7 and S1).

Supplemental Document 1. Appendix: Mathematical derivations and model formulations.

Acknowledgements

This work was supported by NIH grant U54 CA217378. MGC was supported by training grants T32 EB009418 and T32 GM136624. We thank Genti Buzi for helpful advice.

References

1. Getto, P., and Marciniak-Czochra, A. (2015). Mathematical Modelling as a Tool to Understand Cell Self-renewal and Differentiation. Methods Mol Biol *1293*, 247-266. 10.1007/978-1-4939-2519-3_15.
2. Kirouac, D.C., Madlambayan, G.J., Yu, M., Sykes, E.A., Ito, C., and Zandstra, P.W. (2009). Cell-cell interaction networks regulate blood stem and progenitor cell fate. Mol Syst Biol *5*, 293. 10.1038/msb.2009.49.
3. Rodriguez-Brenes, I.A., Komarova, N.L., and Wodarz, D. (2011). Evolutionary dynamics of feedback escape and the development of stem-cell-driven cancers. Proc Natl Acad Sci U S A *108*, 18983-18988. 10.1073/pnas.1107621108.
4. Lander, A.D., Gokoffski, K.K., Wan, F.Y., Nie, Q., and Calof, A.L. (2009). Cell lineages and the logic of proliferative control. PLoS Biol *7*, e15. 10.1371/journal.pbio.1000015.
5. Buzi, G., Lander, A.D., and Khammash, M. (2015). Cell lineage branching as a strategy for proliferative control. BMC Biol *13*, 13. 10.1186/s12915-015-0122-8.
6. Clayton, E., Doupe, D.P., Klein, A.M., Winton, D.J., Simons, B.D., and Jones, P.H. (2007). A single type of progenitor cell maintains normal epidermis. Nature *446*, 185-189. 10.1038/nature05574.
7. Snippert, H.J., van der Flier, L.G., Sato, T., van Es, J.H., van den Born, M., Kroon-Veenboer, C., Barker, N., Klein, A.M., van Rheenen, J., Simons, B.D., and Clevers, H. (2010). Intestinal crypt homeostasis results from neutral competition between symmetrically dividing Lgr5 stem cells. Cell *143*, 134-144. 10.1016/j.cell.2010.09.016.
8. Tallapragada, N.P., Cambra, H.M., Wald, T., Keough Jalbert, S., Abraham, D.M., Klein, O.D., and Klein, A.M. (2021). Inflation-collapse dynamics drive patterning and morphogenesis in intestinal organoids. Cell Stem Cell *28*, 1516-1532 e1514. 10.1016/j.stem.2021.04.002.
9. Rompolas, P., Mesa, K.R., Kawaguchi, K., Park, S., Gonzalez, D., Brown, S., Boucher, J., Klein, A.M., and Greco, V. (2016). Spatiotemporal coordination of stem cell commitment during epidermal homeostasis. Science *352*, 1471-1474. 10.1126/science.aaf7012.
10. Gokoffski, K.K., Kawauchi, S., Wu, H.H., Santos, R., Hollenbeck, P.L.W., Lander, A.D., and Calof, A.L. (2010). Feedback Regulation of Neurogenesis in the Mammalian Olfactory Epithelium: New Insights from Genetics and Systems Biology. In The Neurobiology of Olfaction, A. Menini, ed. (CRC Press). NBK55982 [bookaccession].
11. Manceau, M., Gros, J., Savage, K., Thome, V., McPherron, A., Paterson, B., and Marcelle, C. (2008). Myostatin promotes the terminal differentiation of embryonic muscle progenitors. Genes Dev *22*, 668-681. 10.1101/gad.454408.
12. Marciniak-Czochra, A., Stiehl, T., Ho, A.D., Jager, W., and Wagner, W. (2009). Modeling of asymmetric cell division in hematopoietic stem cells--regulation of self-renewal is essential for efficient repopulation. Stem Cells Dev *18*, 377-385. 10.1089/scd.2008.0143.

