Mathematical Characterization of Drug-Induced Persistence in Cancer

Ji Tae Park[1, 2] and Herbert Levine[1, 2, 3]

[1]Department of Physics, Northeastern University, Boston, MA 02115, USA
[2]Center for Theoretical Biological Physics, Northeastern University, Boston, MA 02115, USA
[3]Department of Bioengineering, Northeastern University, Boston, MA 02115, USA

Abstract

We introduce a general phenomenological framework for understanding how phenotypic plasticity gives rise to drug persisters. These persisters, often quiescent but sometimes which again return to cycling, survive in the presence of treatment and eventually can lead to mutants with true resistance. Our framework builds on recent experimental observations regarding variations between and among single-cell clones and the possible role of the drug itself in enhancing the survival strategy. Predictions of our approach include the existence of an optimum drug concentration as well as an optimum drug holiday schedule to minimize the persistence-based threat.

1 Introduction

Despite ongoing advancements, discoveries, and drug developments in oncology, many patient outcomes remain suboptimal, with refractory and recurrent diseases continuing to result in substantial morbidity and mortality worldwide [1–3]. Although a plethora of mechanisms of resistance to specific therapies have been studied and identified [4], our understanding of the fundamental principles that allow cancer cells to adapt to stressful environments, such as drug exposure, remains incomplete. Traditionally, resistance to anticancer agents has been attributed to stochastic genetic mutations that alter drug targets, create specific and/or non-specific efflux pumps, or induce other advantageous changes in biochemical pathways [5]. However, recent research has highlighted a different phenomenon known as persistence – distinct from resistance – which is believed to be a reversible, drug-tolerant state achieved through phenotypic adaptations, and typically associated with a low proliferation rate [6].

Persistence, first discovered and extensively studied in bacterial systems [7], has only recently garnered attention among cancer biologists. Unlike resistance, whether pre-existent or later acquired by stochastic genetic mutations, persistence is achieved exclusively through shifts in phenotypic states. Indeed, the phenotypic plasticity of cancer cells has long been considered one of the key driving factors behind acquired drug tolerance [8], minimal residual disease (MRD) and dormancy [9], the epithelial-mesenchymal transition (EMT) [10], and cancer stem cell (CSC) formation [11]. A growing body of experimental evidence suggests that upon exposure to a drug, a subpopulation of cancer cells can transiently enter a drug-tolerant state without any alterations in genotype but can regain sensitivity once the drug is removed [12, 13]. Furthermore, although these drug-tolerant persisters (DTPs) are primarily found to be quiescent, a small

fraction of them may re-enter the proliferative cycle, thereby forming "drug-tolerant expanding persisters (DTEP)" also known as "cycling persisters" [12, 14, 15].

There is also increasing evidence linking persistence to adverse treatment outcomes. Hata et al. reported that a subset of drug-tolerant PC9 human non-small cell lung carcinoma (NSCLC) cells subsequently underwent a T790M mutation, conferring resistance. This occurred after prolonged exposure to gefitinib, an anti-EGFR tyrosine kinase inhibitor (TKI)[16]. Similarly, Ramirez et al. demonstrated that extended treatment with erlotinib, another anti-EGFR TKI, can eventually lead to the *de novo* acquisition of resistance by persisting PC9 cells, through diverse mechanisms [17]. Furthermore, Russo et al. suggest persistent cells exhibit stress-induced mutagenesis, increasing their likelihood of acquiring resistance compared to drug-sensitive cells, based on their experiments exposing colorectal cancer cells to either dabrafenib, an anti-BRAF TKI and cetuximab, an anti-EGFR monoclonal antibody, or to cetuximab alone to induce persistence, and then employing the mammalian cell Luria-Delbruck fluctuation assay to quantify rates of resistance development [18]. Notably, their mathematical model, which represents sensitive and persistent cells as two discrete sets of population with drug concentration-dependent transition rates between them, most accurately fit the group's empirical data only after assuming that persistence is largely, if not entirely, drug-induced [18]. This finding contrasts with bacterial systems, where a small subpopulation of persistent cells is thought to be typically maintained through bet-hedging strategies [19]. Collectively, many experimental observations underscore cancer cells' phenotypic plasticity that enables extensive adaptations to various environmental stressors, such as cytotoxic drugs.

In this study, we develop a mathematical framework to describe the dynamic process by which cancer cells enter a persistent state and predict their subsequent fates under treatment. Our approach conceptualizes epigenetic and phenotypic states as one-dimensional continuous variables, and it models the population probability density of these cells using a modified Fokker-Planck equation with diffusion, advection, source (growth), and sink (death) terms dependent on this abstract (epigenetic and phenotypic) space and drug concentration, which itself can be a function of time. This effort builds on previous work by Kessler and Levine [20] who employed a phenomenological model to characterize persistence as observed in Berger et al.'s experiments on PC9 cells treated with gefitinib [21]. In their experiments, Berger et al. measured the chance-to-persist (CTP), defined as the survival fraction on day 7 of the treatment, for each single-cell clone derived from the PC9 cell line. Then, they compared the CTP distributions of naive cells (those not previously treated with gefitinib) and drug-released persisters (DRPs), which were previously exposed to gefitinib and subsequently passaged in a drug-free environment, upon re-exposure to the drug. They found that CTP is a heritable trait within clones and the population exhibits slow gradual reversion to the naive CTP distribution following drug withdrawal [21]. They also identified insulin receptor substrate 1 (IRS1) phosphorylation as a potential molecular determinant of CTP against gefitinib.

In their work, Kessler and Levine developed a model that describes the temporal progression in a one-dimensional epigenetic space regulating a clone's CTP using a partial differential equation (PDE) and the dynamics in the phenotypic space responsible for adaptive survival with the Ornstein-Uhlenbeck process [20]. While their simulations align well with experimental data collected by Berger et al., the model lacks concentration-dependent mechanisms by which drug presence specifically induces persistence. Furthermore, the model contains no mechanism allowing for the re-emergence of proliferation in the population via cycling persisters. To address these limitations, we extend the modeling approach by incorporating intuitive pharmacology-inspired terms, enabling the characterization of cellular states and dynamics across these epigenetic and phenotypic manifolds especially in the presence of a drug. Our approach predicts a critical CTP value for a given drug concentration that can serve as a reference for forecasting a persisting clone's eventual fate – gradual extinction or cycling persistence – and allows us to explore different treatment schemes. In particular, we evaluate the potential benefits of intermittent drug holidays. We emphasize that our framework is not intended to be a set of specific mathematical models designed to fit specific experimental data points. Instead, it provides a broad theoretical framework to investigate more general

and fundamental concepts and principles, such as the role of phenotypic plasticity in persistence, which arise in diverse contexts and cannot be fully captured and understood through data analysis alone. A detailed description of the mathematical model is provided in the subsequent section.

2 The Model

We assume that within a tissue or population of cancer cells sharing the same genotype, each single-cell clone exhibits a distinct probability of persisting under treatment. To quantify this, we adopt an approach inspired by Berger et al. [21] and Kessler and Levine [20], operationally defining CTP, denoted as $x \in (0, 1)$, for a given clone as the surviving fraction of the population after continuous drug treatment at a reference concentration c_p for a duration τ. In our context, we set $c_p = 1$ (in arbitrary units) and $\tau = 7$ days. We interpret x as specifying a long-lived epigenetic state that determines a clone's corresponding CTP, which can dynamically change, allowing cells to explore different epigenetic states (i.e. different values of CTP) at a rate of μ. We also assume that increasing the CTP (x) incurs a cost, which manifests as a reduction in the clone's proliferation, in the absence of any drug.

