

Optimal Treatment for Drug-Induced Cancer Persisters Involves Release Periods and Intermediate Drug Doses

Mattia Corigliano ^{1,*}, Arianna Di Bernardo ², Marco Cosentino Lagomarsino ^{1,2,3,†} and Simone Pompei ^{1,†}

¹*IFOM-ETS, The AIRC Institute of Molecular Oncology, via Adamello 16, 20139 Milan, Italy*

²*Dipartimento di Fisica, Università degli studi di Milano, via Celoria 16, 20133, Milan, Italy*

³*Istituto Nazionale di Fisica Nucleare, Sezione di Milano, via Celoria 16, 20133, Milan, Italy*



(Received 28 November 2024; accepted 21 May 2025; published 26 June 2025)

Targeted cancer therapies often induce a reversible drug-tolerant state in subpopulations of cells, akin to bacterial persistence. Precise characterization of these “cancer persisters” may inform the design of more effective treatment strategies. A previous investigation into the transition to persistence of colorectal cancer cell lines has revealed a distinct dependence on drug presence and concentration, not typical of bacterial systems. Leveraging these findings, this study uses mathematical modeling to explore intermittent treatment protocols aimed at diminishing the long-term fitness of the treated population, thereby enhancing therapeutic efficacy. We adapt a mathematical model originally designed for bacteria to describe colorectal cancer population dynamics in response to a series of treatment and release cycles. The model predicts the long-term increase or reduction of a treated population, as well as its asymptotic composition, leading us to identify success and failure regions within a clinically accessible parameter space, also in combination with hypothetical drugs that act on persisters. Strikingly, our analysis suggests, perhaps counterintuitively, that optimal treatment outcomes may be achieved in correspondence of nonzero recovery periods and lower than currently administered in the clinics drug concentrations. Furthermore, by incorporating patient drug pharmacokinetics in the model, we demonstrate that intermittent dosing strategies currently explored in clinical trials can be optimized to potentially rival the efficacy of continuous dosing regimens. These findings underscore the potential of mathematical models in guiding the design of optimal treatment protocols by fine-tuning nontrivial decisional trade-offs.

DOI: [10.1103/mh5v-lcyw](https://doi.org/10.1103/mh5v-lcyw)

I. INTRODUCTION

Persistence refers to the ability of a subpopulation of cells to survive exposure to a fully inhibitory drug concentration. In bacteria, it has been observed since the early 1940s [1], when researchers first noted that a small fraction of bacterial cells could survive antibiotic treatment while the majority were eliminated. Interestingly, populations grown from these surviving cells exhibited the same behavior, ruling out genetic resistance as the underlying cause. Instead, persistence involves a transient and reversible phenotypic state in which a subset of cells enters a dormant or slow-growing phase, enabling them to evade the lethal effects of antibiotics.

The presence of persister cells poses significant challenges in treating bacterial infections. Their resilience can complicate or undermine treatment, resulting in chronic, recurrent infections, prolonged illness, and the need for more aggressive treatment strategies when such strategies are available

[2]. Crucially, the outcome of a treatment in presence of such phenotypic switches depends sensitively on quantitative parameters of the population dynamics, such as growth, susceptibility, and transition rates. Hence, persistence and similar phenotypic switches have inspired extensive mathematical modeling efforts to identify the conditions that contribute to the emergence of resilient subpopulations [2–6].

In recent years, cancer research leveraging work on cell lines and organoids, and parallels with patient outcomes, has uncovered numerous instances of persistence, revealing a striking similarity between the survival strategies of bacteria and residual tumor cells. Both systems harbor subpopulations of slow-growing persister cells exhibiting remarkable resilience to a wide range of therapeutic agents [7–10]. For example, exposure to certain targeted therapies can reproducibly induce the emergence of *bona fide* persister cells within tumor populations [7,11,12]. Crucially, like bacterial persisters, these tumor cells are not genetically resistant to treatment. Instead, they rely on adaptive survival mechanisms that enable them to withstand otherwise lethal therapies. The resurgence of persister cells can drive tumor regrowth, mirroring the case of treated bacteria. This phenomenon has been observed across various tumor types [8], suggesting it may be a general feature, although specific properties may vary between systems. Moreover, persistence in tumors often coexists with more complex, multicellular strategies that enhance therapy evasion.

*Contact author: mattia.corigliano@ifom.eu

†These authors jointly supervised this work: marco.cosentino-lagomarsino@ifom.eu, simone.pompei@ifom.eu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Colorectal cancer persister cells show different dynamics from typical bacterial persisters. (a) Biphasic survival curve observed in treated bacterial and cancer cells, characterized by an initial rapid decline in cell numbers followed by a slower phase. This two-timescale death kinetics distinctly indicates the emergence of drug-tolerant persister cells during treatment. (b) In the bet-hedging scenario, the transition to persistence occurs independently of the external stress. As a consequence persister cells are already present at the beginning of treatment. (c) In the drug-induced persistence scenario, no persisters are expected to be present at the start of treatment, with persister cells emerging only after drug exposure. (d) A constant transition rate as a function of drug concentration, typical in cases with weak dependence on drug levels. (e) Expected curve of transition rate versus drug concentration observed in colorectal cancer cell lines in Ref. [12], showing a monotonic increase that saturates beyond a critical concentration, capturing the drug dependence of persistence in these systems.

However, under certain conditions and in specific experimental systems, cancer persisters can be isolated and accurately quantified [12]. At a quantitative level, both bacterial and cancer persisters exhibit a similar biphasic survival pattern, characterized by an initial rapid decline in susceptible cell numbers, followed by a much slower decline corresponding to the surviving persister population [Fig. 1(a)]. Mathematical modeling techniques originally developed to study bacterial persistence have proven adaptable for interpreting the survival curves of cancer cells [12,13]. Building on this, we previously quantified the transition-to-persistence dynamics in two colorectal cancer (CRC) cell lines (named WiDr and DiFi) when treated with clinical doses of cetux-

imab, a monoclonal antibody that inhibits EGFR, used in patients with metastatic CRC tumors lacking RAS and BRAF mutations [12]. Cetuximab was administered either alone (in two DiFi cloned populations, called cl. B3 and cl. B6) or in combination with the BRAF inhibitor Dabrafenib (in WiDr cloned populations cl. B5 and cl. B7), used for CRC harboring BRAF mutations (see Appendix A).

Our results demonstrated that in these systems, the persister phenotype was largely induced by drug exposure, rather than pre-existing within the tumor population [12]. By contrast, in bacteria, the transition to a persister state typically occurs spontaneously, reflecting a bet-hedging strategy in face of fast-changing environments [Figs. 1(b) and 1(c)] [5,9,14,15]. Conversely, cancer cells are not evolved to respond to changing environments, and presumably persistence is related to entirely different physiological behaviors [8,16]. Additionally, CRC cell-line data sometimes supported a model where the transition rate from normal to persister cells is dependent on drug concentration [Figs. 1(d) and 1(e)]. Specifically, the WiDr cell lines showed a linear increase in the transition rate with rising doses of dabrafenib (Supplemental Material (SM) Figs. 1(i)–1(l) [17]). Conversely, in the DiFi cell lines, the transition rate was zero in the absence of cetuximab but remained constant across the range of all tested concentrations for treated cells (SM Figs. 1(m) and 1(n) [17]). Notably, the transition rates in both cases can be described by the same mathematical equation: a monotonically increasing function that starts from zero and saturates to a constant value above a critical concentration [Fig. 1(e)]. This critical value is lower than the cetuximab treatment concentrations for the DiFi cell line, while it lies within the range of the investigated dabrafenib concentrations for the WiDr cell line. While these behaviors are intriguing, we currently lack any insight on their long-term impact on treatment.

To address this issue, this study investigates how the quantified features of CRC persisters, particularly their drug-induced nature and the dependence of transition rates on drug concentration, can be leveraged to design optimal drug delivery protocols. This is explored in the context of cyclic treatments, where a series of treatment-release cycles are applied, closely mirroring the therapeutic regimens commonly used in clinical practice [18–20]. Previous work allows us to direct and restrict our analysis to population dynamics regimes and parameters within the observed range of behaviors [12]. Our modeling framework builds upon previous modeling studies in the context of bacterial persisters [4,5,15]. These studies also connect experimental observations to dynamic parameters, giving an insight even when detailed mechanisms remain unclear. Recent work has extended these approaches to a cancer context, though different from the case of persisters we address here [21], highlighting the potential of these tools in guiding clinical treatment strategies. This framework raises several key considerations regarding the design of optimal treatment protocols. For instance, how does the sensitivity of transition rates to drug concentration influence treatment outcomes? Specifically, when transition rates exhibit a weak dependence on drug concentration, it might be possible to optimize the duration of drug release. Alternatively, when transition rates show a strong dependence, one might identify an optimal drug dose to minimize the

long-term fitness of treated cancer cells. Finally, one would wish to apply these insights to clinically relevant scenarios, incorporating within-patient, time-dependent drug profile data, to determine whether cyclic treatment protocols offer advantages over constant dosing strategies.

II. RESULTS

A. A phenotypic-switch model quantifies the long-term efficacy of cancer-cell periodic treatments

The model describes the dynamics of cancer cell populations under intermittent treatment protocols [3–5,15,21]. An intermittent treatment protocol is characterized by the number of treatment-release cycles, q , their respective durations, τ_t and τ_r , and the administered dose profile during treatment, $c(t)$ [Fig. 2(a)]. In this work, we consider protocols that either vary the treatment and release durations, while keeping the dosage constant at clinically relevant levels, or adjust the dosage and the release time, while keeping constant the treatment duration. Since we aim to quantify and compare the efficacy of all these different treatment protocols, our analysis focuses on the long-term dynamics of the treated cancer population [Fig. 2(b)], which, as shown in previous work [4,5], can be used as a valid metric for quantifying protocols' efficacy.

Mathematically, the cancer cell population is represented by the vector $\mathbf{N}(t) = \begin{pmatrix} S(t) \\ P(t) \end{pmatrix}$, where $S(t)$ and $P(t)$ correspond to the number of susceptible and persister cancer cells at time t , respectively. The population undergoes q cycles of treatment and release, with treatment duration τ_t and release duration τ_r . During each phase (treatment or release), the time evolution of the population is governed by the first-order differential equation:

$$\frac{d\mathbf{N}}{dt} = M^{t/r} \mathbf{N}(t), \quad (1)$$

where $M^{t/r}$ is the transition matrix for the treatment (t) and release (r) phases, defined as

$$M^r = \begin{pmatrix} \mu_s^r - \lambda_s^r & \lambda_p^r \\ \lambda_s^r & \mu_p^r - \lambda_p^r \end{pmatrix}$$

$$M^t(c(t)) = \begin{pmatrix} \mu_s^t(c(t)) - \lambda_s^t(c(t)) & \lambda_p^t \\ \lambda_s^t(c(t)) & \mu_p^t - \lambda_p^t \end{pmatrix}. \quad (2)$$

Here $\mu_s^{t/r}$ and $\mu_p^{t/r}$ represent the net growth rates of susceptible (S) and persister (P) cells during the treatment or release phases, while $\lambda_s^{t/r}$ and $\lambda_p^{t/r}$ are the transition rates between the two phenotypes [Fig. 2(c)]. Specifically, $\lambda_s^{t/r}$ is the rate of transition from susceptible to persister cells (the “persistence transition”), and $\lambda_p^{t/r}$ is the rate from persister to susceptible cells (the “switch-back transition”).

We point out two key assumptions made so far. First, we neglected the contribution of resistant cells to the overall population dynamics. In Appendix B we show that, for population sizes up to 10^5 cells (and in particular in the experiments inspiring our study), the impact of resistance remains statistically negligible (SM Figs. 2–4 [17]). Second, we adopted a deterministic description for the dynamics of the treated population, neglecting the potential role of the inherent stochasticity of the cell's transition rates. Appendix C justifies




FIG. 2. A phenotypic-switch model quantifies the long-term efficacy of periodic treatments. (a) A periodic treatment protocol is defined by (i) the number of treatment-release cycles q , (ii) the dose profile $c(t)$ during treatment (constant in the figure), and (iii) the durations of the treatment and release phases, τ_t and τ_r . (b) Illustration of the predicted cell population dynamics over q treatment-release cycles. For fixed durations of treatment and release, the model generates distinct trajectories for different drug concentrations. The long-term behavior, indicated by the asymptotic trend for large q (dashed lines), constitutes a proxy for the efficacy of various protocols across many cycles. (c) Schematic representation of the key cellular transitions considered in the model. Susceptible (purple) and persister (green) cells can either proliferate or die (rates μ_s and μ_p) and transition to the other phenotype (rates λ_s and λ_p , respectively). These parameters differ between the treatment (superscript t) and release (superscript r) phases, resulting in a total of eight distinct parameters. Numerical values of the model parameters inferred in Ref. [12] and used in our analysis are reported in Table I.

this choice and shows that the inclusion of stochastic effects does not alter our main findings.

The model's predicted dynamics depend on eight dynamical rates ($\mu_s^t, \mu_s^r, \mu_p^t, \mu_p^r, \lambda_s^t, \lambda_s^r, \lambda_p^t, \lambda_p^r$) and two timescales (τ_t, τ_r) that define the treatment protocol. Additionally, since several of these parameters may vary with drug concentration [e.g., $\mu_s^t = \mu_s^t(c)$], the models predictions are also shaped by how these rates change with dose. This leads to a rich and potentially complex dynamical behavior.

As shown in Appendix D, it is possible to define appropriately rescaled, dimensionless parameters and identify distinct universal physical regimes where the model exhibits

simplified, approximate behavior. Here, instead, to focus the analysis on an experimentally relevant parameter space, we fixed the parameters and their dose dependencies based on observed dynamics from our recent study of CRC cell lines treated with clinically relevant mono- and combination therapies [12]. This approach allows us to target a biologically grounded scenario, while also providing a flexible modeling framework for future applications where parameters can be experimentally calibrated.