13. Gupta, A., and Khammash, M. (2022). Universal structural requirements for maximal robust perfect adaptation in biomolecular networks. Proc Natl Acad Sci U S A *119*, e2207802119. 10.1073/pnas.2207802119.
14. Kunche, S., Yan, H., Calof, A.L., Lowengrub, J.S., and Lander, A.D. (2016). Feedback, Lineages and Self-Organizing Morphogenesis. PLoS Comput Biol *12*, e1004814. 10.1371/journal.pcbi.1004814.
15. Gokoffski, K.K., Wu, H.H., Beites, C.L., Kim, J., Kim, E.J., Matzuk, M.M., Johnson, J.E., Lander, A.D., and Calof, A.L. (2011). Activin and GDF11 collaborate in feedback control of neuroepithelial stem cell proliferation and fate. Development *138*, 4131-4142. 10.1242/dev.065870.
16. Lo, W.C., Chou, C.S., Gokoffski, K.K., Wan, F.Y., Lander, A.D., Calof, A.L., and Nie, Q. (2009). Feedback regulation in multistage cell lineages. Math Biosci Eng *6*, 59-82.
17. Lander, A.D. (2011). Pattern, growth, and control. Cell *144*, 955-969. 10.1016/j.cell.2011.03.009.
18. Wang, M.X., Lander, A., and Lai, P.Y. (2021). Regulatory feedback effects on tissue growth dynamics in a two-stage cell lineage model. Phys Rev E *104*, 034405. 10.1103/PhysRevE.104.034405.
19. Chen, C.C., Wang, L., Plikus, M.V., Jiang, T.X., Murray, P.J., Ramos, R., Guerrero-Juarez, C.F., Hughes, M.W., Lee, O.K., Shi, S., et al. (2015). Organ-level quorum sensing directs regeneration in hair stem cell populations. Cell *161*, 277-290. 10.1016/j.cell.2015.02.016.
20. Fearon, E.R., and Vogelstein, B. (1990). A genetic model for colorectal tumorigenesis. Cell *61*, 759-767. 10.1016/0092-8674(90)90186-i.
21. Tomasetti, C., Marchionni, L., Nowak, M.A., Parmigiani, G., and Vogelstein, B. (2015). Only three driver gene mutations are required for the development of lung and colorectal cancers. Proc Natl Acad Sci U S A *112*, 118-123. 10.1073/pnas.1421839112.
22. Feinberg, A.P., and Levchenko, A. (2023). Epigenetics as a mediator of plasticity in cancer. Science *379*, eaaw3835. 10.1126/science.aaw3835.
23. Baylin, S.B., and Jones, P.A. (2016). Epigenetic Determinants of Cancer. Cold Spring Harb Perspect Biol *8*. 10.1101/cshperspect.a019505.
24. Hitchins, M.P. (2015). Constitutional epimutation as a mechanism for cancer causality and heritability? Nat Rev Cancer *15*, 625-634. 10.1038/nrc4001.
25. Gueant, J.L., Chery, C., Oussalah, A., Nadaf, J., Coelho, D., Josse, T., Flayac, J., Robert, A., Koscinski, I., Gastin, I., et al. (2018). APRDX1 mutant allele causes a MMACHC secondary epimutation in cblC patients. Nat Commun *9*, 67. 10.1038/s41467-017-02306-5.
26. Tunstrom, K., Katz, Y., Ioannou, C.C., Huepe, C., Lutz, M.J., and Couzin, I.D. (2013). Collective states, multistability and transitional behavior in schooling fish. PLoS Comput Biol *9*, e1002915. 10.1371/journal.pcbi.1002915.
27. Schweisguth, F., and Corson, F. (2019). Self-Organization in Pattern Formation. Dev Cell *49*, 659-677. 10.1016/j.devcel.2019.05.019.
28. Brock, A., Krause, S., and Ingber, D.E. (2015). Control of cancer formation by intrinsic genetic noise and microenvironmental cues. Nat Rev Cancer *15*, 499-509. 10.1038/nrc3959.
29. Liu, J., He, M., Wang, Z., Li, Q., and Xu, B. (2022). Current Research Status of Metronomic Chemotherapy in Combination Treatment of Breast Cancer. Oncol Res Treat *45*, 681-692. 10.1159/000526481.
30. Liu, C.T., Hsieh, M.C., Su, Y.L., Hung, C.M., Pei, S.N., Liao, C.K., Tsai, Y.F., Liao, H.Y., Liu, W.C., Chiu, C.C., et al. (2021). Metronomic vinorelbine is an excellent and safe treatment for advanced breast cancer: a retrospective, observational study. J Cancer *12*, 5355-5364. 10.7150/jca.60682.
31. Gnoni, A., Silvestris, N., Licchetta, A., Santini, D., Scartozzi, M., Ria, R., Pisconti, S., Petrelli, F., Vacca, A., and Lorusso, V. (2015). Metronomic chemotherapy from rationale to clinical studies: a dream or reality? Crit Rev Oncol Hematol *95*, 46-61. 10.1016/j.critrevonc.2015.01.008.
32. Hanahan, D., Bergers, G., and Bergsland, E. (2000). Less is more, regularly: metronomic dosing of cytotoxic drugs can target tumor angiogenesis in mice. J Clin Invest *105*, 1045-1047. 10.1172/JCI9872.
33. Mintz, B., and Illmensee, K. (1975). Normal genetically mosaic mice produced from malignant teratocarcinoma cells. Proc Natl Acad Sci U S A *72*, 3585-3589. 10.1073/pnas.72.9.3585.
34. Thomas, G.A., Williams, D., and Williams, E.D. (1991). Reversibility of the malignant phenotype in monoclonal tumours in the mouse. Br J Cancer *63*, 213-216. 10.1038/bjc.1991.51.
35. Maffini, M.V., Calabro, J.M., Soto, A.M., and Sonnenschein, C. (2005). Stromal regulation of neoplastic development: age-dependent normalization of neoplastic mammary cells by mammary stroma. The American journal of pathology *167*, 1405-1410. 10.1016/S0002-9440(10)61227-8.
36. Brodeur, G.M. (2018). Spontaneous regression of neuroblastoma. Cell Tissue Res *372*, 277-286. 10.1007/s00441-017-2761-2.
37. Risson, E., Nobre, A.R., Maguer-Satta, V., and Aguirre-Ghiso, J.A. (2020). The current paradigm and challenges ahead for the dormancy of disseminated tumor cells. Nat Cancer *1*, 672-680. 10.1038/s43018-020-0088-5.
38. Irvine, K.D., and Shraiman, B.I. (2017). Mechanical control of growth: ideas, facts and challenges. Development *144*, 4238-4248. 10.1242/dev.151902.
39. Carlson, J.M., and Doyle, J. (2002). Complexity and robustness. Proc Natl Acad Sci U S A *99 Suppl 1*, 2538-2545. 10.1073/pnas.012582499.
40. Martincorena, I., Roshan, A., Gerstung, M., Ellis, P., Van Loo, P., McLaren, S., Wedge, D.C., Fullam, A., Alexandrov, L.B., Tubio, J.M., et al. (2015). Tumor evolution. High burden and pervasive positive selection of somatic mutations in normal human skin. Science *348*, 880-886. 10.1126/science.aaa6806.