Sharing the same epigenetic state does not imply complete phenotypic homogeneity among cells. In fact, phenotypic heterogeneity, coupled with plasticity, has been shown to enhance the evolution of cancer cells, driving non-genetically acquired tolerance and eventual resistance even to targeted therapies [22]. To account for this additional variability, we introduce a second continuous variable, $s \in \mathbb{R}$, representing phenotypic states associated with survival under stress. Cells with higher s values display greater survival than those with lower s values. Each epigenetic state x determines an initial clonal population distribution in the phenotypic space s before treatment. Our model allows the treatment to induce phenotypic adaptation and dynamically alter this distribution. Effectively, tumor cells can be viewed as diffusing particles on the x-s plane under a drug-dependent potential and birth and death events. At the population level, this behavior is captured by the following partial differential equation:

$$\frac{\partial P(x,s,t)}{\partial t} = \left(\mu \frac{\partial^2}{\partial x^2} + \sigma^2 \frac{\partial^2}{\partial s^2}\right) P + b \frac{\partial}{\partial s} \left[s - s_0(x) - v(s, c(t))\right] P + \left[r(x, s, c(t)) - k(s, c(t))\right] P \qquad (1)$$

where $c(t)$ is the drug concentration as a function of time.

Partially inspired by Clairambault et al. [23], the advection term $v(s, c(t))$ captures the drug-induced adaptation by shifting cells toward more drug-tolerant states within the phenotypic space. This term highlights the paradoxical effect of treatment: while drugs induce cytotoxicity, they also promote phenotypic adaptations that enable persistence. We chose the generic forms

$$v(s, c(t)) = \frac{v_0}{1 + e^{\beta(s - s_v)}}, \qquad v_0 := \frac{v_{\max} \left[c(t)\right]^h}{c_v + \left[c(t)\right]^h}$$

Here, v_0 is the asymptotic upper limit of $v(s, c(t))$, which saturates sigmoidally toward $v_{\max}$ with increasing drug concentration, $c(t)$, with cooperativity of h. This formulation prevents the advection rate from becoming arbitrarily and unrealistically large at high drug doses. Additionally, $v(s, c(t))$ is a decreasing function of s, reflecting the assumption that cells already with sufficient fitness do not require further advection to gain a survival advantage. The positive parameters, β, s_v, and c_v, modulate the functional shape of $v(s, c(t))$.

The third term in Equation 1 encompasses the proliferation and drug-induced death terms, $r(x, s, c(t))$ and $k(s, c(t))$, respectively. The death rate is a function of the phenotypic state, s, and the drug concentration,

$c(t)$, and is chosen to be

$$k(s, c(t)) = \frac{k_0}{1 + e^{\gamma(s - s_k)}}, \qquad k_0 := \frac{k_{\max}}{1 + \frac{c_{50}}{c(t)}}$$

Analogous to the advection term, the killing rate as a function of s is bounded from above by k_0 which itself asymptotically approaches $k_{\max}$ as $c(t)$ increases. We constructed the form of k_0 to resemble a dose-response curve with c_{50} as the drug concentration at which $k_0 = \frac{1}{2}k_{\max}$. The killing rate diminishes at higher s values, approaching zero as cells become more drug-tolerant. The parameters γ and s_k are both positive constants that determine the shape and steepness of the $k(s, c(t))$.

Lastly, the proliferation term, $r(x, s, c(t))$, is influenced by both epigenetic (x) and phenotypic (s) states (Equation 2). We assume

$$r(x, s, c(t)) = \frac{R_x(x)}{1 + \frac{c(t)}{c_x}} + \frac{R_s(s)}{1 + \frac{c_s}{c(t)}} \tag{2}$$

$$R_x(x) = \xi - ax + r_x \left(1 - \frac{N(t)}{N_{ss}}\right), \qquad N(t) := \int_0^1 dx \int_{-\infty}^{\infty} ds\, P(x, s, t)$$

$$R_s(s) = \frac{r_0}{1 + e^{-\alpha(s - s_r)}}$$

The function $r(x, s, c(t))$ is a slowly increasing logistic function in s, bounded from above by r_0 and parametrized by $\alpha > 0$ and $s_r \geq 0$. This behavior captures the initial quiescence observed in persistent cells during treatment. It is important to note that quiescence is not a consequence of the balancing between birth and death terms ($r - k$), but instead is a state marked by near, if not complete, cell-cycle arrest [24]. The proliferation is also dependent on x, with clones having a baseline growth rate of ξ while incurring a cost for obtaining higher CTP modulated by $a > 0$. Moreover, the total population globally attempts to reach its carrying capacity N_{ss}. Notably, $r(x, s, c(t))$ becomes primarily determined by x when $c(t) \to 0$, while it is predominantly affected by s when $c(t) \to \infty$. With c_x and c_s sufficiently small, $r(x, s, c(t))$ then can be reasonably approximated as:

$$r(x, s, c(t)) \approx \begin{cases} R_x(x), & c(t) = 0 \\ \\ R_s(s), & c(t) \neq 0 \end{cases}$$

We assume that a sufficiently long time has elapsed prior to treatment such that the initial population distribution follows the steady-state solution of Equation 1 in the absence of drug ($c(t) = 0$). P can be decomposed into a steady-state distribution along x and a conditional probability density along s for each given x. That is, at $t = 0$, we can express $P(x, s, 0) = X(x)S(s|x)$. Under this assumption, $X(x)$ and $S(s|x)$ satisfy the following ordinary differential equations:

$$\mu \frac{d^2 X}{dx^2} + R_x(x)X = 0 \quad , \quad N(t) = N_{ss} \tag{3}$$

$$\sigma^2 \frac{d^2 S}{ds^2} + b \frac{d}{ds}\left(s - s_0(x)\right) S = 0 \tag{4}$$

The solutions to Equations 3 and 4 are an Airy distribution with $\bar{x} = \xi/a$ (Equation 5) and a Gaussian distribution centered around $s_0(x)$ (Equation 6), respectively,

$$X(x) = \mathcal{N}\mathrm{Ai}\left(\tilde{a}\left(x - \bar{x}\right)\right), \qquad \tilde{a} = \frac{a}{\mu} \tag{5}$$

$$S(s|x) = \frac{1}{\sqrt{2\pi\sigma^2/b}} e^{-\frac{b}{2\sigma^2}(s - s_0(x))^2} \tag{6}$$

where $\mathcal{N}$ is a scaling constant such that $\int_0^1 dx X(x) = N_{ss}$.

It is worth noting that the motion of these tumor cell "particles" at the two extremes of the epigenetic space ($x = 0$ and $x = 1$) is unclear. Our operational definition of CTP does not allow x to be defined outside of the interval $(0, 1)$ and the exact characterization may vary depending on the chosen boundary conditions. However, we expect that Equation 5 captures behavior away from boundaries reasonably well. To confirm our intuition, we also performed a stochastic simulation, which revealed that the normalized distribution of $X(x)$ closely follows the sharp Airy-like profile with a peak at some low x value (Appendix A), consistent with Equation 5. Meanwhile, $s_0(x)$ is self-consistently determined from x in accordance with the operational definition of CTP

$$x := \int_{-\infty}^{\infty} ds P(x, s, \tau), \qquad (c(t) = c_p) \tag{7}$$

Based on Equations 6 and 7, having a higher x value bestows a single-cell clone with the initial distribution in s with a higher mean phenotype (Figure 1).

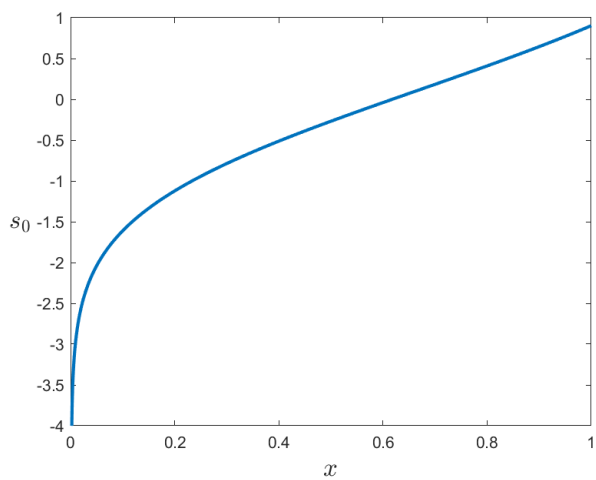

Figure 1: The plot of s_0 as a function of x. The initial distribution in s for a clone with a fixed x is precisely determined by $s_0(x)$.