1. Drug-induced dose-dependent transition to persistence

To accurately capture the observed drug-induced and dose-dependent transition to persistence in our CRC data [12], we defined the switching rates as $\lambda_s^t = \lambda_s^t(c)$, $\lambda_s^r = \lambda_s^r(0) = 0$, $\lambda_p^t = 0$, and $\lambda_p^r = -\mu_p^t$. We also assume that persisters do not grow during the release phase, that is, $\mu_p^r = 0$, and note that the drug-induced hypothesis also implies the initial condition $P(0) = 0$ at the beginning of the first treatment cycle. We further model the dose dependence of the switching rate using a saturating function,

$$\lambda_s^t(c) = \lambda_{\max}(1 - e^{-c/c_\lambda}), \quad (3)$$

where the value of c_λ is derived from measurements in Ref. [12]. As illustrated in SM Fig. 1 [17] of this study, the transition-to-persistence dose dependence for DiFi and WiDr cells is accurately modeled by a monotonically increasing saturating function. In both cases, the transition rate is zero without drug, confirming that the transition is drug induced. DiFi cells exhibit a constant transition rate for all cetuximab concentrations experimentally tested, compatible with Eq. (3) for ($c \gg c_\lambda$), while WiDr cells show a linear increase in transition rate at dabrafenib concentrations that were experimentally tested. Such linear trend is retrieved in Eq. (3) for ($c \ll c_\lambda$). The different behavior in these cell lines, despite the same functional dependence on drug concentration, reflects the different treatments as well as the distinct genetic backgrounds, which influence sensitivity to treatments. For example, WiDr cells with a BRAF mutation only respond to the combination of cetuximab and dabrafenib (with clinical concentrations of $c_{\text{clinical}}^{\text{WiDr}} = 1 \mu\text{M}$ dabrafenib +50 $\mu\text{g/ml}$ cetuximab), while DiFi cells, which are RAS/RAF wild-type, respond to cetuximab alone (with clinical concentrations of $c_{\text{clinical}}^{\text{DiFi}} = 100 \mu\text{g/ml}$ cetuximab).

2. Dose-dependent death dynamics

To model population dynamics with varying drug concentrations, we describe the dose dependency of the death rate of susceptible cells, $\mu_s^t(c)$. From our CRC data [12], this dependency is well captured by the equation

$$\mu_s^t(c) = \mu_s^t(0) - \Delta d_{\max}(1 - e^{-c/c_\mu}), \quad (4)$$

where c_μ is derived from measurements in Ref. [12]. The minimum inhibitory concentration (MIC), defined as the drug concentration at which the proliferation rate becomes zero, is obtained by solving $\mu_s^t(c_{\text{MIC}}) = 0$, leading to

$$c_{\text{MIC}} = -c_\mu \log \left(1 - \frac{\mu_s^t(0)}{\Delta d_{\max}} \right). \quad (5)$$

TABLE I. Model parameters for four cloned CRC cell line populations *WiDr cl. B5,B7* (treated with cetuximab and dabrafenib) and *DiFi cl. B3,B6* (treated with cetuximab). Parameter values are derived from experimental data from Ref. [12] and shown in SM Fig. 1 [17] of this work. All parameters are expressed in units of days^{-1} . Dose-dependent parameters are given at the clinically relevant concentrations $c_{\text{clinical}} = 1 \mu\text{M}$ dabrafenib for WiDr clones and $c_{\text{clinical}} = 0.69 \mu\text{M}$ cetuximab for DiFi clones.

	WiDr B5	WiDr B7	DiFi B3	DiFi B6
Release^a				
μ_s^r	0.70 d ⁻¹	0.68 d ⁻¹	0.32 d ⁻¹	0.44 d ⁻¹
μ_p^r	0	0	0	0
λ_s^r	0	0	0	0
λ_p^r	0.10 d ⁻¹	0.07 d ⁻¹	0.07 d ⁻¹	0.04 d ⁻¹
Treatment^b				
c_{cli}	1 μM	1 μM	0.69 μM	0.69 μM
$\mu_s^t(c_{\text{cli}})$	-0.67 d ⁻¹	-0.99 d ⁻¹	-0.41 d ⁻¹	-0.61 d ⁻¹
μ_p^t	0 d ⁻¹	-0.07 d ⁻¹	-0.07 d ⁻¹	-0.04 d ⁻¹
$\lambda_s^t(c_{\text{cli}})$	0.01 d ⁻¹	0.14 d ⁻¹	0.23 d ⁻¹	0.23 d ⁻¹
λ_p^t	0	0	0	0
Dose trend^b				
c_μ	0.16 μM	0.34 μM	2.05 nM	2.47 nM
c_λ	75.8 μM	36.5 μM	0.2 nM	0.2 nM
$\mu_s^t(0)$	0.15 d ⁻¹	0.05 d ⁻¹	0.02 d ⁻¹	0.03 d ⁻¹
$\lambda_{\max}$	1 d ⁻¹	5 d ⁻¹	0.23 d ⁻¹	0.23 d ⁻¹
$\Delta d_{\max}$	0.82 d ⁻¹	1.09 d ⁻¹	0.43 d ⁻¹	0.64 d ⁻¹

^aEstimated from fit of growth dynamics data in Ref. [12].

^bEstimated from fit of treatment and dose-response data [12].

This condition will be important when integrating the model with patient-derived pharmacokinetic data. Above the MIC, the proliferation rate becomes negative, indicating that susceptible cells die on average, with the death rate increasing with drug concentration but eventually saturating. We note that the dose-dependent behavior of both the proliferation and transition rates during treatment (illustrated in SM Fig. 1 [17]) is a distinctive feature that differentiates our CRC-specific mathematical framework from previous work.

Table I reports a summary of the numerical values of the model parameters derived from data in Ref. [12] and used in our analyses. For simplicity, the remainder of the main text will focus on showing the results for one representative cloned population WiDr cl. B7, leaving to the SM [17] additional information and specific insights into the behavior of the other studied cloned populations. Appendix E presents a sensitivity analysis testing the robustness of our results to perturbations in the model parameters around their values reported in Table I and shows that the results remain robust across a wide range of parameter variations.

3. Quantification of a protocol's long-term efficacy

Having specified the framework, our goal is twofold: (1) to characterize the conditions under which periodic treatments lead to the extinction of a treated population and (2) to optimize treatment protocols for accelerating cancer extinction dynamics. Both objectives are addressed by quantifying two key observables derived from our model.

First, we evaluate the population's "long-term fitness," Λ , defined as the asymptotic per-cycle exponential rate of population growth or decay [5],

$$\Lambda \equiv \lim_{q \rightarrow \infty} \frac{1}{q} \log N(q). \quad (6)$$

A value of $\Lambda < 0$ ($\Lambda > 0$) indicates the extinction (survival) of the population and thus corresponds to a "treatment success" ("treatment failure") outcome. Importantly, we note that if Λ is only slightly negative, then extinction may take very long. The characteristic time of the population long-term dynamics is given by Λ^{-1} in units of number of treatment-release cycles or by $\Lambda^{-1}(\tau_r + \tau_t)$ in units of days.

The *per-cycle* long-term fitness, Λ , can be estimated in three distinct, but equivalent, ways (see Appendix F for details). First, one can simulate the dynamical system of Eq. (1) over many treatment-release cycles and estimate the long-term fitness as the asymptotic (large q) slope $N(q)/N(q-1) \sim e^\Lambda$ [Fig. 2(b)]. Alternatively, one can diagonalize the matrix $G = e^{M^t \tau_r} e^{M^t \tau_t}$, governing the evolution of the solution of the dynamical system Eq. (1) according to

$$\mathbf{N}(q) = G^q \mathbf{N}(0) = (e^{M^t \tau_r} e^{M^t \tau_t})^q \mathbf{N}(0), \quad (7)$$

and estimate the long-term fitness as

$$\Lambda = \log(g_+), \quad (8)$$

where g_+ is the largest eigenvalue of the evolution matrix G . Computation of the eigenvalues of G can either be done numerically or analytically (see Appendix F and Refs. [3–5,15,21]).

In this study, we first validated each independent estimation method, ensuring they produced consistent results (SM Fig. 5 [17]) and then selected the most suitable method for each analysis. Further details on the long-term fitness quantification, including additional analytical considerations, are provided in Appendices D–F.

The second observable we used to characterize cancer long-term dynamics, "asymptotic population composition," tracks the long-term imbalance between susceptible (S) and persister (P) cells in the total population:

$$\Psi \equiv \lim_{q \rightarrow \infty} \frac{S(q) - P(q)}{S(q) + P(q)}. \quad (9)$$

We chose to focus on the quantity $\Psi \equiv (S - P)/(S + P)$ rather than reporting the individual fractions $\Psi_s \equiv S/(S + P)$ and $\Psi_p \equiv P/(S + P)$, as Ψ provides a symmetric metric centered around zero that more intuitively captures the balance between susceptible and persister cells in the population. The individual fractions can still be easily recovered from Ψ using the relations $\Psi_s = (1 + \Psi)/2$ and $\Psi_p = (1 - \Psi)/2$. A value of $\Psi = 1$ indicates a fully susceptible population ($\Psi_s = 1, \Psi_p = 0$), $\Psi = -1$ corresponds to a population composed entirely of persister cells ($\Psi_s = 0, \Psi_p = 1$), and $\Psi = 0$ denotes an equal mixture of the two cell types ($\Psi_s = \Psi_p = 1/2$).

The "asymptotic population composition," Ψ , can be estimated through simulations or by solving Eq. (7). For large q , the asymptotic composition is given by (see Appendix F)

$$\Psi = \frac{g_1^+ - g_2^+}{g_1^+ + g_2^+}, \quad (10)$$

where g_1^+ and g_2^+ are the components of the leading eigenvector $\mathbf{g}^+$.

It is important to note that both long-term fitness and asymptotic population composition depend on the treatment protocol parameters, i.e., $\Lambda = \Lambda(\tau_t, \tau_r, c(t))$ and $\Psi = \Psi(\tau_t, \tau_r, c(t))$. This raises the question of whether a treatment protocol can be designed to optimize outcomes, such as rapidly extinguishing the cancer population. The following two sections address this question.

While our quantification of treatment protocol efficacy is based on the *per-cycle* long-term fitness consistent with prior work in bacteria [3,5,15] we acknowledge that alternative definitions and proxies can also be considered. In what follows, we continue to adopt the *per-cycle* long-term fitness as our main metric, but we refer interested readers to Appendix G and the associated SM [17] figures for a more comprehensive discussion. For example, one alternative measure is $\tilde{\Lambda} \equiv \Lambda/(\tau_r + \tau_t)$, which expresses long-term fitness in units of inverse time and has recently been proposed as a proxy for treatment efficacy [22]. In Appendix G, we present arguments supporting our choice of *per-cycle* units for the types of optimization problems considered here, and we compare the two approaches, highlighting their differences and implications.

B. Conditions for the design of successful intermittent protocols at constant clinical doses

To investigate how treatment and release durations can be effectively tuned to drive the treated cell population to extinction, we first compared periodic treatment protocols with varying treatment (t_t) and release (t_r) durations, assuming a fixed drug dose profile $c(t) = c$ and a fixed number of treatment cycles q , ideally large in our analysis. This approach allows us to explore whether adjusting the duration and spacing of treatment cycles can enhance efficacy within standard clinical frameworks, where the dose profile and total number of cycles are typically fixed [18–20].

Figure 3 recapitulates our findings for constant-dose protocols with varying release and treatment durations at fixed clinically relevant concentrations of dabrafenib and cetuximab. We present here results for the clone WiDr B7, leaving in SM Figs. 6–8 [17] the results for the other tested CRC cell lines (WiDr B5, DiFi B3-B6), which are all consistent.

Figure 3(a) shows long-term fitness values for different constant-dose protocols and identifies survival ($\Lambda > 0$) and extinction ($\Lambda < 0$) regions. In these protocols, the constant dose is set equal to the clinical dabrafenib dose, as reported in Table I, and treatment and release times vary over a 1-week range [Fig. 3(a)]. By positioning a given protocol relative to the phase boundary $\Lambda = 0$ [solid black line in Fig. 3(a)], which separates survival and extinction regions in the phase plane (τ_t, τ_r), our model predicts whether a clinically plausible protocol will effectively eradicate the population.

Moreover, it predicts the characteristic timescale of the extinction (or survival) dynamics, quantified as the number of treatment-release cycles, $\log 2/|\Lambda|$, or the number of days, $\log 2/|\Lambda|(\tau_r + \tau_t)$, needed to halve (or double) the treated population. For example, a protocol where clinical doses of dabrafenib are administered for $\tau_t = 5$ days every week ($\tau_r = 2$ days) falls in the extinction region, with long-term




FIG. 3. Conditions for the success and failure of intermittent treatment protocols with constant doses. (a) Heat-map showing long-term fitness values, Λ (unit 1/cycles), for periodic treatment protocols with varying treatment and release times on a 0- to 7-d scale and drug concentration fixed to the clinical reference value $c_{\text{cli}} = 1 \mu\text{M}$. The long-term fitness was computed numerically using Eq. (8). The solid black line represents the $\Lambda = 0$ boundary, which separates extinction (blue, $\Lambda < 0$) from survival (red, $\Lambda > 0$) regions. (b) Heat-map showing asymptotic population composition, Ψ , values in correspondence of the same protocols as before. The asymptotic population composition was computed numerically using Eq. (10). Regions where $\Psi > 0$ (purple) indicate that dynamics is driven by a predominantly susceptible population, while regions where $\Psi < 0$ (green) indicate a predominantly persister population. The solid black line indicates the $\Psi = 0$ boundary, where susceptible and persister fractions are equal. (c) The “success-failure (SF) phase diagram” recapitulating both long-term fitness and asymptotic population composition information. (d) Examples of the model’s predicted dynamics in the susceptible-driven failure ($\tau_t = 2$ days, $\tau_r = 4$ days), the susceptible-driven success ($\tau_t = 5$ days, $\tau_r = 4$ days) and the persister-driven success ($\tau_t = 5$ days, $\tau_r = 1$ days) regions of the protocol’s space. WiDr cl. B7 parameters (Table I).

fitness value $\Lambda \approx -0.5 \text{ cycles}^{-1}$, implying a characteristic extinction timescale of about $\log 2/|\Lambda| \approx 1.4$ cycles, or 10 days. In this case, neglecting transient dynamics, a tumor of size $10^6 \approx 2^{20}$ cells would require approximately 28 cycles, or 193 days, of intermittent treatment to be fully eradicated.