41. Lee-Six, H., Olafsson, S., Ellis, P., Osborne, R.J., Sanders, M.A., Moore, L., Georgakopoulos, N., Torrente, F., Noorani, A., Goddard, M., et al. (2019). The landscape of somatic mutation in normal colorectal epithelial cells. Nature *574*, 532-537. 10.1038/s41586-019-1672-7.
42. Poole, A.J., Li, Y., Kim, Y., Lin, S.C., Lee, W.H., and Lee, E.Y. (2006). Prevention of Brca1-mediated mammary tumorigenesis in mice by a progesterone antagonist. Science *314*, 1467-1470. 10.1126/science.1130471.
43. Hollern, D.P., Contreras, C.M., Dance-Barnes, S., Silva, G.O., Pfefferle, A.D., Xiong, J., Darr, D.B., Usary, J., Mott, K.R., and Perou, C.M. (2019). A mouse model featuring tissue-specific deletion of p53 and Brca1 gives rise to mammary tumors with genomic and transcriptomic similarities to human basal-like breast cancer. Breast Cancer Res Treat *174*, 143-155. 10.1007/s10549-018-5061-y.
44. Retsky, M., Demicheli, R., Hrushesky, W., Baum, M., and Gukas, I. (2010). Surgery triggers outgrowth of latent distant disease in breast cancer: an inconvenient truth? Cancers (Basel) *2*, 305-337. 10.3390/cancers2020305.
45. Retsky, M.W., and Demicheli, R. (2017). Perioperative Inflammation as Triggering Origin of Metastasis Development. In Perioperative Inflammation as Triggering Origin of Metastasis Development, M.W. Retsky, and R. Demicheli, eds. (Springer International Publishing), pp. 19-54. 10.1007/978-3-319-57943-6.
46. Dulucq, S., Nicolini, F.E., Rea, D., Cony-Makhoul, P., Charbonnier, A., Escoffre-Barbe, M., Coiteux, V., Lenain, P., Rigal-Huguet, F., Liu, J., et al. (2022). Kinetics of early and late molecular recurrences after first-line imatinib cessation in chronic myeloid leukemia: updated results from the STIM2 trial. Haematologica *107*, 2859-2869. 10.3324/haematol.2022.280811.
47. Rueda, O.M., Sammut, S.J., Seoane, J.A., Chin, S.F., Caswell-Jin, J.L., Callari, M., Batra, R., Pereira, B., Bruna, A., Ali, H.R., et al. (2019). Dynamics of breast-cancer relapse reveal late-recurring ER-positive genomic subgroups. Nature *567*, 399-404. 10.1038/s41586-019-1007-8.
48. Kay, K., Dolcy, K., Bies, R., and Shah, D.K. (2019). Estimation of Solid Tumor Doubling Times from Progression-Free Survival Plots Using a Novel Statistical Approach. AAPS J *21*, 27. 10.1208/s12248-019-0302-5.
49. Ruggeri, A., Labopin, M., Ciceri, F., Mohty, M., and Nagler, A. (2016). Definition of GvHD-free, relapse-free survival for registry-based studies: an ALWP-EBMT analysis on patients with AML in remission. Bone Marrow Transplant *51*, 610-611. 10.1038/bmt.2015.305.
50. Marconi, L., Sun, M., Beisland, C., Klatte, T., Ljungberg, B., Stewart, G.D., Dabestani, S., Choueiri, T.K., and Bex, A. (2021). Prevalence, Disease-free, and Overall Survival of Contemporary Patients With Renal Cell Carcinoma Eligible for Adjuvant Checkpoint Inhibitor Trials. Clin Genitourin Cancer *19*, e92-e99. 10.1016/j.clgc.2020.12.005.
51. Zisman, A., Pantuck, A.J., Dorey, F., Said, J.W., Shvarts, O., Quintana, D., Gitlitz, B.J., deKernion, J.B., Figlin, R.A., and Belldegrun, A.S. (2001). Improved prognostication of renal cell carcinoma using an integrated staging system. J Clin Oncol *19*, 1649-1657. 10.1200/JCO.2001.19.6.1649.