Although the population distribution, $P(x, s, t)$, across epigenetic and phenotypic spaces can, in principle, evolve in both x and s, we assume that changes in s occur on a much shorter timescale than changes in x. The assumption is justified as phenotypic transitions typically occur rapidly whereas epigenetic changes are inherently slower. To reflect this, we impose the condition, $\mu \ll \sigma^2$, ensuring that the s-dynamics dominate in the presence of a drug ($c(t) > 0$) while the x-dynamics play a more prominent role in the absence of a drug ($c(t) = 0$). Under these assumptions and for cases where $c(t)$ is either some positive constant or zero (e.g. similar to an *in vitro* experiment), Equation 1 can be approximated as:

$$\partial_t P(x, s, t) \approx \begin{cases} \sigma^2 \partial_s^2 P + b\partial_s \left(s - s_0(x) - v(s, c)\right) P + \left[R_s(s) - k(s, c)\right] P, & c(t) > 0 \\ \\ \left(\mu \partial_x^2 + \sigma^2 \partial_s^2\right) P + b\partial_s \left(s - s_0(x)\right) P + R_x(x) P, & c(t) = 0 \end{cases} \tag{8}$$

When $c(t) = 0$, the solution of Equation 8 will relax to the initial conditions, i.e. the Airy distribution in x and Gaussian distribution in s, as described in Equations 5 and 6, respectively. When the drug is present ($c(t) > 0$), the advection rate $v(s)$, birth rate $R_s(s)$, and death rate $k(s)$ depend solely on the

phenotypic state s. As previously alluded, $R_s(s)$ was chosen as a slowly increasing monotonic function. Consequently, in the presence of a drug, the population initially declines, but the probability density shifts toward higher s values, enabling the cells to acquire greater fitness and enhanced tolerance to the treatment. If a subpopulation starts with a very high $s_0(x)$ and/or is advected sufficiently further in the positive s-direction, these cells have the potential to become cycling persisters.

3 Continuous Treatment

Using Equation 8, we simulated various treatment schemes to evaluate their efficacy and predict outcomes. For a fixed set of parameters (Appendix B), we first performed a 90-day continuous treatment simulation with $c(t) = 1$ (in arbitrary units) and generated survival curves $A(x, t)$,

$$A(x,t) = \int_{-\infty}^{\infty} ds\, P(x,s,t) \tag{9}$$

for fixed values of x (Figure 2(a)). We then also generated survival curves for the entire population $N(t)$ at different constant drug concentrations (ranging from 10^{-4} to 10^{1}) with an assumed initial distribution given by Equations 5 and 6 (Figure 2(b)).

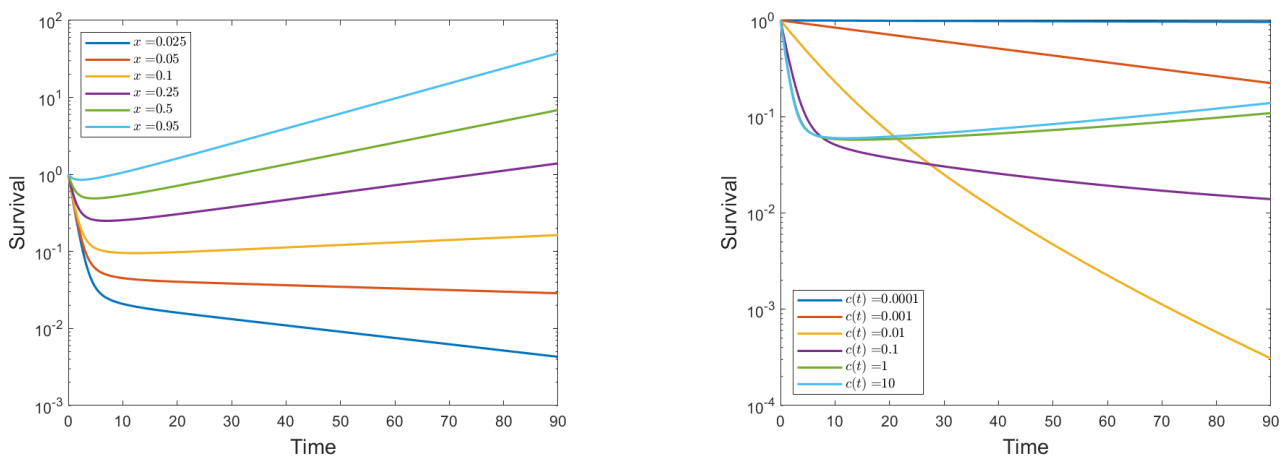

(a) (b)

Figure 2: (a) Survival curves for single-cell clones with different x values (CTP) throughout the treatment with $c(t) = 1$. (b) The total population curves at various drug concentrations ($c(t)$) assuming the initial distribution in x and s is given by Equations 5 and 6.

The first set of simulations with $c(t) = 1$ (Figure 2(a)) reveals that subpopulations with different CTPs asymptotically converge to one of two outcomes: gradual extinction or cycling persistence (Figure 2(a)). When the size of each x-clone (i.e. a single-cell clone with a given x) is weighted by the initial CTP distribution (Equation 5), the total population initially decrease sharply during treatment. However, the population ultimately begins to regrow, albeit at a relatively slow rate (2(b)). This behavior reflects the near extinction of low-CTP cells, which initially dominated the population distribution, followed by the emergence of cycling persistence among survivors through advection in the s-phenotypic space. Increasing the drug concentration enhances the initial killing rate, but also promotes extensive phenotypic adaptations, potentially leading to cycling persistence (2(b)). Conversely, reducing $c(t)$ excessively can result in insufficient killing and suboptimal long-term outcomes (2(b)). It is important to note that during most *in*

vitro drug-dose response experiments, viability is usually measured 24-72 hours after the initiation of the treatment. As shown in Figure 2(a), cells with low x values are predestined for extinction, whereas cells with sufficiently high x values may regain the ability to cycle. This suggests the existence of a critical CTP value, x_c for a given constant drug concentration, such that a solution to $\frac{dA(x,t)}{dt} = 0$ exists for $x > x_c$ but not for $x < x_c$. With the parameters chosen for the simulation (Appendix B), we numerically estimated $x_c \approx 0.071$ for $c(t) = 1$.

Under stronger assumptions, x_c can be approximated through a more analytical approach. Suppose the advection rate can be reasonably approximated by a constant value v_c, and the difference between growth and death rate is small. That is,

$$v(s) \approx v_c, \qquad R_s(s) - k(s) \approx 0$$

Of course, this approximation breaks down in a short timescale or where $|r(s) - k(s)|$ becomes significant. However, at near steady-state regimes, where the population experiences minimal change, we believe these assumptions are qualitatively valid. Under these conditions and with $c(t) = 1$, Equation 8 simplifies to

$$\partial_t P = \sigma^2 \partial_s P + b\partial_s(s - s_0(x) - v_c)P \tag{10}$$

From Equation 9 and the $c(t) > 0$ case of Equation 8, one can easily derive

$$\frac{dA(x,t)}{dt} = \int_{-\infty}^{\infty} ds\ [R_s(s) - k(s)]\, P(x,s,t) \tag{11}$$

Meanwhile, Equation 10 can be solved by integrating the product of its initial steady-state solution, $P^{ss}(s|x)$, and the Green's function $G(s, s', t, t'|x)$ for the Ornstein-Uhlenbeck process given by:

$$G(s,s',t,t'|x) = \sqrt{\frac{b}{2\pi\sigma^2\left(1 - e^{-2b(t-t')}\right)}}\, e^{-\frac{b}{2\sigma^2}\frac{\left[s - s_0(x) - v_c - (s' - s_0(x) - v_c)e^{-b(t-t')}\right]^2}{1 - e^{-2b(t-t')}}}$$