Additionally, increasing the treatment duration [moving horizontally in Fig. 3(a)] generally decreases the long-term fitness, thus reducing the average number of cycles needed for complete extinction. For instance, if the treatment duration is extended from $\tau_t = 5$ to $\tau_t = 7$ days while keeping the same release duration of $\tau_r = 2$ days, then the long-term fitness decreases to $\Lambda \approx -0.63 \text{ cycles}^{-1}$, which implies a characteristic timescale of approximately 1 cycle. In this case, the 10^6 -cell tumor would require around 20 cycles and 180 days to be eradicated. However, it is important to note that while the characteristic timescale in cycles may decrease, the total

time does not necessarily decrease, since the duration of each treatment-release cycle also increases with treatment time.

Interestingly, changing the release duration, while keeping the dose and treatment duration fixed [moving vertically in Fig. 3(a)], leads to more complex scenarios. This can result in either improved or worsened outcomes, reflecting a “reentrant” behavior in the extinction region, i.e., by increasing the release time at fixed treatment time, the long-term fitness isolines intersect twice due to the fitness minimum [Fig. 3(a)]. This behavior highlights a trade-off in choosing the release duration, which can be optimized to minimize the long-term fitness. We will discuss this optimization feature in the next section.

Figure 3(b) shows the asymptotic population composition, Ψ , which captures the dynamic interplay between susceptible and persister cells under constant-dose intermittent treatments. Similarly to the analysis of long-term fitness, we observe that the treatment-release phase plane divides into two distinct regions: one where susceptible cells predominantly drive the population dynamics ($\Psi > 0$) and another where the “keystone species” are instead persister cells ($\Psi < 0$). In other words, the dynamics can be piloted by both susceptible and persister cells depending on the protocol. For example, under a protocol where clinical doses of dabrafenib are administered for $\tau_t = 5$ days every week ($\tau_r = 2$ days), the dynamics is predominantly driven by persister cells. Indeed, the population composition value for this protocol is $\Psi \approx -0.5$ [Fig. 3(b)], corresponding to 75% persister cells ($\%P = 75\%$).

When the release duration increases while keeping treatment duration and drug concentration constant [moving vertically in Fig. 3(b)], the population dynamics shifts from being persister-driven at low release durations to being increasingly dominated by susceptible cells at higher release durations. This transition is due to the switch-back process during the release phase: At low release times, only a small fraction of persisters can switch back to the susceptible state. As the release duration increases, a larger fraction of persisters switch back to the susceptible state, resulting in a more susceptible-driven dynamic. Eventually, when the release duration is sufficiently long, all persisters present the treatment phase switch back during the release phase, leading to a completely susceptible-dominated population. Similarly, when the treatment duration is increased while maintaining constant drug concentration and release time [moving horizontally in Fig. 3(b)], the dynamics become increasingly dominated by persister cells. This shift reflects the transition to persistence during treatment, as longer treatment phases provide more time for susceptible cells to transition to the persister phenotype.

Integrating information from both long-term fitness and asymptotic population composition values, we define a comprehensive “success-failure phase diagram” [Fig. 3(c)], recapitulating all possible outcomes under constant-dose intermittent treatment protocols. We define three relevant regions (which are seen consistently across all studied cell lines, SM Figs. 6–8 [17]). The “susceptible-driven failure” (SDF) region contains those treatment protocols that fail in eradicating the tumor population due to a surge in susceptible cell growth, predominantly during release intervals [top region in Fig. 3(d)]. On the other hand, successful extinction can rely on

both susceptible [“susceptible-driven success” (SDS) region in Fig. 3(d)] and persister cells [“persister-driven success” (PDS) region in Fig. 3(d)]. As previously noted, short release intervals favor eradication via elimination of persister cells, whereas longer intervals shift the extinction dependence to susceptible cells.

We finally note that, in principle, a fourth “persister-driven failure” (PDF) region, where survival dynamics is predominantly driven by persister cells can also exist, although not present in Fig. 3. This regime appears considering DiFi cl. B6 parameters, as reported in SM Fig. 9 [17]. Interestingly, this region is narrowly located in correspondence of short treatment and release durations (0–2 days), making it a rather unstable regime, as small variations in both release and treatment time disrupt it. In such cases, during treatment-release cycles the tumor grows on average with a predominant fraction of persister cells. This is possible due to a combination of factors including high transition rates (which makes susceptible cells transition to the persister state) and lower treatment and release durations (which decreases death and switch-back of persisters).

C. The most effective treatment protocols optimize release time and drug concentration jointly

Having understood the conditions underlying successful treatment protocols, we now address the problem of determining the best possible (single-drug) treatment strategy, in terms of minimizing the long-term fitness (thus minimizing residual tumor population and accelerating extinction dynamics). Mathematically, this is an optimization problem over the space of all possible intermittent treatment protocols parametrized by (τ_t, τ_r, c) . The optimal treatment strategy is defined as the protocol $(\tau_t^*, \tau_r^*, c^*)$ in correspondence of which the long-term fitness is minimum:

$$(\tau_t^*, \tau_r^*, c^*) = \operatorname{argmin}\{\Lambda(\tau_t, \tau_r, c)\}. \quad (11)$$

We numerically solved Eq. (11) over the parameter range $[0, 7] \text{ d} \times [0, 7] \text{ d} \times [10^{-8}, 10^{-4}] \text{ M}$, and found that the optimal strategy for the WiDr B7 clone is

$$\tau_t^* = 7 \text{ d}, \quad \tau_r^* = 3.3 \text{ d}, \quad c^* = 3.8 \times 10^{-7} \text{ M}. \quad (12)$$

Consistently with our previous observations, we note that the best choice for treatment duration is the maximum achievable value. This is because, as we already discussed, the general effect of treatment duration is to decrease the per-cycle long-term fitness, that is, minimize the average number of treatment-release cycles needed to extinguish the cancer population. On the other hand, the optimal values for the release time and drug concentration are less trivial and require further investigation. The remainder of this section focuses on discussing the optimization of these two parameters. Again, we initially consider parameters from the clone WiDr cl. B7, where a clear dependency on drug concentration of both proliferation and transition rates is seen (SM Fig. 1 [17]), leaving the discussion of all other studied CRC cell lines to the end of this section.

To investigate in further detail how varying drug concentration and release duration affects the long-term efficacy of treatment protocols, we set a constant treatment duration

$\tau_t = 7$ days (the optimal value set by the boundary in our exploration). As described in the previous section, we computed long-term fitness and population composition varying release durations and drug concentrations to reconstruct treatment outcomes in the release time - drug concentration phase plane (Figs. 4(a) and 4(b) and SM Figs. 10(a) and 10(b) [17]). Our analysis reveals again that the release time-drug concentration phase plane can be divided into three regions: SDF, SDS, and PDS, as shown in SM Figs. 10(b) and 10(c) [17]. Increasing release time shifts the population towards more susceptible cells, as longer release durations allow more persister cells to switch back and proliferate. In contrast, higher drug concentrations increase the fraction of persister cells by both accelerating the elimination of susceptible cells and inducing more persisters, resulting in a persister-dominated population.

Figure 4(a) shows that increasing release duration as well as drug concentration affects long-term fitness. For protocols with fixed treatment and release times while varying drug concentration [moving along horizontal lines in Fig. 4(a)], long-term fitness initially decreases at low drug concentrations, reaching negative values around the clinical dabrafenib concentration ($c = 1 \mu\text{M}$), before increasing again. Similarly, for protocols with fixed treatment duration and drug concentration while varying release time [moving along vertical lines in Fig. 4(a)], long-term fitness first decreases, and then increases, eventually reaching positive values for sufficiently high release durations. This behavior suggests the presence of an optimal protocol [red star in Fig. 4(a)], which minimizes long-term fitness by jointly optimizing drug dose and release time. Figure 4(b) shows how these findings are paralleled by population composition. For the specific WiDr B7 CRC cloned population used here, the optimal protocol balances the fractions of persister and susceptible cells [red star in Fig. 4(a) and 4(b)]. This balance in persister and susceptible cell dynamics is crucial for achieving optimal extinction outcomes. However, this feature is not universal (and thus is likely to be patient specific), as different parameters may lead to varying optimal population compositions, as seen in SM Figs. 11–13 [17], which show results for other CRC clones.

Figure 4(c) shows the dose dependence of long-term fitness at the optimal release time $\tau_r^* = 3.3$ days, where the optimal drug concentration is $c^* = 3.8 \times 10^{-7} \text{ M}$, less than half the clinical reference value ($c = 1 \mu\text{M}$). This finding suggests that lower drug doses may reduce toxicity and improve therapeutic efficacy by increasing extinction speed. This dose optimization is a consequence of the dabrafenib dose-dependent behavior of the WiDr cell line, as both death and transition-to-persistence rates increase monotonically with drug concentration (depicted in SM Fig. 1 [17] and recapitulated by Eqs. (3) and (4)). There is a trade-off between using lower drug concentrations, where both death and transition rates are low, and higher doses, where both rates are fast [Fig. 4(e)]. As shown in our previous study [12], the optimal concentration balances this trade-off by maximizing killing dynamics while maintaining a moderate transition rate to persistence. Suboptimal concentrations occur when lower doses fail to increase the death rate sufficiently, while higher doses overly increase the transition rate with little additional effect on death rate, which is already at its maximum [Fig. 4(e)].




FIG. 4. The optimal treatment protocol minimizes long-term fitness through joint optimization of release time and drug concentration. [(a) and (b)] Heat-maps showing long-term fitness (a) and asymptotic population composition values (b) as a function of drug concentration and release duration for a fixed treatment time of 7 days. Optimal protocol (red star) is achieved at values $\tau_r^* = 3.3$ days and $c^* = 3.8 \times 10^{-7}$ M. The colorbar in panel (a) reports values of Λ in unit of $1/\text{cycles}$. (c) The long-term fitness as a function of dabrafenib concentration is minimal (achieving a value of $\approx -0.66 \text{ cycles}^{-1}$) at the optimal drug concentration $c^* = 3.8 \times 10^{-7}$ M. (d) Long-term fitness as a function of release duration is minimum at the optimal time of $\tau_r^* = 3.3$ days. [(e) and (f)] Illustration of the dose (release time) trade-off defining an optimal drug concentration (release duration). Both too low (short) and too high (long) doses (release times) yield higher final population sizes than the optimal drug concentration (release time). (e) Below (above) the optimal dose, drug concentration can be increased (decreased) to increase (decrease) the death rate (transition rate) of susceptible cells without increasing (decreasing) the transition-to-persistence (death) rate. (f) Below (above) the optimal release, release duration can be increased (decreased) to increase (decrease) the fraction of persister cells resuming proliferation. WiDr cl. B7 parameters (Table I).

Similarly, Fig. 4(d) shows the long-term fitness dependence on release time at the optimal drug concentration $c^* = 3.8 \times 10^{-7}$ M. We note that, within clinically relevant concentration ranges, which lie around the optimal concentration value, neither significantly shortening nor excessively extending the release time yields optimal outcomes, suggesting a precise balance between dosing intervals. This timing-based optimization highlights the interplay between persister cell reversion to a vulnerable state and subsequent drug efficacy. As shown in Fig. 4(f), if the release time is extended, then more persister cells revert to a drug-sensitive state, enhancing the population's response in the following cycle. However, if the release time is too short, then most treated cells remain in a persister state, reducing the overall impact of the next cycle.

We replicated the analyses discussed so far with the other CRC cloned populations. SM Figs. 11–13 [17] show that release time optimization is found in both our CRC cell lines across both cloned populations, underscoring the importance of precisely timed drug administration to maximize therapeutic efficacy. On the other hand, dose dependency of the persister transition is specific of WiDr clones treated with different doses of dabrafenib, hence resulting in a optimal dose minimizing the long-term fitness (for both WiDr B5 and B7). Conversely DiFi populations (both B3 and B6) do not show

an increased transition rate to persistence for increasing drug concentrations, leading to a plateau of optimal concentration values for long-term fitness (SM Figs. 12 and 13(c) [17])

In Appendix D, we derive the following analytical approximation for the long-term fitness, valid in the vicinity of the optimum:

$$\Lambda \approx \log \left(\frac{\epsilon_-^r \epsilon_+^t [(1 + AB)\epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t]}{\epsilon_+^r \epsilon_-^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t) - \epsilon_-^r \epsilon_+^t} \right), \quad (13)$$

where the following shorthand notations have been introduced: $\epsilon_+^r = e^{\mu_s^r \tau_r}$, $\epsilon_-^r = e^{-\lambda_p^r \tau_r}$, $\epsilon_+^t = e^{\mu_p^t \tau_t}$, $\epsilon_-^t = e^{(\mu_s^t - \lambda_p^t) \tau_t}$, $A = \frac{\lambda_s^r}{\mu_s^t - \lambda_s^t - \mu_p^r}$, $B = \frac{\lambda_p^r}{\mu_s^t + \lambda_p^r}$ (see Appendix D for further details). SM Figs. 14(a), 14(c), 14(e), and 14(g) [17] show the good agreement between the approximation in Eq. (13) (blue line) and the numerically computed exact long-term fitness (black line), confirming the validity of the analytical expression near the optimum. Importantly, by solving the conditions $\frac{\partial \Lambda}{\partial c} = 0$ and $\frac{\partial \Lambda}{\partial \tau_r} = 0$, we find that the optimal drug concentration and release time satisfy the following two balance relations:

$$\frac{d\epsilon_-^r/dc}{dA/dc} = K_c, \quad \frac{d\epsilon_+^r/d\tau_r}{d\epsilon_-^r/d\tau_r} = K_r, \quad (14)$$

where K_c and K_r are parameter-dependent constants. At optimality, these constants take the values $K_c \in \{0.58, 1.39, 0.71, 0.57\}$, $K_r \in \{-84, -177, -34, -8.7\}$ for WiDr cl. B7, WiDr cl. B5, DiFi cl. B6, and DiFi cl. B3, respectively. Thus, our analytical approximation provides a mathematical description of the optimal treatment strategy, which simultaneously balances: (i) the rate of susceptible cell elimination against their transition to persisters during treatment, and (ii) the regrowth of susceptible cells during the drug-free phase against the rate at which persisters revert to susceptibility. The approximate optimal solutions derived from the balance conditions in Eq. (14) are shown in SM Figs. 14(b), 14(d), 14(f), and 14(h) [17], alongside the exact numerical optima, further validating the accuracy and usefulness of our analytical approach near the optimum.