52. Brookman-Amissah, S., Kendel, F., Spivak, I., Pflanz, S., Roigas, J., Klotz, T., and May, M. (2009). Impact of clinical variables on predicting disease-free survival of patients with surgically resected renal cell carcinoma. BJU Int *103*, 1375-1380. 10.1111/j.1464-410X.2008.08233.x.
53. Tran, B.V., Moris, D., Markovic, D., Zaribafzadeh, H., Henao, R., Lai, Q., Florman, S.S., Tabrizian, P., Haydel, B., Ruiz, R.M., et al. (2023). Development and validation of a REcurrent Liver cAncer Prediction ScorE (RELAPSE) following liver transplantation in patients with hepatocellular carcinoma: Analysis of the US Multicenter HCC Transplant Consortium. Liver Transpl *29*, 683-697. 10.1097/LVT.0000000000000145.
54. Sakon, M., Umeshita, K., Nagano, H., Eguchi, H., Kishimoto, S., Miyamoto, A., Ohshima, S., Dono, K., Nakamori, S., Gotoh, M., and Monden, M. (2000). Clinical significance of hepatic resection in hepatocellular carcinoma: analysis by disease-free survival curves. Arch Surg *135*, 1456-1459. 10.1001/archsurg.135.12.1456.
55. Kornacker, M., Stumm, J., Pott, C., Dietrich, S., Sussmilch, S., Hensel, M., Nickelsen, M., Witzens-Harig, M., Kneba, M., Schmitz, N., et al. (2009). Characteristics of relapse after autologous stem-cell transplantation for follicular lymphoma: a long-term follow-up. Ann Oncol *20*, 722-728. 10.1093/annonc/mdn691.
56. Silverman, M.K., Kopf, A.W., Grin, C.M., Bart, R.S., and Levenstein, M.J. (1991). Recurrence rates of treated basal cell carcinomas. Part 1: Overview. J Dermatol Surg Oncol *17*, 713-718. 10.1111/j.1524-4725.1991.tb03424.x.
57. Maeda, R., Yoshida, J., Hishida, T., Aokage, K., Nishimura, M., Nishiwaki, Y., and Nagai, K. (2010). Late recurrence of non-small cell lung cancer more than 5 years after complete resection: incidence and clinical implications in patient follow-up. Chest *138*, 145-150. 10.1378/chest.09-2361.
58. Zalaudek, I., Horn, M., Richtig, E., Hodl, S., Kerl, H., and Smolle, J. (2003). Local recurrence in melanoma in situ: influence of sex, age, site of involvement and therapeutic modalities. Br J Dermatol *148*, 703-708. 10.1046/j.1365-2133.2003.05155.x.
59. Morton, D.L., Thompson, J.F., Cochran, A.J., Mozzillo, N., Nieweg, O.E., Roses, D.F., Hoekstra, H.J., Karakousis, C.P., Puleo, C.A., Coventry, B.J., et al. (2014). Final trial report of sentinel-node biopsy versus nodal observation in melanoma. N Engl J Med *370*, 599-609. 10.1056/NEJMoa1310460.
60. Ureshino, H. (2021). Treatment-free remission and immunity in chronic myeloid leukemia. Int J Hematol *113*, 642-647. 10.1007/s12185-021-03117-7.

61. Saussele, S., Richter, J., Guilhot, J., Gruber, F.X., Hjorth-Hansen, H., Almeida, A., Janssen, J., Mayer, J., Koskenvesa, P., Panayiotidis, P., et al. (2018). Discontinuation of tyrosine kinase inhibitor therapy in chronic myeloid leukaemia (EURO-SKI): a prespecified interim analysis of a prospective, multicentre, non-randomised, trial. The Lancet. Oncology *19*, 747-757. 10.1016/S1470-2045(18)30192-X.