Thus, the full solution for $P(x,s,t)$ is:

$$\begin{aligned} P(x,s,t) &= \int_{-\infty}^{\infty} ds' G(s,s',t,0|x) P^{ss}(s'|x) \\ &= \sqrt{\frac{b}{2\pi\sigma^2}}\, e^{-\frac{b}{2\sigma^2}\left[(s - s_0(x) - v_c)e^{bt} + v_c\right]^2} e^{-2bt} \end{aligned} \tag{12}$$

By substituting Equation 12 into Equation 11, we can obtain an explicit expression for $\frac{dA}{dt}$. This allows us to plot the curve and determine the critical value x_c, where a root exists for $x > x_c$ but does not for $x < x_c$ (Figure 3(a)). However, it suffices to focus on the asymptotic behavior at steady-state. At $t \to \infty$, the existence of a root for Equation 11 is determined entirely by $s_0(x)$ (Equation 13).

$$\lim_{t\to\infty} \frac{dA(x,t)}{dt} = \sqrt{\frac{b}{2\pi\sigma^2}} \int_{-\infty}^{\infty} ds\ [R_s(s) - k(s)]\, e^{-\frac{b}{2\sigma^2}(s - s_0(x) - v_c)^2} \tag{13}$$

For a given constant $c(t) = c_0$, we chose $v_c = v(s_{rk}, c_0)$ where $R_s(s_{rk}) - k(s_{rk}, c_0) = 0$. Since $\frac{dA}{dt} < 0$ at small t for all $x < 1$, the critical value x_c precisely corresponds to the root of Equation 13. For example, with $c(t) = 1$, we calculated $v_c = 3.56$ which yields $x_c = 0.11$ (Figure 3(b)). This value is close to the critical value determined by manually solving the fully time-dependent Equation 8 ($x_c = 0.071$). While this approximation method does not yield exact values of x_c for every c_0, it still qualitatively replicates the correct behavior (Figure 4).

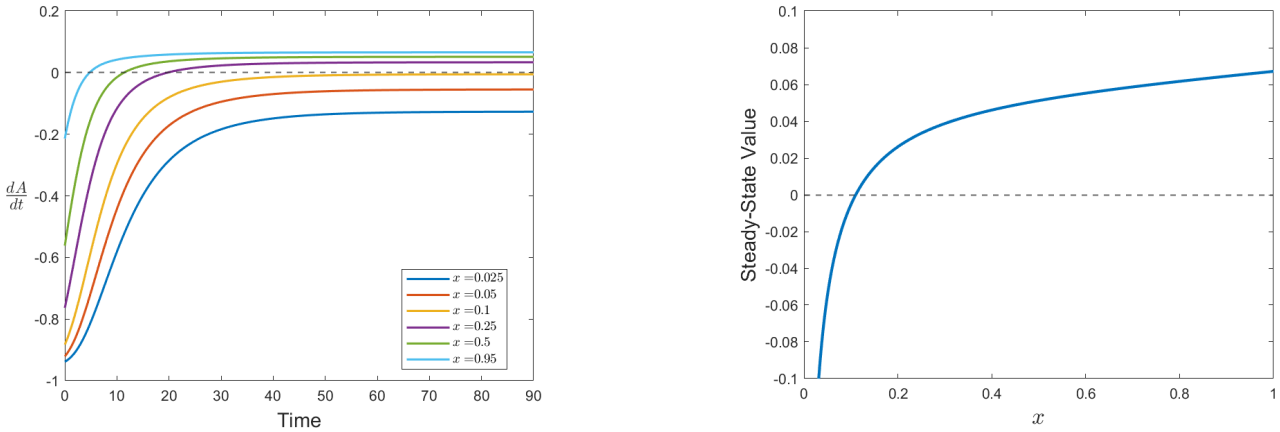

(a) (b)

Figure 3: Plots of (a) $\frac{dA}{dt}$ for different values of x (CTP) and (b) the asymptotic $\frac{dA}{dt}$ (Equation 13) as a function of x for $c(t) = 1$.

Determining x_c allows us to predict the fate of single-cell clones with a given CTP, analogous to the experimental work by Oren et al., where different lineages of PC9 cell lines were treated with osimertinib, an EGFR TKI, for 14 days and subsequently classified as "drug sensitive", "persisters with small progeny", and "persisters colonies" [15]. In our context, "persisters with small progeny" refer to DTPs with minimal or no cell division, while "persisters colonies" denote cycling persisters. Therefore, the fraction of the population likely to attain cycling persistence can be estimated using an initial CTP distribution. Finding x_c can also aid in optimizing therapeutic doses. Selecting a dose with high x_c reduces the likelihood of cycling persistence while enhancing cell killing. Our method suggests that the optimal constant dose yielding the maximum x_c is $c_0 = 6.5 \times 10^{-3}$, while solving the PDE yields $c_0 = 4.9 \times 10^{-3}$ (Figure 4). This condition is consistent with the data presented in Figure 2(b), where the concentration that is most effective is seen to be around 0.01.

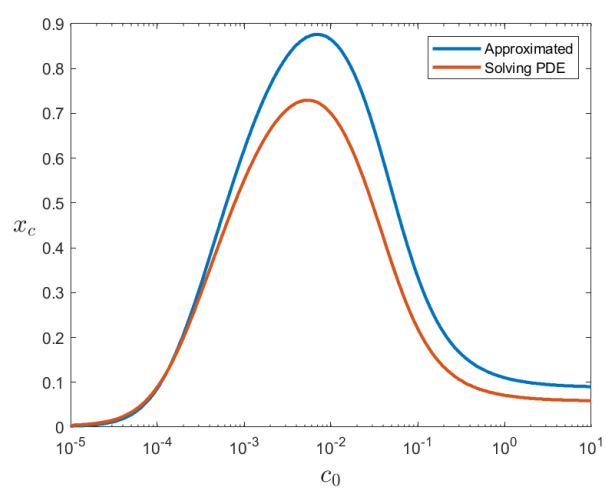

Figure 4: Values of critical CTP (x_c) calculated for different $c(t) = c_0$ using our approximation method and by solving the PDE (Equation 8).

4 Drug Holidays

A potential weakness of continuous therapy as demonstrated by previous studies [25, 26] and real-world clinical data [27, 28] as well as our simulations above is the induction of DTPs and cycling persisters, which may acquire resistance after a prolonged survival and ultimately lead to treatment failure. Clinicians have long observed the emergence of acquired persistence and resistance due to inadvertent adaptation mechanisms and the unintended selection pressures caused by cytotoxic therapeutics [29, 30]. In response, concepts such as drug holidays [24, 31, 32] and adaptive therapy [33] are gaining traction and being increasingly explored in clinical settings. The rationale behind drug holidays is that they allow more sensitive cells to repopulate and occupy a larger fraction of the tumor population during periods of drug withdrawal, instead of continuously administering maximally tolerated dose (MTD) or other high doses to maximize short-term cell death. Promising results from early clinical trials, particularly in prostate cancer, suggest that adaptive therapies with intermittent drug holidays can delay disease progression more effectively than continuous treatment [34]. Mathematical models have also been developed to predict the potential benefit of adaptive therapy [35]. However, most of these models consider two discrete states – sensitive and resistant – without accounting for the continuous phenotypic transitions observed in persistence.