To quantify the gain in treatment efficacy achieved through the joint optimization of drug concentration and release duration, we compare the performance of the optimal protocol against (i) a range of reasonable alternatives and (ii) a benchmark protocol defined by continuous administration of the clinical reference dose, $c_{\text{cli}} = 1 \mu\text{M}$, with no release period ($\tau_r = 0$). We first note that, comparing the efficacy of all protocol combinations that lead to a negative long-term fitness, our analysis also includes protocols yielding only a modest (i.e., near-zero) long-term fitness (Fig. 4(a), SM Figs. 10(a) and 10(b) [17]). For instance, administering clinical doses of dabrafenib ($c = 1 \mu\text{M}$) for $\tau_t = 7 \text{ d}$ every 9 d ($\tau_r = 2 \text{ d}$) results in a long-term fitness of approximately -0.6 , whereas extending the interval between treatments to $\tau_r \approx 7 \text{ d}$ yields a fitness close to zero. Hence, the predicted gain in extinction time through protocol optimization can, in principle, become arbitrarily large across the space of effective treatment strategies. Even when compared to the clinically relevant benchmark, the efficacy of the optimal protocol remains substantially higher. Our model predicts, perhaps counterintuitively, that lower drug doses combined with longer treatment-free intervals can significantly improve outcomes by reducing the number of treatment cycles needed for complete tumor eradication. Specifically, the gain in long-term fitness is approximately 40% (from ≈ -0.5 to ≈ -0.7 ; SM Fig. 10(d) [17]). In addition, this gain is associated with a marked shift in population composition from a persister-dominated state to one with nearly equal fractions of susceptible and persister cells (SM Figs. 10(b) and 10(c) [17]). Under the benchmark protocol, long-term population dynamics are eventually driven by persister cells, and the corresponding per-cycle long-term fitness becomes $\Lambda_{\text{benchmark}} = -0.073 \times 7 \approx -0.5$, reported as a horizontal dashed line in SM Fig. 10(d) [17]. The improved long-term fitness observed under the optimal protocol suggests that allowing a fraction of persister cells to switch back to a susceptible state during treatment-free intervals can accelerate extinction dynamics. To further substantiate these predictions, we performed stochastic simulations of a population of 10^4 cells undergoing both the optimal and benchmark treatment protocols (SM Fig. 10(e) [17]) and computed the associated extinction time distributions (SM Fig. 10(f) [17]). These simulations confirm that the optimal protocol, despite involving lower doses and less frequent treatments, outperforms the benchmark by reducing

the average extinction time by more than 30% (10 cycles versus 15).

Finally, we reiterate that the optimization results discussed in this section are valid with respect to our choice of using the *per-cycle* long-term fitness as a metric for treatment efficacy and within the clinically relevant and experimentally motivated protocol bounds $[0, 7] \text{ d} \times [0, 7] \text{ d} \times [10^{-8}, 10^{-4}] \text{ M}$. While this approach is both natural and practical for the problems considered in this work (see Appendix G for detailed discussions), alternative methods have also been proposed, sometimes yielding similar results, but in other cases leading to different conclusions [22]. In Appendix G, we compare the outcomes obtained using various time units and efficacy proxies. We conclude that, while redefining long-term fitness in different time units may affect the specific optimal solutions, considering alternative proxies such as tumor size at a fixed time point (SM Fig. 15 [17]) or the probability of treatment success (SM Fig. 16 [17]) does not significantly alter our main findings.

D. Optimal intermittent dosing of dabrafenib rivals the efficacy of continuous treatments

To further assess the potential clinical applicability of our quantitative framework, we sought to connect it to the ongoing clinical debate surrounding continuous versus intermittent dosing regimens [20]. To this end, we transitioned from the simplified assumption of constant drug treatment profiles to explore models reflecting the dose kinetics observed in patients treated with similar doses of the same targeted drugs we used with our CRC cell lines (Table II). This required shifting from a constant concentration to a time-dependent profile that decreases after each administration. We assumed that the reduction in drug concentration directly correlates with decreased death rates and transition rates for susceptible cells, as measured for our cell lines.

Inspired by experimental observations [19], we model the in-patient dose profile as a double exponentially decreasing function of the form:

$$c(t) = \theta(t - t_0)\xi(e^{-k_e(t-t_0)} - e^{-k_a(t-t_0)}), \quad (15)$$

which, albeit very simple, recapitulates known key aspects of drug pharmacokinetics, such as temporal delays (t_0), maximum induced plasma concentration (c_{max}), and both dose absorption (rate k_a) and elimination (rate k_e) dynamics [Fig. 5(a)].

Importantly, we used the experimental measurements from the clinical trial in Ref. [19], characterizing the population pharmacokinetics of dabrafenib, to fix the parameters in Eq. (15). Specifically, we use reported values for the absorption rate, $k_a = 1.88 \text{ h}^{-1}$ ($\tau_a = \log(2)/k_a \approx 0.36 \text{ h}$), maximum induced (peak absorption) plasma concentration, $c_{\text{max}} = 1512 \text{ ng ml}^{-1}$, lag time, $t_0 = 0.48 \text{ h}$, and area under the curve, $\text{AUC} = 8820 \text{ ng ml}^{-1} \text{ h}$ to estimate ξ and k_e as (see Appendix H for further details)

$$k_e = -k_a \frac{W\left(-\frac{c_{\text{max}}}{k_a \text{AUC}} \log \frac{k_a \text{AUC}}{c_{\text{max}}}\right)}{\log \frac{k_a \text{AUC}}{c_{\text{max}}}} \approx 0.23 \text{ h}^{-1}, \quad (16)$$




FIG. 5. Exploiting drug pharmacokinetics to optimize intermittent dosing of dabrafenib in patients. (a) Illustration of the typical time course of drug concentration in patients, comprising absorption and elimination dynamics. Reported values for dabrafenib dose (150 mg BID), peak plasma concentration ($c_{\max} = 1512 \text{ ng ml}^{-1}$), absorption ($k_a = 1.88 \text{ h}^{-1}$, $\tau_a = 0.36 \text{ h}$), and elimination ($k_e = 0.23 \text{ h}^{-1}$, $\tau_e = 3 \text{ h}$) rates, have been derived from study [19] and used to fix the plasma concentration profile reported in panel (b). (b) Inferred plasma concentration profile, showing a rapid absorption phase ($\approx 0.36 \text{ h}$), followed by a slower exponential decline of drug concentration over time ($\approx 3 \text{ h}$). After 1 day the plasma concentration reaches the minimum inhibitory concentration of dabrafenib ($\approx 8 \text{ ng/ml}$, dashed horizontal line). (c) Constant versus intermittent dosing of dabrafenib, as recently compared in clinical trial [20]. (d) Comparison of predicted cancer decline for constant and intermittent dosing of dabrafenib. Red line shows dynamics under the optimal intermittent dosing strategy. (e) Illustration showing optimal intermittent dosing parameters. (f) Long-term fitness as a function of release time showing nonoptimality of both continuous and intermittent dosing strategies considered in study [20]. Optimal release duration, defining the optimal intermittent strategy, is equal to $\tau_r = 2.6 \text{ days}$ (red star).

with $\tau_e \approx 3 \text{ h}$, and where W is the Lambert function, and

$$\xi = \frac{k_e k_a \text{AUC}}{(k_a - k_e)} \approx 2301 \text{ ng ml}^{-1}. \quad (17)$$

The resulting dose profile predicted for a patient,

$$c(t) = \begin{cases} 0, & t < 0.48 \\ 2301[e^{-0.23(t-0.48)} - e^{-1.88(t-0.48)}], & t \geq 0.48 \end{cases} \quad (18)$$

is plotted in Fig. 5(b) over a daily time span. We note that in about one day, the inferred dabrafenib plasma concentration reaches the minimum inhibitory concentration (MIC), $c_{\text{MIC}}^{\text{dab}} \approx 8 \text{ ng ml}^{-1}$ [dashed horizontal line in Fig. 5(b)], as estimated from Eq. (5). Below the MIC, dabrafenib concentrations are no longer capable of inducing susceptible cell death, leading us to set $\mu_s^*(c_{\text{MIC}}) = 0$ in our model, and motivating the need of renovating dabrafenib administration on a daily timescale [19,20]. Furthermore, we note that the time dependence of the drug concentration makes the dose-dependent parameters in our model (susceptible cells' death and transition-to-persistence rates), dependent on time. SM Fig. 17 [17] shows the specific time dependence of our model parameters under these conditions.

Equipped with these additional features, we simulated our model to investigate the problem of whether intermittent dosing of dabrafenib might be preferable to constant dosing regimens in a patient. We note that a recent clinical study [20] addressing this same question reported no evidence in favor of an intermittent treatment strategy. Here, by comparing the model's predicted cancer dynamics under the same continuous and intermittent dosing protocols considered in Ref. [20], depicted in Fig. 5(c), we revised this comparison from the more quantitative point of view of assessing the long-term dynamics of the treated cell population. For simplicity, we computed the long-term dynamics under both continuous and intermittent dosing protocols further approximating Eq. (18) by considering an instantaneous absorption dynamics.

Consistently with the findings of the clinical-trial [20], an intermittent dosing (ID) strategy where dabrafenib is administered for 21 days every 28 days (release time of 1 week) shows inferior performance in our model compared to a continuous dosing (CD) strategy in which dabrafenib is administered every day [Fig. 5(d)]. Interestingly, however, when we exploit our model's flexibility to compute how long-term fitness changes as a function of the release time interval between consecutive dabrafenib administrations, we see that neither the CD nor the ID considered in the clinical trial in Ref. [20] are predicted to be optimal choices, when aiming at minimizing the long-term fitness of the treated cell population. Instead, we find that an intermittent protocol with a release time interval of about 3 days, $\tau_r^* = 2.6 \text{ days}$ [Fig. 5(e)], minimizes the long-term fitness [Fig. 5(f)], yielding the fastest achievable extinction dynamics of the total cancer population [red line in Fig. 5(d)]. Thus, our analysis leads us to surmise that an optimized intermittent protocol, particularly when its release timing is finely tuned, may be able to yield outcomes that not only rival but may even exceed those of a continuous dosing regimen.

E. Combination therapies modulating persister dynamics lead to differential treatment outcomes

Persister cells play a critical role in enabling tumors to rapidly tolerate treatments. Our previous results also high-




FIG. 6. Quantification of the long-term fitness effects of combination therapies targeting persister cells dynamics. [(a) and (b)] Explored combination therapies include additional drugs that either reduce the transition-to-persistence rate [type I, panel (a)], or increase persisters' death rate [type II, panel (b)]. [(c) and (d)] Heat-maps showing the difference in long-term fitness between combination therapies, either type I [panel (c)] or type II [panel (d)], and single-drug treatment protocols with varying treatment and release times on a 0- to 7-day scale and drug concentration fixed to clinical reference levels. Red (blue) regions correspond to an increase (decrease) in long-term fitness upon addition of the auxiliary drug. (e) Phase diagram integrating information from panels (c) and (d), individuating three main possible effects upon drug addition: detrimental or negligible effects (case 1), beneficial only if type II drugs are added (case 2) and beneficial under either drug (case 3). (f) Models predicted dynamics under protocols from each of the three cases, confirming the expected different behavior. WiDr cl. B7 parameters.

lighted their potential influence on the long-term outcome of periodic treatment protocols. For instance, a large, poorly tuned fraction of persister cells switching back to a proliferative, drug-sensitive state can ultimately lead to treatment failure. In addition, persisters may serve as a reservoir for the emergence of resistant clones over time (although resistance is not explicitly modeled here), further compromising treatment success. Building on this idea, we hypothesized that combination therapies incorporating drugs specifically targeting persister cell dynamics could improve treatment efficacy and potentially accelerate tumor eradication. To test this, we conducted a proof-of-principle analysis using our quantitative framework, comparing the long-term efficacy of hypothetical combination therapies aimed at targeting cancer persistence with that of single-drug regimens at clinical doses of cetuximab and dabrafenib (Table I). We considered two potential mechanisms of action for the auxiliary drug: (1) reducing the transition-to-persistence rate (type I) or (2) increasing the death rate of persister cells (type II). These mechanisms were incorporated into our model as a percentage decrease or increase in the respective rates [Figs. 6(a) and 6(b)]. For instance, when a molar concentration, c' , of a type I auxiliary drug is added, we assumed that the transition-to-persistence rate switch to a value

$$\lambda_s^t(c, c') = f(c')\lambda_s^t(c) = f(c')\lambda_{\max}(1 - e^{-c/c_s}), \quad (19)$$

with $f(c') \in (0, 1)$. Alternative modeling choices are possible, such as modifying the characteristic dose scale c_s , but both approaches ultimately decrease $\lambda_s^t(c, c')$. Given the absence of specific data to favor one mechanism over another, we adopted the simpler form described in Eq. (19). For subsequent analyses, we assumed an illustrative value of $f(c') = 0.5$ [Figs. 6(a) and 6(b)]. However, we acknowledge that a more comprehensive analysis would require precise experimental characterization of $f(c')$. We simulated the model before and after introducing the auxiliary drug and calculated the differences in long-term fitness between single-drug and combination therapy protocols. We varied treatment and release times within a 0- to 7-day range as in previous analyses, keeping the drug dose constant at clinical reference values. Interestingly, our analysis revealed distinct long-term fitness outcomes depending on whether a type I or type II drug was used [Fig. 6(c) and 6(d)]. Specifically, adding drugs that increase persister cell mortality (type II) generally reduced long-term fitness across protocols [Fig. 6(d)], with a varying degree of efficacy. In contrast, combination therapies targeting the transition-to-persistence rate (type I) showed both beneficial and detrimental effects, depending on the specific treatment protocol [Fig. 6(c)].