Using our model (Equation 8), we simulated various treatment regimens incorporating drug holidays over a total duration of 90 days to investigate their impact on phenotypic adaptations and overall cell survival. These results were compared with continuous constant-dose administration at $c(t) = c_0$. For all simulations, the total drug exposure was constrained to match the area under the curve (AUC) of the respective continuous treatment, $c_0 T$, where $T = 90$ days represents the total therapy duration. Under a regimen with n isochronal drug holidays between $t = 0$ and $t = T$ (starting and ending with drug administration), the dose during the treatment phase then is $\frac{2n+1}{n+1} c_0$, and $c(t)$ is given by

$$c(t) = \begin{cases} c_0, & \text{Continuous} \\ \dfrac{2n+1}{n+1} c_0 \displaystyle\sum_{i=0}^{n} \Theta\left(t - \frac{2iT}{2n+1}\right) \Theta\left(-t + \frac{(2i+1)T}{2n+1}\right), & \text{With } n \text{ holidays} \end{cases} \tag{14}$$

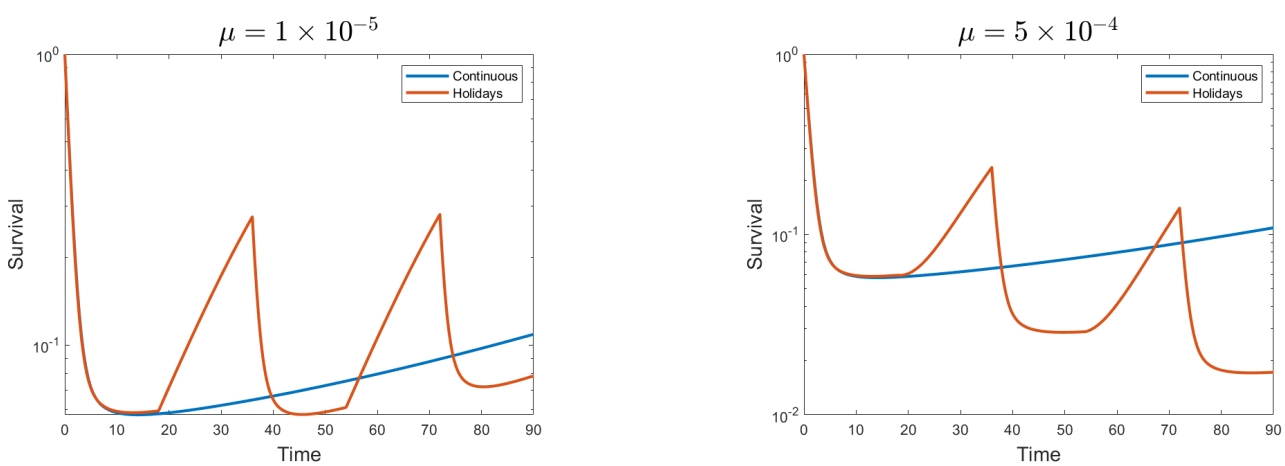

(a) (b)

Figure 5: The peaks and troughs of the total tumor population during treatment with two drug holidays compared to the curve during continuous therapy with $c_0 = 1$ and (a) $\mu = 1 \times 10^{-5}$ and (b) $\mu = 5 \times 10^{-4}$

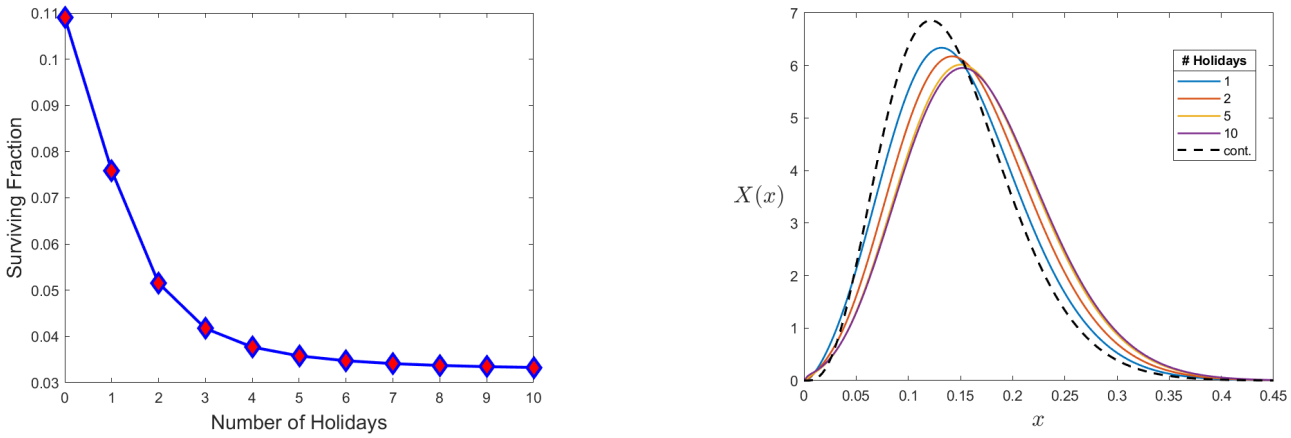

(a) (b)

Figure 6: (a) The overall surviving fractions of tumor population and (b) the CTP distributions ($X(x)$) after $T = 90$ days of treatment with a fixed total dosage but different numbers of holidays.

During the "drug-on" period ($c(t) > 0$), cells adapt to drug-induced stress, primarily through s-dynamics. In contrast, during the drug holiday phase ($c(t) = 0$), the population seeks to relax toward its initial steady-state distributions in both s (Equation 6) and x (Equation 5). These relaxation dynamics reduce fitness in the phenotypic space by reversing drug-induced advection in s and increasing the prevalence of cells with lower CTP (x). For simplicity, the drug-on and holiday durations were kept equal, implying higher drug-on doses compared to the continuous regimen to achieve equivalent total exposure (Equation 14).

Simulations revealed that the efficacy of drug holidays heavily depended on the rate of epigenetic diffusion, μ. Higher values of μ enable the population to recover low-CTP cells more effectively during holiday periods, leading to greater tumor reduction when the drug is reapplied. For instance, with two holidays and $c_0 = 1$, higher μ values consistently resulted in better outcomes, as shown in Figure 5. Conversely, lower μ values slowed the relaxation toward the initial distribution in x, diminishing the benefits of drug holidays. These results suggest that drug holidays may be particularly effective against tumors with more "malleable" epigenetic landscapes.

Additional simulations examined the effect of varying the number of holidays (n) while keeping total drug exposure fixed. In these simulations, μ and c_0 were kept at 3.6×10^{-5} and 1, respectively. The results interestingly revealed that arbitrarily increasing the number of holidays eventually produces diminishing returns, as shown in Figure 6(a). With more holidays, the drug-on doses are higher (to preserve AUC equivalence), intensifying selection pressure and phenotypic adaptations. At the same time, shorter holiday intervals provide insufficient time for the population to revert to initial distributions. In fact, even a single holiday is sufficient to shift the CTP distribution, $X(x)$, further to the right compared to the continuous treatment (Figure 6(b)). Of course, implementing excessively many holidays within a 90-day treatment period is neither feasible nor practical in clinical settings, but the simulations provide mathematical confirmation of our intuitive expectations. Overall, these findings underscore the importance of carefully balancing drug-on periods, which exert cytotoxicity but also drive adaptation and selection, with holiday durations sufficient to allow the repopulation of less fit and low-CTP cells.

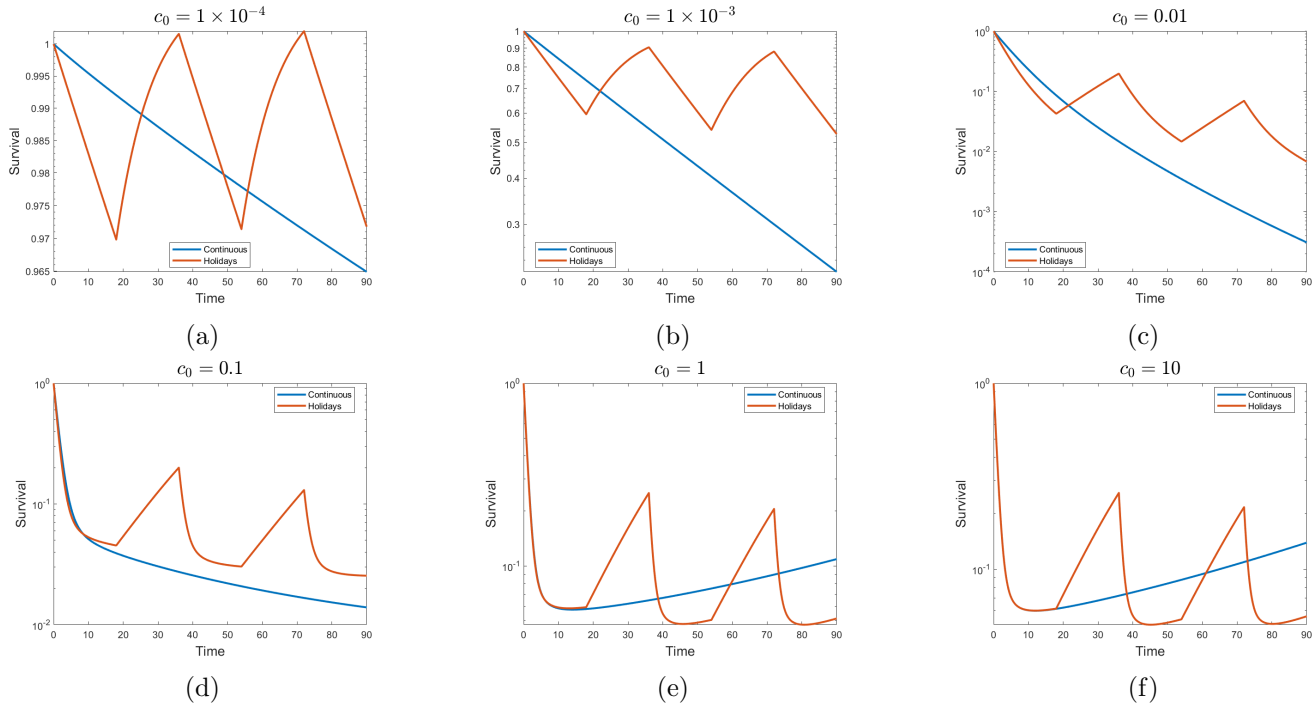

Figure 7: Simulated total population curves for both continuous and two-holiday therapies with (a) $c_0 = 0.0001$, (b) $c_0 = 0.001$, (c) $c_0 = 0.01$, (d) $c_0 = 0.1$, (e) $c_0 = 1$, (f) $c_0 = 10$

In addition to holiday frequency, the outcomes also depended on the drug dose (c_0). Simulations with two holidays ($n = 2$) and varying c_0 revealed a nuanced relationship between dose and the efficacy of holidays, as illustrated in Figure 7. At low doses, drug holidays were ineffective and offered no real benefit compared to continuous treatment (Figures 7(a)–(d)). However, at sufficiently high doses capable of inducing significant drug tolerance over longer timescales, the two drug holidays improved the treatment outcome at day 90 (Figures 7(e)–(f)). Together with other simulations mentioned earlier, this finding highlights the complex interplay between tumor plasticity, treatment schedules, and drug dosing. In particular, optimizing drug holiday regimens requires consideration of not only total drug exposure and timing but also the pharmacokinetics/pharmacodynamics (PK/PD) of the specific drug and the adaptive capacity of the tumor population.

5 Discussion

Drug persistence, as distinct from resistance, is receiving growing attention among clinicians and cancer biologists. Previous efforts to experimentally probe and theoretically model the phenotypic plasticity of tumor cells, which leads to the emergence of DTPs and cycling persisters, confirm the reversibility of the phenomenon, albeit with a likely significant time lag [12, 18, 21]. These studies suggest that persistence is likely to be drug-induced at least in cancer systems. Motivated by these findings, we developed a mathematical model to describe cancer cells as a population distributed across a probability density in an abstract space representing epigenetic (x) and phenotypic (s) degrees of freedom. This model characterizes how drug exposure enables certain cells to withstand cytotoxicity, enter a quiescent state, and even restore the ability to cycle and proliferate. In our theoretical framework, the presence of a drug drives the cells to advect in s, gaining fitness and tolerance to the drug's pharmacological effect. Further advection to higher s may lead to cycling persistence. Conversely, in the absence of the drug, the population gradually regains

drug sensitivity by reversing the advection in s and approaching the initial steady-state distribution in x.

Overall, our model captures several key characteristics and properties of persistence. First, it does not presume the existence of persistent cells prior to treatment. However, different cells have varying probabilities of becoming persistent, represented by the variable x (CTP), which we consider epigenetically determined. The fate of these cells depends on their initial positions in s, advection rate ($v(s, c(t))$), and the net proliferation/death rate ($r(x, s, c(t)) - k(s, c(t))$). Second, our model accounts for the reversibility and allows cells to reclaim the sensitivity once the drug is removed. This reversal occurs through two mechanisms: (a) the cessation of advection, effectively removing v from Equation 8, and (b) the gradual return of the x-distribution to its steady-state Airy form (Equation 5). However, this process can be slow, especially for the x-dynamics, if μ is small, making it difficult to observe substantial changes unless drug holidays are sufficiently long. Lastly, the model takes into account the possible emergence of cycling persistence. Importantly, the initial quiescence of persistent cells does not result solely from the balance between birth and death rates. Experimental evidence confirms the existence of states resembling actual cell-cycling arrests among DTPs [24]. Our model enables a subset of DTPs to transition into cycling persisters by advecting to higher s values, where $R_s(s)$, initially nearly zero, becomes sufficiently large to support net positive growth.

Simulations of various treatment schedules elucidated several interesting and critical factors to consider when designing and optimizing therapeutic interventions. Under continuous treatment with a constant concentration ($c(t) = c_0$), our model predicts the existence of a critical CTP, x_c, that can be used to determine the fates of single-cell clones – cycling persistence or eventual extinction. We demonstrated that our semi-analytical method for approximating x_c produces results that qualitatively align with those obtained by solving the fully time-dependent PDE (Equation 8). Moreover, calculating the values of x_c for different c_0 enabled us to identify the optimal continuous dose that maximizes the CTP threshold for cycling persistence. We also showed how drug holidays could mitigate the induction of (cycling) persistence by allowing cells to revert toward $s_0(x)$ in s and repopulating cells with low CTP. However, several crucial caveats must be noted. The efficacy of holidays was significantly enhanced in populations with higher μ, suggesting that epigenetic "malleability" plays a critical role in determining treatment outcomes under intermittent regimens. Furthermore, arbitrarily increasing the number of holidays resulted in diminishing returns, accompanied by a shift in $X(x)$ at the treatment endpoint toward higher x values. Lastly, the benefits of drug holidays were only apparent at sufficiently high baseline doses, emphasizing the need for dose optimization.

Several key terms in Equation 1 depend on drug concentration $c(t)$, as the model characterizes persistence as a phenotypic adaptation to the drug. Surprisingly, there is limited experimental and theoretical literature that frames persistence as a continuous spectrum of phenotypic traits or explores how varying drug concentrations and administration schedules influence the induction of these traits. While higher levels of stress often jeopardize cell survival, they can also trigger and activate more extensive adaptation mechanisms. Thus, increasing the drug dose enhances the killing rate but may simultaneously induce a greater number of DTPs and cycling persisters over longer timescales. Finding an optimal balance between eliminating sensitive cells and minimizing the generation of persisters is crucial for designing treatment algorithms that maximize therapeutic efficacy and reduce the risk of limited response. Additionally, adverse effects and potential toxicities must also be considered especially in *in vivo* and clinical settings. Combining our theoretical work with available PK/PD data has the potential to facilitate highly efficacious and safe therapy optimization.

Our model, however, operates under several assumptions and has certain limitations. For example, our phenotypic space governs both cell survivability and proliferation, as r and k are functions of s. As a result, DTPs are treated as a required precursor state before cells can transition into cycling persistence. This approach differs from the model proposed by Clairambault et al., which separates survivability and

the ability to proliferate into two distinct phenotypic dimensions, thereby allowing direct transitions from drug-sensitive states to cycling persistence [23]. Moreover, the physiological mechanisms underlying our simulation results derived from an abstract mathematical model remain uncertain. We did aim to treat persistence as a pharmacologically driven and drug concentration-dependent process, using functional forms grounded and commonly applied in biological studies. As noted earlier, the goal of our work is to convince cancer biologists that it is possible to conceptualize persistence at a theoretical level and that many cleverly designed experiments are warranted not only to refine our model but also to obtain a deeper understanding of stress-mediated phenotypic adaptations at a cellular or molecular level. In fact, growing evidence suggests that alterations in tumor metabolism may be intimately linked to cell survival under drug treatment [15, 36]. As a result, we plan to investigate and identify relevant biological and metabolic pathways affected by the PK/PD of the drug of interest, and perturb them in simulations to validate the behaviors predicted by our current model.