We integrated these findings into a phase diagram [Fig. 6(e)], recapitulating the long-term fitness effects of combination therapies across protocols with varying treatment

and release durations, while maintaining fixed clinical doses. Remarkably, our analysis divides the (τ_t, τ_r) phase plane into three distinct regions [labeled “case 1,” “case 2,” and “case 3” in Fig. 6(e)], showing different responses to combination therapies. Specifically, protocols lying in the region where susceptible cells dominate the long-term dynamics (case 1, comparing with Fig. 3) show no significant benefits from adding an auxiliary drug targeting persister cells, with type I drugs often having detrimental effects. Consequently, drugs increasing persister cell mortality had minimal impact, while those reducing the transition-to-persistence rate could inadvertently increase the number of susceptible cells. The top panel of Fig. 6(f) illustrates the population dynamics in this region, ($\tau_t = 2$ days, $\tau_r = 6$ days, $c = 1$ μ M). On the other hand, protocols belonging to the region where persister cells predominantly drive the extinction of the cancer population [case 2, comparing Figs. 6(e) and 3] largely benefit from type II drugs targeting persister cell mortality, while type I drugs have little to no effect [Figs. 6(c)–6(e)]. In this region, interventions that increase persister cell death are particularly effective and reducing transition to persistence has little effect [see the middle panel of Fig. 6(f) for an example of the predicted dynamics for the protocol, with parameters $\tau_t = 7$ days, $\tau_r = 1$ days, $c = 1$ μ M]. Finally, we find empirically an intermediate region [case 3 in Fig. 6(e)] region, where both combination therapies yield comparable beneficial effects on the long-term fitness [Figs. 6(c)–6(e) illustrated in bottom panel of Fig. 6(f) for parameters $\tau_t = 7$ days, $\tau_r = 6$ days, $c = 1$ μ M]. In summary our analysis indicates that hypothetical combination therapies targeting the transition-to-persistence rate or persister cell mortality can enhance treatment efficacy very differently, and their impact should depend heavily on the specific treatment protocol. The next section returns to single-drug protocols to further explore optimal treatment design.

III. DISCUSSION

Targeted therapies based on EGFR inhibitors like cetuximab, and their combination with dabrafenib for BRAF-mutated tumors, were shown to extend median survival by a few months in colorectal cancer patients when added to standard chemotherapy, compared to chemotherapy alone [23]. This is a modest but non-negligible outcome, and it is important to understand how it can be improved upon. While a limitation to the efficacy is certainly due to the emergence of resistance, phenotypic transitions and persistence also likely play a role [8,11,24]. Given the limitations of current targeted therapies and the potential role of mechanisms beyond resistance, we explored the efficacy of periodic treatment-release protocols. In this analysis, we have assumed that resistance does not develop (see Appendix B for a discussion). While of course we acknowledge that resistance is a significant clinical challenge, we have shown that for sufficiently small population sizes (up to 10^5 for the regime of the parameter space motivating this study) it can be safely neglected. In a more clinical horizon, population size will be essential to set whether persisters play a role or pre-existing and *de novo* mutations make treatment strategies useless. Note, however, that recent work suggests that even in the presence of resis-

tant clones, treatment release periods can still be beneficial if competitive nonresistant clones remain [21]. Additionally, in the case of resistance emergence due to drug-induced mutagenesis, mathematical modeling has also shown that less aggressive treatments could potentially lead to improved outcomes [25]. However, while these lines of evidence point to the idea that periods of treatment release can generally be an advantage, they also show that such release periods need to be precisely targeted on the specific parameters of a tumor population (and adapted if these parameters change), or they risk to be detrimental. Since the precise knowledge of these parameters in a clinical setting is scarce, the practical applicability of these concepts remains very limited, although as personalized-medicine assays advance, one could envisage that such parameters may be used to optimize treatment time-administration protocols, together with choice of treatments (see below).

To avoid these complications, we have rooted our analysis on observations and parameter values measured in CRC cell lines [12]. Importantly, we note that the parameters and behaviors were different in the two tested CRC cell lines and treatments, supporting the idea that patient specificity might be relevant. Under the radical (but necessary) assumption that the parameters measured on cell lines may guide us to reveal relevant scenarios, our findings suggest a nonintuitive landscape for the effectiveness of periodic treatments in controlling and potentially eradicating treated cancer cell populations with persisters. Specifically, we show how long-term fitness defines the boundaries between successful therapy and treatment failure and how population composition in susceptible vs persisters is essential for understanding optimal treatment protocols, due to a delicate balance between treatment (which kills susceptible cells but produces persisters) and release phases (which clear persisters but allow susceptible cells to proliferate). A consequence of this complex balance is the observation of reentrant behavior in extinction dynamics (Fig. 3), where increasing the duration of release periods can counterintuitively optimize the eradication process, by providing sufficient time to clear persisters (a result unattainable under continuous treatment). In scenarios where treatment fails, unchecked proliferation of susceptible cells during release intervals prevails. Conversely, effective treatment regimens capitalize on either efficient elimination of susceptible cells or control of persister generation, achieving optimal outcomes when they can accomplish both. This insight underscores a delicate interplay dependent on parameters such as the persister transition rates that may be challenging to access in realistic situations, where slight adjustments could be critical in tipping the balance between therapy failure and success.

Our framework also allows for exploring hypothetical combination therapies targeting persister death and the transition from susceptible cells, which can significantly improve the outcome of periodic treatment protocols. If drug dose impacts on both the death rate of susceptible cells and their transition rate to persisters, as in WiDr cells treated with dabrafenib and cetuximab, then the scenario is even more complex, due to the quantitative trade-off between effective killing and persister production. As shown previously [12], an increasing dose ramp decreases persister production in a

single treatment compared to the same total dose administered continuously. While this work focuses on constant alternated treatments, we speculate that combining optimal dose ramps with optimized protocols could further improve patient outcomes if persisters are relevant. For periodic treatments at constant doses, we identified an optimal dose near (but lower) the clinically relevant dose of dabrafenib, where the trade-off between proliferation suppression and persistence induction yields the best outcomes. Collectively, these findings suggest that specifically targeting the persistence mechanism could enhance therapeutic efficacy.

We note that the concept of intermittently resuming treatment to allow persisters to resume proliferation is well-established in microbiology such as in bacterial spore eradication [26] and in pulse dosing antibiotic treatments optimization [27–30]. In this work, we adapted these concepts to mammalian systems and cancer, where the dynamics are more complex. Specifically, while bacterial persistence has been extensively modeled as a bet-hedging strategy in fluctuating environments [31], our study explores a completely different parameter regime, a drug-induced and dose-dependent persistence, yielding rich and nontrivial phase diagrams, that are fundamentally distinct from the ones of standard microbial models. Our approach also contrasts with existing cancer literature [25,32–34], which typically considers the benefits of drug holidays in the context of genetic heteroresistance. In such scenarios, resistant clones coexist with susceptible cells and are subject to fitness trade-offs, allowing drug-sensitive cells to dominate during treatment breaks. Persistence, by contrast, involves reversible, drug-induced phenotypic switching without genetic alterations allowing persister cells to revert to drug sensitivity, a possibility absent in genetically resistant populations.

Turning to more general implications of our findings, we have shown (Appendix D) how the long-term fitness of a cell population under treatment can be universally characterized by nine dimensionless (scale-invariant) parameters that capture the relative growth/decay rates and transition timescales between treatment and release phases, allowing for a model-independent understanding of the system's behavior. The evolution of the population can be represented by a universal matrix form (whose entries are linear combinations of the dimensionless parameters) whose highest eigenvalue determines the long-term fitness, revealing fundamental trade-offs that govern optimal treatment strategies, such as balancing the killing rate of susceptible cells with the transition rate to persister cells. By identifying these universal aspects, we provide a framework to understand and optimize treatment strategies across different systems, without needing to simulate every specific scenario, and in particular to identify the trade-offs in the relevant situation of drug-induced and dose-dependent persister transitions [35]. In other words, we identify a general “drug-induced” scenario, transcending the specific parameters analyzed here and different from the regimes investigated previously with our approach [5,15]. Hence, while using experimental parameters to provide a firm motivation for the relevant regimes to consider, our framework transcends a specific mathematical model designed to fit specific experimental data, and provides generalizable theoretical guiding principles

to investigate the role of phenotypic plasticity in persistence across diverse contexts.

We recognize that there are significant challenges in translating these highly theoretical optimizations into practical clinical applications. Key factors such as pharmacokinetics and pharmacodynamics and effective behavior of a treated tumor *in vivo*, where several factors related to the microenvironment can play an adverse or favorable role and may play a crucial role in determining the effective dose experienced by a tumor, introducing variability that must be carefully considered [18–20]. Moreover, the fact that cetuximab is typically administered in combination with chemotherapy regimens like FOLFIRI or FOLFOX, rather than as a monotherapy, adds another layer of complexity [36–38]. Encouragingly, a recent study showed that intermittent administration of FOLFIRI and panitumumab achieves better outcomes than continuous treatment [39]. However, substantial interpatient variability in plasma cetuximab concentrations [18] further complicates the prediction of treatment efficacy. Spatial effects, including heterogeneous drug distribution within the body, toxicity concerns, and the influence of the tumor microenvironment, must also be accounted for when considering the clinical implementation of our findings. To take a first step towards mitigating this gap of knowledge, we have simulated our model with time-dependent doses and constraints derived from pharmacokinetics data on dabrafenib [19,20], and therefore with realistic values for the time-dependent plasma concentrations of the treatment but still assuming that the death-rate and transition-rates were those observed in cell lines. Under these assumptions, the model prediction is that treatment-release periods of about 2 days would still outperform continuous treatment (but much longer release times would not).

However, given the complexity of the problem and the extremely limited available knowledge, our work serves merely as a proof of concept, offering a foundation for future research to validate and refine these hypotheses. One promising approach could involve experimental comparisons of treatment regimens in cell lines, patient-derived organoids, and mouse xenografts [40,41], providing increasingly sophisticated and controlled models for assessing the efficacy of various treatment strategies. These experimental systems would allow for testing the predictive power of our simplified modeling framework in more clinically relevant environments. Alternatively, a more empirical approach might involve directly designing clinical trials to test various time-dependent protocols in patients, such as maintaining an equal total dose while varying the administration schedule [18–20]. Such trials could help verify the broader hypothesis that “keystone” persisters may be relevant for therapy outcomes and could be guided by mathematical models based on *in vitro* systems and measured pharmacokinetic parameters. If successful, then these protocols could ideally extend beyond CRC to other cancers, such as lung cancer treated with osimertinib, where persister transitions have been observed *in vitro* [7,42]. However, to motivate these protocols for other cancer treatments, a more advanced quantitative *in vitro* characterization of persister behavior, similar to that performed for CRC, would be essential [12]. This approach, combining theoretical insights with empirical

validation, may effectively reach conclusions within a more reasonable timeframe.

In conclusion, these insights offer a roadmap for the rational design of personalized treatment schedules based on precise measurements and mathematical modeling and suggest that a deeper understanding of the dynamic interplay between drug action, phenotypic plasticity, and recovery phases during treatment release could help isolate optimal conditions for treatment success, to inform more effective therapy strategies.

ACKNOWLEDGMENTS

This work was supported by Fondazione AIRC per la ricerca sul cancro ETS (AIRC IG 2019 Grant No. 23258 and AIRC IG 2024 Grant No. 30391—P.I. Marco Cosentino Lagomarsino). M.C. was supported by Fondazione AIRC per la ricerca sul cancro ETS (ID. 28177). S.P. was supported by Fondazione Umberto Veronesi (FundRef: 501100004710). We thank Edo Kussell, Johannes Berg, Einar Bjarki Gunnarsson, Mariangela Russo, Alberto Bardelli, Ciro Mercurio, Andrea Bertotti, and Silvia Marsoni for feedback and useful discussions.

M.C.L. and S.P. conceived the study and contributed with key ideas at different stages. M.C. and A.D.B. performed data analysis, model simulations, and analytical calculations. M.C., M.C.L., and S.P. wrote the paper. All authors read and approved the final version of the paper.

DATA AVAILABILITY

The data that support the findings of this article are openly available [43].

APPENDIX A: CONSIDERED COLORECTAL CANCER CELL LINES

In Ref. [12], the population dynamics of two microsatellite-stable human CRC cell lines were analyzed: the RAS/RAF wild-type “DiFi” and the BRAF V600E-mutated “WiDr.” These cell lines were treated with specific drug regimens: cetuximab alone at a dose of $100 \mu\text{g ml}^{-1}$ for DiFi, consistent with its approval for metastatic CRC patients lacking RAS and BRAF mutations, and a combination of cetuximab ($50 \mu\text{g ml}^{-1}$) with the BRAF inhibitor dabrafenib ($1 \mu\text{M}$) for WiDr, reflecting its potential effectiveness in BRAF-mutant CRC cases. For both models, individual clones were isolated, and their growth and drug sensitivity closely matched those of the parental populations.

In this study, we evaluated our model in four different parameter configurations, derived from fitting population dynamics data for four cloned populations: WiDr cl. B5, WiDr cl. B7, DiFi cl. B3, and DiFi cl. B6, based on analyses in Refs. [11,12].

To determine the parameters of untreated sensitive cells, we measured growth rates using growth assays under standard cell-culture conditions. For parameters under drug treatment, instead, we analyzed data from two types of drug-response

TABLE II. Adopted CRC cell lines and treatments [12].

CRC cell lines	Sensitivity	Treatment
DiFi ^a B3-B6	EGFR blockade	cetuximab ($100 \mu\text{g/ml}$)
WiDr ^b B5-B7	EGFR blockade	cetuximab ($50 \mu\text{g/ml}$)
	BRAF inhibition	dabrafenib ($1 \mu\text{M}$)

^aRAS/RAF wild-type.