In summary, we present a mathematical model that captures the induction of DTPs and cycling persisters in a drug concentration-dependent manner, aligning with various trends and properties in experimental and clinical data. As discussed, mapping relevant biochemical and metabolic processes to the outcomes predicted by the model would further confirm our theory. Identifying specific molecules or physiological subnetworks could uncover viable therapeutic targets to prevent or mitigate persistence. Efforts like ours could significantly advance the understanding of persistence and acquired resistance, helping clinicians design optimal personalized treatments for cancer patients. Moreover, we believe that our mathematical approach can be generalized to characterize not only how tumor cells survive treatment but also how other organisms adapt to stimuli and stress beyond cytotoxic agents. It is well-documented that bacteria and yeasts, for example, exhibit remarkable cellular plasticity under diverse forms of stress, with these adaptations often enhancing their ability to respond to entirely different challenges, including even the development of persistence against antimicrobial treatments [19, 37–39]. Our framework may offer a promising tool for investigating the interplay between adaptation mechanisms for various stress factors and how they are coupled to collectively endow cells with survival advantages under different conditions. Ultimately, we aim to continue developing robust theoretical models that predict tumor responses to treatment and propose strategies to combat cancer cells by disrupting their adaptability under drug exposure.

Authors' Disclosures

The authors declare no conflict of interest and report no disclosures.

Authors' Contributions

The authors have equally contributed to this work.

Acknowledgements

The authors acknowledge the support of the NSF through the Center for Theoretical Biological Physics, grant no. PHY-2019745.

Appendix A Verification of Airy Distribution in $X(x)$

Suppose a sufficient amount of time had passed to allow $P(x, s, t)$ to reach its steady-state before treatment. In the absence of a drug, the cells occupy the probability density in s given by Equation 6. Under these conditions, we focus exclusively on the x-dynamics described by:

$$\partial_t P(x,t) = \mu \partial_x^2 P + \left[\xi - ax + r_x \left(1 - \frac{N(t)}{N_{ss}} \right) \right] P \tag{A1}$$

where

$$N(t) = \int_0^1 dx\, P(x,t)$$

Equation A1 can be rewritten in terms of the normalized probability density $p(x,t) := \frac{P(x,t)}{N(t)}$.

$$\partial_t p(x,t) = \mu \partial_x^2 p - a\left(x - \langle x \rangle(t)\right) p \tag{A2}$$

where

$$\langle x \rangle(t) = \int_0^1 dx\, x p(x,t) = \frac{1}{N(t)} \int_0^1 dx\, x P(x,t)$$

The steady-state solution of Equation A2 is indeed an Airy distribution identical to Equation 5. However, the Airy function is defined on all reals whereas our operational definition of CTP confines x to the interval $(0, 1)$. Also, the dynamics near the boundaries $x = 0$ and $x = 1$ on this abstract epigenetic space is ambiguous as it is biologically unclear whether the motion of a cell's state is absorbed into or reflected from these limits. Furthermore, note that $s_0 \to -\infty$ as $x \to 0$. Hence, to validate our assumption about the initial condition of $P(x, s, t)$ prior to treatment, we performed a stochastic simulation using the Gillespie algorithm.

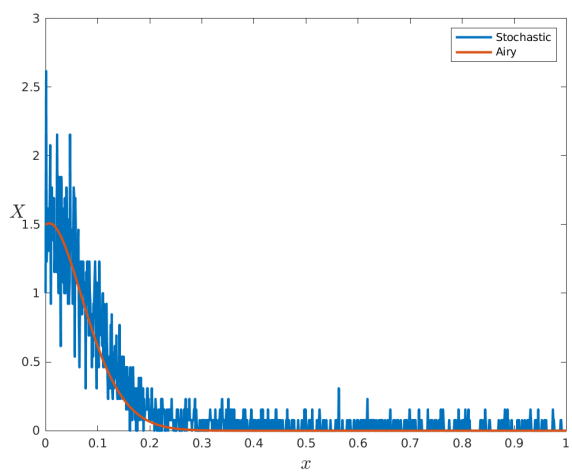

Figure 8: Stochastic simulation of the x-distribution compared with the analytical steady-state solution of Equation 3

The detailed scheme of the algorithm is as follows. First, the epigenetic space $x \in (0, 1)$ is discretized into units of Δx. Starting with an initial population $N(0) = N_0$, where all cells occupy the same state ($x = x_0$), the simulation is run for a sufficiently long duration $\mathcal{T}$, to allow the distribution to approach a near steady-state. During the simulation, each cell can move up or down in x by Δx at a rate of $\frac{\mu}{2(\Delta x)^2}$ while dying at a rate of ax. At the boundaries ($x = 0$ and $x = 1$), cells are constrained to move up or down,

respectively. Because we are interested in the normalized distribution, we maintain a constant population size by placing a new cell at a random value of x whenever a death event occurs. Using $N_0 = 2500$, $x_0 = 0.1$, $\mathcal{T} = 4000$, and $\Delta x = 0.001$ (with μ and a values consistent with Appendix B), we successfully replicated a sharp Airy-like distribution in x, closely resembling the curve generated from Equation 5 (Figure 8).

Appendix B Values of Parameters Used

The table below lists the parameter values used in the simulations, unless otherwise specified in the text.

Parameter	Value
μ	3.6×10^{-5}
σ	0.3
b	0.09
$v_{\max}$	10
h	1
c_v	0.05
β	0.5
s_v	0
$k_{\max}$	1
c_{50}	0.05

Parameter	Value
γ	4
s_k	0
r_0	0.08
α	1
s_r	2.5
ξ	0.012
r_x	0.1
a	0.18
c_s	≈ 0
c_x	≈ 0

References

[1] F. Bray, M. Laversanne, H. Sung, J. Ferlay, R. L. Siegel, I. Soerjomataram, and A. Jemal, CA: A Cancer Journal for Clinicians **74**, 229 (2024).

[2] D. A. Mahvi, R. Liu, M. W. Grinstaff, Y. L. Colson, and C. P. Raut, CA: A Cancer Journal for Clinicians **68**, 488 (2018).

[3] E. B. Maldonado, S. Parsons, E. Y. Chen, A. Haslam, and V. Prasad, Future Science OA **6**, `10.2144/fsoa-2020-0024` (2020).

[4] S. U. Khan, K. Fatima, S. Aisha, and F. Malik, Cell Communication and Signaling **22**, 1 (2024).

[5] N. Vasan, J. Baselga, and D. M. Hyman, Nature **575**, 299 (2019).

[6] S. Shen, S. Vagner, and C. Robert, Cell **183**, 860 (2020).

[7] G. L. Hobby, K. Meyer, and E. Chaffee, Proceedings of the Society for Experimental Biology and Medicine **50**, 281 (1942).

[8] J. He, Z. Qiu, J. Fan, X. Xie, Q. Sheng, and X. Sui, Signal transduction and targeted therapy **9**, 209 (2024).

[9] E. Dalla, A. Sreekumar, J. A. Aguirre-Ghiso, and L. A. Chodosh, Cold Spring Harbor Perspectives in Medicine **13**, 1 (2023).

[10] C. L. Chaffer, B. P. San Juan, E. Lim, and R. A. Weinberg, Cancer and Metastasis Reviews **35**, 645 (2016).

[11] A. P. Thankamony, K. Saxena, R. Murali, M. K. Jolly, and R. Nair, Frontiers in Molecular Biosciences **7**, 1 (2020).

[12] S. V. Sharma, D. Y. Lee, B. Li, M. P. Quinlan, F. Takahashi, S. Maheswaran, U. McDermott, N. Azizian, L. Zou, M. A. Fischbach, K. K. Wong, K. Brandstetter, B. Wittner, S. Ramaswamy, M. Classon, and J. Settleman, Cell **141**, 69 (2010).