^bBRAF V600E-mutated.

growth assays. The first was a dose-response assay, where CRC clones were exposed to increasing concentrations of targeted therapies. This assay was used to evaluate growth curves, defined as the number of live cells over time and drug concentration, and to quantify the interplay between cell growth and the transition to a persister state during drug treatment. The second was a single-dose assay, involving three weeks of exposure to a constant drug concentration. This assay revealed a biphasic killing curve, characterized by an initial rapid decline in the number of sensitive cells, followed by a slower phase indicative of persister emergence, resembling patterns observed in bacterial populations. The surviving persister fraction showed a gradual but measurable decline over time, suggesting a slow tendency for persisters to die.

All treatment data and the model fits are presented in SM Fig. 1 [17]. Inferred model’s parameters are summarized in Table I.

APPENDIX B: QUANTIFICATION OF RESISTANCE EFFECTS

This Appendix discusses the role of resistant cells in our model due to their potential impact on treatment outcomes. We adopt the stringent definition that resistant variants are entirely unaffected by treatment, with no fitness costs or competition with susceptible cells. This conservative approach excludes heteroresistance and competitive interactions during drug holidays, allowing for a clearer analysis of resistance dynamics. The emergence of resistance in our experimentally motivated setup can either occur during the free expansion of susceptible cells before the first treatment-release round (pre-existing resistance) or during the actual treatment-release cycles (*de novo* resistance). Both effects are strongly influenced by population size. Through analytical and stochastic simulations, we show that the emergence of resistance does not inevitably occur in periodic treatment-release cycles but rather depends on factors such as initial cell count, treatment duration, and drug concentration. Our analysis demonstrates that, for population sizes up to $N \approx 10^4$ – 10^5 cells, resistance is statistically negligible. The dominant factors shaping treatment efficacy remain the dynamics of susceptible and persister cells. Furthermore, along the optimal treatment strategies identified in this study, the likelihood of resistance developing remains extremely low, ensuring that our conclusions are robust to its presence. However, in clinically relevant scenarios, such as large tumors ($N > 10^5$ cells with our parameters) or heterogeneous subpopulations, resistance could become a key factor, introducing trade-offs between competing cell populations. Addressing these complexities lies beyond the scope

of this study but remains an important direction for future research.

1. Evaluation of the risk of pre-existing resistance

Under our hypotheses, if resistant cells pre-exist at treatment start, then resistance will inevitably take over, making any protocol optimization ineffective. We computed the risk of this event within our experimental parameters as a function of initial tumor size (see below), setting an upper bound to ensure pre-existing resistance remains statistically negligible. This risk is defined as the probability of having at least one resistant cell at treatment start (SM Fig. 2(a) [17])

$$P(T) = 1 - H(0, T), \quad (\text{B1})$$

where $H(0, T)$ is the probability of having no resistant cells emerging before time T . This probability can be computed analytically as a function of treatment size under given assumptions [44–46]. Assuming population growth is exponential and that resistant cells have the same fitness as untreated sensitive cells one can show that the probability of having resistant cells when treating a tumor of size N_0 reads

$$P(N_0) = 1 - \exp \left\{ -\gamma_s N_0 \frac{b}{d} \log \left(\frac{1 - \frac{d/b}{1-d/b} N_0^{-1}}{1 - d/b} \right) \right\}, \quad (\text{B2})$$

where b, d are the birth and death rates of resistant cells and γ_s is the mutation rate from sensitive to resistant cells. For fixed γ_s, b, d , the risk of pre-existing resistance increases with tumor size N_0 (SM Fig. 2(b) [17]). As can be seen, tumors below 10^5 cells have a risk under 1%, while for $N_0 \geq 10^7$, resistance is inevitable. This motivated our previous *in vitro* study at $\approx 10^4$ tumor cells [12].

2. Evaluation of the risk of *de novo* resistance

Even without pre-existing resistant cells, *de novo* resistance can develop during treatment-release cycles from mutations in susceptible and persister cells. While resistance will eventually arise in any treatment protocol with a positive long-term fitness, it may not always occur over long timescales when the protocol is effective ($\Lambda < 0$). In those cases, resistance can only develop if susceptible and persister cells mutate before the population goes extinct. Indeed, although the long-term fitness analysis requires $q \gg 1$, extinction can still occur within a finite (possibly large) q , depending on the initial tumor size. Thus, when $\Lambda < 0$, two competing timescales emerge: the extinction timescale, τ_e , given by

$$\langle \tau_e \rangle = \frac{\tau_r + \tau_t}{\Lambda} \log N_0, \quad (\text{B3})$$

and the mutation timescale, τ_m , which represents the average time for mutations to occur in susceptible and persister cells:

$$\langle M(\tau_m) \rangle = \int_0^{\tau_m} dt (\gamma_s S + \gamma_p P) = 1. \quad (\text{B4})$$

This mutation timescale is an underestimate, as mutations must also lead to the growth of a resistant colony. Here, we consider only the time for the first emergence and not

the drift dynamics before its establishment. We estimated the extinction and mutation timescales via numerical integration of the dynamical system in Eq. (1), incorporating mutation dynamics. Specifically, we calculated the characteristic extinction and mutation timescales of Eqs. (B3) and (B4) as the first “deterministic” passage time for $N = 1$ and $M = 1$, respectively, across the diagram in Fig. 3(a), starting with 10^4 cells (SM Fig. 3 [17]). We found that in all horizontal sections (fixed τ_r), either mutation or extinction dominates, with the winner population determined by the sign of the long-term fitness, in line with Fig. 3(a). Thus, from this deterministic analysis, we conclude that when $\Lambda < 0$, the development of resistance has no deterministic effect, motivating us to neglect resistant cells from the deterministic model in Eq. (1). Since resistance development is a stochastic process, we further addressed it using a stochastic approach. *De novo* resistance arises from mutations in susceptible and persister cells at rates γ_s and γ_p , respectively, with $\gamma_p = 10 \times \gamma_s$ (stress-induced mutagenesis) assumed conservatively (SM Fig. 4(a) [17]). Under these assumptions, we simulated stochastic trajectories $\{S(t), P(t), R(t), M(t)\}$ using the Gillespie algorithm [47], drawing transition rates from an exponential distribution and updating system states accordingly. Stochastic trajectories for parameters ($\tau_r = 2$ d, $\tau_t = 7$ d, $c = 1$ μM , $N_0 = 10^4$) are shown in SM Fig. 4(b) [17]. Consistent with a predicted negative long-term fitness [Fig. 3(a)], susceptible and persister cells go extinct in most realizations; however, in 2 of 300 simulated realizations, resistance developed, leading to tumor relapse. When *de novo* resistance occurs, the long-term fitness estimate no longer accurately describes the population dynamics. Thus, we focused on two observables: (i) the probability of developing *de novo* resistance before extinction, estimated as the fraction of realizations where resistance developed, with a threshold of 5% for negligible resistance, and (ii) the expected number of resistance mutations before extinction, estimated as the average number of mutations across realizations. Our analysis shows that within the relevant parameter range, resistance development is statistically negligible and does not significantly affect the long-term fitness estimate when $\Lambda < 0$. For instance, in SM Figs. 4(c) and 4(d) [17], the risk of resistance increases with initial tumor size [Fig. 4(c)] but remains below 5% for sizes up to 20–30k cells. For larger sizes, such as 10^5 cells, the risk reaches 20%, leading to a slight deviation in long-term fitness (one in five realizations). SM Fig. 4(d) [17] confirms that treatment duration does not clearly affect resistance risk. Further validation in SM Figs. 4(e) and 4(f) [17] shows that the expected number of resistant mutations remains below 1 across all parameters tested, emphasizing that resistance is statistically negligible and should be excluded from a deterministic model for cancer evolution in these conditions. To conclude, while we acknowledge the potential implications of resistance on treatment outcomes, we have provided sufficient evidence supporting the idea that in the adverse scenario of stress-induced mutagenesis and no competitive disadvantage of resistance, resistance has little effect as long as population size is lower than 10^4 – 10^5 cells. Indeed, whenever $\Lambda > 0$, resistance only accelerates the adverse scenario, and where $\Lambda < 0$, where our optimization analysis applies, resistance is statistically negligible.

APPENDIX C: QUANTIFICATION OF STOCHASTIC EFFECTS

This Appendix addresses the role of stochasticity in our predictions. Our estimation method of the long-term fitness is based on deterministic population dynamics. However, since the underlying transitions (including cell birth-death events) are inherently stochastic, the long-term fitness is itself a stochastic variable and trajectory dependent. These stochastic effects should emerge for small populations. To quantify the variability in the estimate of the long-term fitness introduced by these stochastic effects, we used the Gillespie algorithm to sample several stochastic trajectories of the Markov process $\{S(t), P(t)\}_t$ defined by the following transitions:

Release:

$(S, P) \rightarrow (S + 1, P)$; susceptible growth, rate: $\mu_s^r S$

$(S, P) \rightarrow (S, P + 1)$; persisters growth, rate: $\mu_p^r P$

$(S, P) \rightarrow (S + 1, P - 1)$; switch-back, rate: $\lambda_p^r P$

Treatment:

$(S, P) \rightarrow (S - 1, P)$; susceptible death, rate: $\mu_s^t S$

$(S, P) \rightarrow (S, P - 1)$; persisters growth, rate: $\mu_p^t P$

$(S, P) \rightarrow (S - 1, P + 1)$; S-to-P transition, rate: $\lambda_s^t S$

Stochastic trajectories generated with this method are shown in SM Fig. 16(a) [17], where they are compared to the deterministic dynamics (solid black line). Stochastic effects become significant for small populations, especially under treatment conditions, but they do not qualitatively affect the overall behavior (SM Figs. 16(a)–16(d) [17]). The relative error in estimating long-term fitness remains below 10% for the population sizes considered and decreases with larger populations (SM Fig. 16(d) [17]). Therefore, the conclusions drawn regarding treatment success and failure regions in the phase diagram remain robust. Additionally, we evaluated the probability of treatment success, defined as the fraction of realizations leading to extinction within a given time frame. For protocols with negative long-term fitness, this probability increases over time (SM Fig. 16(e) [17]). Importantly, by varying parameters like drug concentration and release timing, we identified optimal values that maximize treatment success, which align closely with the predictions from the deterministic model (SM Figs. 16(f) and 16(g) [17]). In conclusion, while the inherent stochasticity in the transition rates that define our dynamical system [Eq. (1)] introduces variability in the outcomes, it does not qualitatively change the key findings of our study. The deterministic model provides a reliable approximation of the systems behavior in the parameter ranges we tested.

APPENDIX D: UNIVERSAL ASPECTS OF THE MODEL

This Appendix provides a more general description of our model, identifying “universal” features that extend beyond the specific model parameters studied in this work. In doing this, we also contextualize our assumptions and findings within such broader picture.

1. Rescaled dimensionless parameters

First, we note that the evolution matrix, $G = e^{M^r \tau_r} e^{M^t \tau_t}$, which governs the system's dynamics for a given drug concentration c [Eq. (7)], is a function of the eight rates $(\mu_s^r, \lambda_s^r, \mu_p^r, \lambda_p^r, \mu_s^t, \lambda_s^t, \mu_p^t, \lambda_p^t)$, defining the dynamical transitions in our model, and the two timescales (τ_r, τ_t) , defining the treatment-release protocol. Given the structure of the problem, 10 parameters and one dimension (time), we anticipate 9 dimensionless parameters: 8 rescaled parameters and the long-term fitness. These 8 parameters correspond to specific combinations of the 10 original parameters. Although determining the exact combinations may not be trivial in principle, in the following, we will show that the decomposition of the evolution matrix G in terms of the eigenvalues $(m_\pm^r, m_\pm^t)$ and the eigenvectors $(R = (\mathbf{r}^+, \mathbf{r}^-), T = (\mathbf{t}^+, \mathbf{t}^-))$ of the rest and treatment matrices will provide a natural basis for the rescaled dimensionless parameters. Specifically, we note that the evolution matrix can be decomposed as

$$G = R \begin{pmatrix} e^{m_+^r \tau_r} & 0 \\ 0 & e^{m_-^r \tau_r} \end{pmatrix} R^{-1} T \begin{pmatrix} e^{m_+^t \tau_t} & 0 \\ 0 & e^{m_-^t \tau_t} \end{pmatrix} T^{-1}, \quad (D1)$$

with

$$m_\pm^{r/t} = \frac{(\mu_s^{r/t} - \lambda_s^{r/t} + \mu_p^{r/t} - \lambda_p^{r/t})}{2} \pm \frac{|\mu_s^{r/t} - \lambda_s^{r/t} - \mu_p^{r/t} + \lambda_p^{r/t}|}{2} \times \sqrt{1 + \frac{4\lambda_s^{r/t}\lambda_p^{r/t}}{(\mu_s^{r/t} - \lambda_s^{r/t} - \mu_p^{r/t} + \lambda_p^{r/t})^2}} \quad (D2)$$

and

$$R \equiv \begin{pmatrix} \frac{1}{\sqrt{1+\gamma_+^r}} & \frac{1}{\sqrt{1+\gamma_-^r}} \\ \frac{-\gamma_+^r}{\sqrt{1+\gamma_+^r}} & \frac{-\gamma_-^r}{\sqrt{1+\gamma_-^r}} \end{pmatrix}, \quad T \equiv \begin{pmatrix} \frac{1}{\sqrt{1+\gamma_+^t}} & \frac{1}{\sqrt{1+\gamma_-^t}} \\ \frac{-\gamma_+^t}{\sqrt{1+\gamma_+^t}} & \frac{-\gamma_-^t}{\sqrt{1+\gamma_-^t}} \end{pmatrix}, \quad (D3)$$

where

$$\gamma_\pm^{r/t} = \frac{\lambda_s^{r/t}}{\mu_p^{r/t} - \lambda_p^{r/t} - m_\pm^{r/t}} = \frac{\mu_s^{r/t} - \lambda_s^{r/t} - m_\pm^{r/t}}{\lambda_p^{r/t}}. \quad (D4)$$

As a result, the evolution matrix G can be fully expressed in terms of eight rescaled, dimensionless parameters:

$$\epsilon_\pm^{r/t} \equiv e^{m_\pm^{r/t} \tau_{r/t}}, \quad (D5)$$

$$\gamma_\pm^{r/t} \equiv \frac{\lambda_s^{r/t}}{\mu_p^{r/t} - \lambda_p^{r/t} - m_\pm^{r/t}} = \frac{\mu_s^{r/t} - \lambda_s^{r/t} - m_\pm^{r/t}}{\lambda_p^{r/t}}. \quad (D6)$$

Specifically, we found

$$G_{11} = \frac{1}{\gamma_+^r - \gamma_-^r} \frac{1}{\gamma_+^t - \gamma_-^t} [\epsilon_+^r \epsilon_+^t (\gamma_-^r - \gamma_+^t) \gamma_-^t - \epsilon_+^r \epsilon_-^t (\gamma_-^r - \gamma_-^t) \gamma_+^t - \epsilon_-^r \epsilon_+^t (\gamma_+^r - \gamma_+^t) \gamma_-^t + \epsilon_-^r \epsilon_-^t (\gamma_+^r - \gamma_-^t) \gamma_+^t], \quad (D7)$$

$$G_{12} = \frac{1}{\gamma_+^r - \gamma_-^r} \frac{1}{\gamma_+^t - \gamma_-^t} [\epsilon_+^r \epsilon_+^t (\gamma_-^r - \gamma_+^t) - \epsilon_+^r \epsilon_-^t (\gamma_-^r - \gamma_-^t) - \epsilon_-^r \epsilon_+^t (\gamma_+^r - \gamma_+^t) + \epsilon_-^r \epsilon_-^t (\gamma_+^r - \gamma_-^t)], \quad (D8)$$

$$G_{21} = \frac{1}{\gamma_+^r - \gamma_-^r} \frac{1}{\gamma_+^t - \gamma_-^t} [-\epsilon_+^r \epsilon_+^t (\gamma_-^r - \gamma_+^r) \gamma_+^r \gamma_-^t + \epsilon_+^r \epsilon_-^t (\gamma_-^r - \gamma_+^r) \gamma_+^r \gamma_+^t + \epsilon_-^r \epsilon_+^t (\gamma_+^r - \gamma_-^r) \gamma_-^r \gamma_-^t - \epsilon_-^r \epsilon_-^t (\gamma_+^r - \gamma_-^r) \gamma_-^r \gamma_+^t], \quad (\text{D9})$$

$$G_{22} = \frac{1}{\gamma_+^r - \gamma_-^r} \frac{1}{\gamma_+^t - \gamma_-^t} [-\epsilon_+^r \epsilon_+^t (\gamma_-^r - \gamma_+^r) \gamma_+^r + \epsilon_+^r \epsilon_-^t (\gamma_-^r - \gamma_+^r) \gamma_+^r + \epsilon_-^r \epsilon_+^t (\gamma_+^r - \gamma_-^r) \gamma_-^r - \epsilon_-^r \epsilon_-^t (\gamma_+^r - \gamma_-^r) \gamma_-^r]. \quad (\text{D10})$$

We note that the rescaled parameters defined in Eqs. (D5) and (D6) are related to the four distinct dynamical modes during the release/treatment phases, i.e., $e^{m_{\pm}^{r/t}} \left[\frac{1}{\sqrt{1+\gamma_{\pm}^{r/t}}} \frac{-\gamma_{\pm}^{r/t}}{\sqrt{1+\gamma_{\pm}^{r/t}}} \right]^t$. During each treatment or release phase the population vector, $\mathbf{N} = [S, P]^t$, eventually scales as $\mathbf{N} \sim e^{m_{\pm}^{r/t}} \left[\frac{1}{\sqrt{1+\gamma_{\pm}^{r/t}}} \frac{-\gamma_{\pm}^{r/t}}{\sqrt{1+\gamma_{\pm}^{r/t}}} \right]^t$, so that $m_{\pm}^{r/t}$ represent the asymptotic exponential growth rate, whereas $\gamma_{\pm}^{r/t}$ is the asymptotic persister-to-susceptible subpopulations ratio [48].

2. Identification of distinct physical regimes

Each entry of the G matrix [Eqs. (D7)–(D10)] can be written as a linear combination of products of exponential factors, $\epsilon_{\pm}^r \epsilon_{\pm}^t$, with coefficients being given functions of the $\gamma_{\pm}^{r/t}$ parameters, that is,

$$G_{ij} = f_{ij}^{++}(\gamma_{\pm}^{r/t}) \epsilon_+^r \epsilon_+^t + f_{ij}^{+-}(\gamma_{\pm}^{r/t}) \epsilon_+^r \epsilon_-^t + f_{ij}^{-+}(\gamma_{\pm}^{r/t}) \epsilon_-^r \epsilon_+^t + f_{ij}^{--}(\gamma_{\pm}^{r/t}) \epsilon_-^r \epsilon_-^t. \quad (\text{D11})$$

Importantly, although $\epsilon_+^r \epsilon_+^t$ is the highest exponential factor, the leading term in each entry also depends on the coefficients $f_{ij}^{\pm\pm}(\gamma_{\pm}^{r/t})$, which are given functions of the lengths of the treatment and release phases, as well as of the value of the drug concentration. In other words, the leading term in each entry of the evolution matrix, G , is *a priori* protocol dependent. Therefore, in the general case where all terms in the linear combination [Eq. (D11)] are non-negligible, the evolution matrix exhibits highly nontrivial behavior across the space of all possible protocols defined by (τ_t, τ_r, c) . This complexity presents a significant challenge for a fully analytical treatment of G . Consequently, previous studies have identified distinct physical regime-specific regions in parameter space where appropriate approximations and simplifications of G enable analytical tractability in problems like the estimation of the long-term fitness from G (Ref. [5], see also Appendix F). Here we limit ourselves to report in SM Figs. 18 and 19 [17] the complex and variable behavior of each entry of the evolution matrix, G , over the protocol spaces $(\tau_t, \tau_r, c = \text{const.})$ and $(\tau_t = \text{const.}, \tau_r, c)$ in the drug-induced limit discussed below, where some of the coefficients f_{ij} in the linear combination are set to zero, canceling the contribution of the corresponding exponential factor.

3. The drug-induced limit

To reflect the experimentally observed behavior of CRC persister cells [12], we assumed in this work that $\lambda_s^r = \lambda_p^r = 0$ and that persisters do not grow during the release phase, i.e., $\mu_p^r = 0$. In this limit, which we call the “drug-induced limit,” the rescaled parameters read [assuming $\mu_s^t(c) - \lambda_s^t(c) - \mu_p^t(c) < 0$]

$$\epsilon_+^r = e^{\mu_s^r \tau_r}, \quad \epsilon_-^r = e^{-\lambda_p^r \tau_r}, \quad \epsilon_+^t = e^{\mu_p^t \tau_t}, \quad \epsilon_-^t = e^{(\mu_s^t - \lambda_s^t) \tau_t}$$

$$\gamma_+^r = 0, \quad \gamma_-^r = \frac{\mu_s^r + \lambda_p^r}{\lambda_p^r} \equiv B^{-1}, \quad \gamma_+^t \rightarrow \infty, \quad (\text{D12})$$

$$\gamma_-^t = \frac{\lambda_s^t}{\mu_p^t - (\mu_s^t - \lambda_s^t)} \equiv -A.$$

Therefore, the drug-induced regime corresponds to the limit $\gamma_+^r \rightarrow 0$ and $\gamma_+^t \rightarrow \infty$ in the eight-dimensional space of rescaled parameters. In this regime, two parameters λ_s^r and λ_p^t are set to zero, reducing the system’s dynamics to the six remaining rescaled parameters: $\epsilon_{\pm}^{r/t}, A, B$. It is worth noting that, although similar dynamical regimes have been explored previously by varying treatment and release durations [5], the drug-induced limit considered here is conceptually distinct, as it arises from specific assumptions about the parameter values rather than from considering specific treatment protocols. By taking the limit $\gamma_+^r \rightarrow 0, \gamma_+^t \rightarrow \infty$ in Eqs. (D7)–(D10), we find that the expression of the evolution matrix for the drug-induced scenario reads

$$G_{11}^{\text{DI}} = \epsilon_+^r \epsilon_-^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t + \epsilon_-^r \epsilon_-^t)$$

$$G_{12}^{\text{DI}} = B(\epsilon_+^r \epsilon_+^t - \epsilon_-^r \epsilon_+^t)$$

$$G_{21}^{\text{DI}} = -A(\epsilon_-^r \epsilon_+^t - \epsilon_-^r \epsilon_-^t)$$

$$G_{22}^{\text{DI}} = \epsilon_-^r \epsilon_+^t. \quad (\text{D13})$$

4. Behavior close to optimality

Building upon the general expression for the evolution matrix derived in the drug-induced regime [Eq. (D13)], we now focus on the close-to-optimality behavior of the long-term fitness, Λ^{DI} , defined as the logarithm of the highest eigenvalue of G^{DI} . Specifically, we consider a parameter configuration close to the optimal protocol solution outlined in the main text. These rescaled optimal parameters are mathematically defined by the numerical values reported in Table III. In the following, we will show that the specific values assumed by these parameters can be exploited to derive an approximated analytical expression for the long-term fitness valid near optimality. As can be seen by looking at Table I, close to the optimal parameters’ configuration the term $\epsilon_-^r \epsilon_-^t = 0$ can be set to 0 without any loss of information. Furthermore, since $\delta \approx 1$ for all clones (Table I), in the evaluation of the eigenvalues of G^{DI} we can approximate $\sqrt{1 + 4G_{12}G_{21}/(G_{11} - G_{22})^2} \approx (1 + 2G_{12}G_{21}/(G_{11} - G_{22})^2)$, yielding the two eigenvalues

$$g_1 = G_{11}^{\text{DI}} + \frac{G_{12}^{\text{DI}} G_{21}^{\text{DI}}}{G_{11}^{\text{DI}} - G_{22}^{\text{DI}}}, \quad g_2 = G_{22}^{\text{DI}} - \frac{G_{12}^{\text{DI}} G_{21}^{\text{DI}}}{G_{11}^{\text{DI}} - G_{22}^{\text{DI}}}. \quad (\text{D14})$$

TABLE III. Rescaled dimensionless parameters at optimality for the four CRC cell lines.

Cell line	ϵ_+^r	ϵ_-^r	ϵ_+^t	ϵ_-^t	B	A	AB
WiDr cl. B7	9.2	0.8	0.6	0.006	0.1	-0.08	-0.008
WiDr cl. B5	25	0.6	1	0.012	0.12	-0.01	-0.001
DiFi cl. B6	3.3	0.9	0.8	0.003	0.08	-0.29	-0.02
DiFi cl. B3	2.2	0.8	0.6	0.011	0.17	-0.40	-0.07
Cell line	$\epsilon_+^r \epsilon_+^t$	$\epsilon_+^r \epsilon_-^t$	$\epsilon_-^r \epsilon_+^t$	$\epsilon_-^r \epsilon_-^t$	δ^a		
WiDr cl. B7	5.5	0.06	0.48	0.005	0.98		
WiDr cl. B5	25	0.3	0.6	0.007	0.96		
DiFi cl. B6	2.6	0.01	0.72	0.003	0.99		
DiFi cl. B3	1.3	0.02	0.48	0.009	0.97		

$$^a \delta \equiv \sqrt{1 + 4G_{12}G_{21}/(G_{11} - G_{22})^2}/(1 + 2G_{12}G_{21}/(G_{11} - G_{22})^2).$$

As a consequence, we can write the following near-optimum analytical approximation of the long-term fitness

$$\Lambda^{\text{DI}} \approx \log(\max\{g_1, g_2\}) \\ = \log\left(\frac{\epsilon_-^r \epsilon_+^t [(1 + AB)\epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t]}{\epsilon_+^r \epsilon_-^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t) - \epsilon_-^r \epsilon_+^t}\right), \quad (\text{D15})$$

where the last equality holds since close to the optimum, g_2 is always the higher eigenvalues, Eq. (D15) is reported in SM Figs. 14(a), 14(c), 14(e), and 14(g) (blue line) [17] together with the exact long-term fitness estimate (black line), confirming the validity of the approximation around the optimum. Importantly, we can use Eq. (D15) to compute $\frac{\partial \Lambda}{\partial c} = 0$ and $\frac{\partial \Lambda}{\partial \tau_r} = 0$ and gain some physical intuition behind the optimal parameters' configuration. We find that the optimal values of the drug concentration and the release time are such that

$$\frac{d\epsilon_-^t/dc}{dA/dc} = K_c, \quad (\text{D16})$$

and

$$\frac{d\epsilon_+^t/d\tau_r}{d\epsilon_-^t/d\tau_r} = K_r, \quad (\text{D17})$$

where K_c and K_r are parameter-dependent constants taking the following values at optimality, $K_c \in \{0.58, 1.39, 0.71, 0.57\}$ and $K_r \in \{-84, -177, -34, -8.7\}$ for WiDr cl. B7, WiDr cl. B5, DiFi cl. B6, and DiFi cl. B3, respectively. Equation (D16) expresses a balance between the effective rate of susceptible cell elimination and their transition to persisters. Likewise, Equation (D17) describes a trade-off between the effective growth rate of susceptible cells and the rate at which persisters switch back. The optimal solutions defined by the balance conditions [Eqs. (D16) and (D17)] are reported in SM Figs. 14(b), 14(d), 14(f), and 14(h) [17] together with the exact optimal solutions, supporting the validity of our general analytical approach around optimality.