[13] Y. Goyal, G. T. Busch, M. Pillai, J. Li, R. H. Boe, E. I. Grody, M. Chelvanambi, I. P. Dardani, B. Emert, N. Bodkin, J. Braun, D. Fingerman, A. Kaur, N. Jain, P. T. Ravindran, I. A. Mellis, K. Kiani, G. M. Alicea, M. E. Fane, S. S. Ahmed, H. Li, Y. Chen, C. Chai, J. Kaster, R. G. Witt, R. Lazcano, D. R. Ingram, S. B. Johnson, K. Wani, M. C. Dunagin, A. J. Lazar, A. T. Weeraratna, J. A. Wargo, M. Herlyn, and A. Raj, Nature **620**, 651 (2023).

[14] R. H. Chisholm, T. Lorenzi, and J. Clairambault, Biochimica et Biophysica Acta - General Subjects **1860**, 2627 (2016).

[15] Y. Oren, M. Tsabar, M. S. Cuoco, L. Amir-Zilberstein, H. F. Cabanos, J. C. Hütter, B. Hu, P. I. Thakore, M. Tabaka, C. P. Fulco, W. Colgan, B. M. Cuevas, S. A. Hurvitz, D. J. Slamon, A. Deik, K. A. Pierce, C. Clish, A. N. Hata, E. Zaganjor, G. Lahav, K. Politi, J. S. Brugge, and A. Regev, Nature **596**, 576 (2021).

[16] A. N. Hata, M. J. Niederst, H. L. Archibald, M. Gomez-Caraballo, F. M. Siddiqui, H. E. Mulvey, Y. E. Maruvka, F. Ji, H.-e. C. Bhang, V. K. Radhakrishna, G. Siravegna, H. Hu, S. Raoof, E. Lockerman, A. Kalsy, D. Lee, C. L. Keating, D. A. Ruddy, L. J. Damon, A. S. Crystal, C. Costa, Z. Piotrowska, A. Bardelli, A. J. Iafrate, R. I. Sadreyev, F. Stegmeier, G. Getz, L. V. Sequist, A. C. Faber, and J. A. Engelman, Nature Medicine **22**, 262 (2016).

[17] M. Ramirez, S. Rajaram, R. J. Steininger, D. Osipchuk, M. A. Roth, L. S. Morinishi, L. Evans, W. Ji, C. H. Hsu, K. Thurley, S. Wei, A. Zhou, P. R. Koduru, B. A. Posner, L. F. Wu, and S. J. Altschuler, Nature Communications **7**, 1 (2016).

[18] M. Russo, S. Pompei, A. Sogari, M. Corigliano, G. Crisafulli, A. Puliafito, S. Lamba, J. Erriquez, A. Bertotti, M. Gherardi, F. Di Nicolantonio, A. Bardelli, and M. Cosentino Lagomarsino, Nature Genetics **54**, 976 (2022).

[19] L. P. Morawska, J. A. Hernandez-Valdes, and O. P. Kuipers, WIREs Mechanisms of Disease **14**, 1 (2022).

[20] D. A. Kessler and H. Levine, Physical Review Letters **129**, 108101 (2022).

[21] J. A. Berger, E. Gigi, L. Kupershmidt, Z. Meir, N. Gavert, Y. Zwang, A. Prior, S. Gilad, U. Harush, I. Haviv, S. M. Stemmer, G. Blum, E. Merquiol, M. Mardamshina, S. K. Strauss, G. Friedlander, J. Bar, I. Kamer, Y. Reizel, T. Geiger, Y. Pilpel, Y. Levin, A. Tanay, B. Barzel, H. Reuveni, and R. Straussman, Nature Cancer **2**, 1055 (2021).

[22] Y. Su, M. Bintz, Y. Yang, L. Robert, A. H. Ng, V. Liud, A. Ribas, J. R. Heath, and W. Wei, PLoS Computational Biology **15**, 1 (2019).

[23] R. H. Chisholm, T. Lorenzi, A. Lorz, A. K. Larsen, L. N. De Almeida, A. Escargueil, and J. Clairambault, Cancer Research **75**, 930 (2015).

[24] Y. Pu, L. Li, H. Peng, L. Liu, D. Heymann, C. Robert, F. Vallette, and S. Shen, Nature Reviews Clinical Oncology **20**, 799 (2023).

[25] A. Nam, A. Mohanty, S. Bhattacharya, S. Kotnala, S. Achuthan, K. Hari, S. Srivastava, L. Guo, A. Nathan, R. Chatterjee, M. Jain, M. W. Nasser, S. K. Batra, G. Rangarajan, E. Massarelli, H. Levine, M. K. Jolly, P. Kulkarni, and R. Salgia, Biomolecules **12**, 1 (2022).

[26] T. Kuosmanen, J. Cairns, R. Noble, N. Beerenwinkel, T. Mononen, and V. Mustonen, PLoS Computational Biology **17**, 1 (2021).

[27] G. Tonini, M. Imperatori, B. Vincenzi, A. M. Frezza, and D. Santini, Journal of Experimental and Clinical Cancer Research **32**, 1 (2013).

[28] S. S. Skanland and A. R. Mato, Blood Advances **5**, 334 (2021).

[29] M. Mikubo, Y. Inoue, G. Liu, and M. S. Tsao, Journal of Thoracic Oncology **16**, 1798 (2021).

[30] H. Isozaki, R. Sakhtemani, A. Abbasi, N. Nikpour, M. Stanzione, S. Oh, A. Langenbucher, S. Monroe, W. Su, H. F. Cabanos, F. M. Siddiqui, N. Phan, P. Jalili, D. Timonina, S. Bilton, M. Gomez-Caraballo, H. L. Archibald, V. Nangia, K. Dionne, A. Riley, M. Lawlor, M. K. Banwait, R. G. Cobb, L. Zou, N. J. Dyson, C. J. Ott, C. Benes, G. Getz, C. S. Chan, A. T. Shaw, J. F. Gainor, J. J. Lin, L. V. Sequist, Z. Piotrowska, B. Y. Yeap, J. A. Engelman, J. J. K. Lee, Y. E. Maruvka, R. Buisson, M. S. Lawrence, and A. N. Hata, Nature **620**, 393 (2023).

[31] A. Hong, G. Moriceau, L. Sun, S. Lomeli, M. Piva, R. Damoiseaux, S. L. Holmen, N. E. Sharpless, W. Hugo, and R. S. Lo, Cancer Discovery **8**, 74 (2018).

[32] X. Sun, L. F. Wu, S. J. Altschuler, and A. N. Hata, Nature Cancer **5**, 1298 (2024).

[33] J. West, F. Adler, J. Gallaher, M. Strobl, R. Brady-Nicholls, J. Brown, M. Roberson-Tessi, E. Kim, R. Noble, Y. Viossat, D. Basanta, and A. R. Anderson, eLife **12**, 1 (2023).

[34] J. Zhang, J. Cunningham, J. Brown, and R. Gatenby, eLife **11**, 1 (2022).

[35] K. Gallagher, M. A. Strobl, D. S. Park, F. C. Spoendlin, R. A. Gatenby, P. K. Maini, and A. R. Anderson, Cancer Research **84**, 1929 (2024).

[36] B. N. Gantner, F. R. Palma, M. R. Pandkar, M. J. Sakiyama, D. Arango, G. M. DeNicola, A. P. Gomes, and M. G. Bonini, Trends in Cancer **xx**, 1 (2024).

[37] P. L. Freddolino, J. Yang, A. Momen-Roknabadi, and S. Tavazoie, eLife **7**, 1 (2018).

[38] S. Stern, T. Dror, E. Stolovicki, N. Brenner, and E. Braun, Molecular Systems Biology **3**, 1 (2007).

[39] Y. Kaplan, S. Reich, E. Oster, S. Maoz, I. Levin-Reisman, I. Ronin, O. Gefen, O. Agam, and N. Q. Balaban, Nature **600**, 290 (2021).