APPENDIX E: ROBUSTNESS OF RESULTS TO PARAMETER PERTURBATIONS

This Appendix describes the sensitivity analysis we performed to test the robustness of our conclusions. Specifically,

we estimated the long-term fitness under small perturbations of each parameter around its reported value in Table I. We modeled the perturbations as additive Gaussian random variables with zero mean and standard deviation equal to 20% of the reference parameter value. For parameters equal to 0 in the model, we set the noise to a level compatible with the typical values assumed by the other model parameters. This approach accounts for variability in key transitions, such as persistence induction and switch-back, potentially shifting the analysis from a strictly drug-induced response towards a qualitatively distinct bet-hedging-like scenario. SM Fig. 20 [17] shows the results of this analysis, reported as the behavior of the long-term fitness as a function of release duration and drug concentration in both the reference condition (SM Figs. 20(a) and 20(b) [17]) and parameter-perturbed conditions (SM Figs. 20(c) and 20(d) [17]). Importantly, we find that our findings remain robust over the described perturbations of the model parameters. We reasoned that the robustness of our results should be due to the fact that fluctuations in the parameters were such that $\frac{4\lambda_s^{r/t}\lambda_p^{r/t}}{(\mu_s^{r/t} - \lambda_s^{r/t} - \mu_p^{r/t} + \lambda_p^{r/t})^2} \ll 1$ in the eigenvalue equation [Eq. (D2)], implying a behavior of the system close to the one observed under the drug-induced hypothesis ($\lambda_s^r = \lambda_p^t = 0$), which sets this quantity exactly equal to 0. We confirmed this reasoning by computing the quantity $\frac{4\lambda_s^{r/t}\lambda_p^{r/t}}{(\mu_s^{r/t} - \lambda_s^{r/t} - \mu_p^{r/t} + \lambda_p^{r/t})^2}$ for all tested parameter realizations and showing that even strong perturbations did not violate this condition (SM Figs. 20(e) and 20(f) [17]).

APPENDIX F: ESTIMATION OF LONG-TERM FITNESS AND ASYMPTOTIC POPULATION COMPOSITION

This Appendix discusses different equivalent ways of evaluating the treated population's long-term fitness:

$$\Lambda \equiv \lim_{q \rightarrow \infty} \frac{1}{q} \log N(q), \quad (\text{F1})$$

and asymptotic composition:

$$\Psi \equiv \lim_{q \rightarrow \infty} \frac{S(q) - P(q)}{S(q) + P(q)}. \quad (\text{F2})$$

1. Method 1: Estimation via dynamics simulations

One first possibility for estimating the long-term fitness, exploiting the actual definition of this parameter, is to simulate the dynamical system of Eq. (1) over many treatment-release cycles and then compute the long-term fitness as the asymptotic slope of the $\log N$ vs q plot:

$$\Lambda^{\text{simul.}} = \log\left(\frac{N(q)}{N(q-1)}\right), \quad q \gg 1. \quad (\text{F3})$$

Indeed, in the long-term limit $q \gg 1$, $N(q) \propto e^{\Lambda q}$, $N(q-1) \propto e^{\Lambda(q-1)}$, so that $N(q)/N(q-1) \propto e^{\Lambda}$. Asymptotic population composition is straightforwardly obtained from the definition, Eq. (F2), as

$$\Psi^{\text{simul.}} = \frac{S(q) - P(q)}{S(q) + P(q)}, \quad q \gg 1. \quad (\text{F4})$$

We used the `scipy.integrate` library to integrate the dynamical system of Eq. (1) over many treatment-release cycles. We

then used Eqs. (F3) and (F4) to evaluate long-term fitness and asymptotic population composition values.

2. Method 2: Estimation via numerical diagonalization

This method considers the solution of the dynamical system of Eq. (1), that is Eq. (7), and aims to find the eigenvalues g_+ , g_- and eigenvectors $\mathbf{g}^+ = (g_1^+, g_2^+)$, $\mathbf{g}^- = (g_1^-, g_2^-)$ of the evolution matrix $G \equiv e^{M^r \tau_r} e^{M^t \tau_t}$ to write $\mathbf{N}(q) = G^q \mathbf{N}(0)$ in the form

$$\begin{aligned} \mathbf{N}(q) &= \begin{pmatrix} g_1^+ & g_1^- \\ g_2^+ & g_2^- \end{pmatrix} \begin{pmatrix} g_+^q & 0 \\ 0 & g_-^q \end{pmatrix} \begin{pmatrix} g_1^+ & g_1^- \\ g_2^+ & g_2^- \end{pmatrix}^{-1} \begin{pmatrix} S_0 \\ 0 \end{pmatrix} \\ &= \frac{S_0 g_+^q}{g_1^+ g_2^- - g_2^+ g_1^-} \begin{pmatrix} g_1^+ g_2^- - g_1^- g_2^+ \left[\frac{g_-}{g_+} \right]^q \\ g_2^+ g_2^- - g_2^- g_1^+ \left[\frac{g_-}{g_+} \right]^q \end{pmatrix}, \end{aligned} \quad (\text{F5})$$

where we used the drug-induced hypothesis $P_0 = 0$. In the long-term limit $q \gg 1$ one has:

$$S(q) = S_0 g_+^q \frac{g_1^+ g_2^-}{g_1^+ g_2^- - g_2^+ g_1^-}, \quad (\text{F6})$$

$$P(q) = S_0 g_+^q \frac{g_2^+ g_2^-}{g_1^+ g_2^- - g_2^+ g_1^-}, \quad (\text{F7})$$

$$N(q) = S_0 g_+^q \frac{(g_1^+ + g_2^+) g_2^-}{g_1^+ g_2^- - g_2^+ g_1^-}. \quad (\text{F8})$$

$$G^{\text{DI}} = \begin{pmatrix} \epsilon_+^r \epsilon_-^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t + \epsilon_-^r \epsilon_-^t) & B(\epsilon_+^r \epsilon_+^t - \epsilon_-^r \epsilon_+^t) \\ -A(\epsilon_-^r \epsilon_+^t - \epsilon_-^r \epsilon_-^t) & \epsilon_-^r \epsilon_+^t \end{pmatrix}. \quad (\text{F11})$$

Eigenvalues and eigenvectors of G can therefore be computed as:

$$g_{\pm} = \frac{1}{2} \text{Tr}(G) \pm \frac{1}{2} \sqrt{\text{Tr}(G)^2 - 4 \det(G)}, \quad (\text{F12})$$

and

$$\mathbf{g}^{\pm} = \begin{pmatrix} \frac{1}{\sqrt{1+\gamma_{\pm}^2}} \\ \frac{-\gamma_{\pm}}{\sqrt{1+\gamma_{\pm}^2}} \end{pmatrix}, \quad (\text{F13})$$

where

$$\text{Tr}(G) = \epsilon_+^r \epsilon_-^t + \epsilon_-^r \epsilon_+^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t + \epsilon_-^r \epsilon_-^t), \quad (\text{F14})$$

$$\det(G) = \epsilon_+^r \epsilon_-^t \epsilon_-^r \epsilon_+^t - AB(\epsilon_+^r \epsilon_+^t \epsilon_-^r \epsilon_-^t - \epsilon_+^r \epsilon_-^t \epsilon_-^r \epsilon_+^t), \quad (\text{F15})$$

and

$$\gamma_{\pm} \equiv \frac{\epsilon_+^r \epsilon_-^t - AB(\epsilon_+^r \epsilon_+^t - \epsilon_+^r \epsilon_-^t - \epsilon_-^r \epsilon_+^t + \epsilon_-^r \epsilon_-^t) - \epsilon_+^r \epsilon_+^t \omega_{\pm}}{B(\epsilon_+^r \epsilon_+^t - \epsilon_-^r \epsilon_+^t)}. \quad (\text{F16})$$

Therefore, substituting these quantities into Eqs. (F9) and (F10) we find

$$\Lambda^{\text{theor.}} = \log \left(\frac{1}{2} \text{Tr}(G) + \frac{1}{2} \sqrt{\text{Tr}(G)^2 - 4 \det(G)} \right), \quad (\text{F17})$$

and

$$\Psi^{\text{theor.}} = \frac{1 + \gamma_+}{1 - \gamma_+}. \quad (\text{F18})$$

Therefore, using Eqs. (F1) and (F2), one finds

$$\Lambda^{\text{numer.}} = \log g_+, \quad (\text{F9})$$

which corresponds to Eq. (8) in the main text and

$$\Psi^{\text{numer.}} = \frac{g_1^+ - g_2^+}{g_1^+ + g_2^+} \quad (\text{F10})$$

which is Eq. (10) in the main text.

We used the `numpy.linalg` library to compute the eigenvectors and eigenvalues of the matrix $G \equiv e^{M^r \tau_r} e^{M^t \tau_t}$ and then used Eqs. (F9) and (F10) to evaluate long-term fitness and asymptotic population composition values.

3. Method 3: Analytical estimation

From the analytical point of view, diagonalizing directly G is a daunting task, mostly because G is defined via exponentiation of treatment and release matrices. However, we showed in Appendix D how this problem can be tackled analytically by exploiting the eigenvalues and the eigenvectors of the release and treatment evolution matrices. In the drug-induced limit describing our experiments, G can be written as [Eq. (D13)]

We point out that the analytical estimates derived in this section [Eqs. (F17) and (F18)] are exact. In fact, we verified that the different long-term fitness and asymptotic population composition estimation methods discussed in this section, summarized by Eqs. (F3), (F9), and (F17) and Eqs. (F4), (F10), and (F18), are all consistent and provide the same quantities (SM Fig. 5 [17]).

4. Relation with previous work

An exact analytical expression for the long-term fitness was previously derived by Patra and Klumpp [3], using a different analytical approach. Since Eq. (F17) presented here is also exact, it can be considered an equivalent solution specific to our drug-induced scenario. Approximated analytical solutions valid in specific physical regimes have also been reported in earlier studies. For instance, Kussell and co-workers [5] provided an approximated analytical solution for the long-term fitness valid for large populations, referred to as the “deterministic limit” in Ref. [5]).

We already discussed in Appendix D how distinct physical regimes can be individuated in our framework by looking at the leading terms in each entry of the evolution matrix G . An equivalent approximated solution, valid for large populations and for the drug-induced scenario, can therefore be derived factoring out the term $\epsilon_+^r \epsilon_+^t$ from Eq. (D13), which for long treatment and release durations, $\tau_t, \tau_r \gg 1$, dominate all other

contributions. Put it differently, we can approximate G^{DI} as

$$G^{\text{DI}} \approx \epsilon_+^r \epsilon_+^t \begin{pmatrix} -AB & B \\ 0 & 0 \end{pmatrix}, \quad (\text{F19})$$

yielding $g_+ \approx -\epsilon_+^r \epsilon_+^t AB$. The approximated formula for the long-term fitness therefore reads

$$\begin{aligned} \Lambda^{\text{approx}} &= \log(\epsilon_+^r \epsilon_+^t) + \log(-AB) \\ &= \mu_s^r \tau_r + \mu_p^t \tau_t + \log(\lambda_s^t \lambda_p^r) \\ &\quad + \log\left(\frac{1}{\mu_p^t + \lambda_s^t - \mu_s^r \mu_s^r + \lambda_p^r}\right), \end{aligned} \quad (\text{F20})$$

which is equivalent to Eq. (8) in Ref. [5] within our model regime. We verified the validity of this approximation by comparing it with exact estimation methods.

APPENDIX G: ALTERNATIVE METRICS FOR TREATMENT EFFICACY

This Appendix discusses alternative criteria and time units for evaluating treatment outcome. In the main text, we considered the per-cycle long-term fitness as the primary measure of treatment efficacy. Here we explore how our conclusions are affected when alternative metrics, such as long-term fitness measured in d^{-1} [22] (rather than 1/cycles) or the average tumor size after a fixed period, defined either in treatment cycles or absolute time, are considered instead of the per-cycle long-term fitness.

1. Justification of the choice of using treatment cycles as time units

Several reasons support the use of treatment cycles as the preferred unit: (i) Absolute time is typically not a clinical constraint, as treatment is often discussed in terms of therapy cycles, and (ii) therapy cycles are controllable in clinical settings, where drug pharmacokinetics are predefined for each patient and cannot be altered at will. For example, one cannot arbitrarily choose the duration of drug exposure in a patient if the pharmacokinetics dictate a longer duration. However, it is possible to wait for the drug's effect to wane, thus defining the end of a treatment cycle; (iii) the cycle unit allows for a natural comparison of intermittent treatments with continuous ones. Specifically, comparing a continuous treatment protocol and an intermittent one with the same treatment time and dose by comparing protocols with $\tau_r = 0$ (continuous) and $\tau_r > 0$ (intermittent) ensures that the overall drug exposure remains constant, even when the waiting time between administrations varies. Moreover, if the delivered drug dose, then c is the same for both protocols (as in the analyses reported in Fig. 3), the cumulative drug dose delivered to the patient remains constant (as illustrated in Fig. 3). Thus, the “per-cycle comparison” makes sense, even though it may appear asymmetric in absolute time. In contrast, comparing in absolute time by comparing a continuous protocol with a treatment time of $\tau_t' = \tau_t + \tau_r > \tau_t$ would imply that the continuous protocol delivers a higher overall drug exposure and cumulative dose. This is exactly the case in the analyses of Fig. 5, where we compare continuous and intermittent treat-

ment protocols with the same total dose profile but different drug administration modes. In this situation, the “per-cycle” comparison feels more appropriate, as the continuous protocol would deliver substantially more drugs in absolute time.

2. Comparison of results using units of time

We performed a detailed analysis of the differences in the results between using the $\frac{1}{\text{cycle}}$ versus $\frac{1}{\text{time}}$ units. SM Fig. 15 [17] compares the results of analyses from Figs 3(a) (SM Figs. 15(a)–15(c) [17]) and 4(a) (SM Figs. 15(e)–15(g) [17]) of the main text, using both time unit choices for the WiDr cl. B7 clone, highlighting similarities and differences in optimizing long-term fitness across the two units. As expected, the unit choice does not affect the identification of treatment failure ($\Lambda > 0$) and success ($\Lambda < 0$) regions. However, using absolute time may have different implications for determining the optimal protocol. This is particularly true for the WiDr B7 clone, as shown in SM Fig. 15 [17], but not for other clones (see SM Fig. 21 [17]). For example, comparing SM Figs. 15(a)–15(e) [17] with SM Figs. 15(c)–15(g) [17], we see that in absolute time there is no clear optimal release time. Indeed, solving Eq. (11) for long-term fitness in the $\frac{1}{\text{time}}$ unit, using WiDr cl. B7 parameters, yields the optimal protocol ($\tau_r = 7$ d, $\tau_r = 0$ d, $c = 2 \times 10^{-7}$ M), whereas for the $\frac{1}{\text{cycle}}$ unit, the optimal protocol is ($\tau_t = 7$ d, $\tau_r = 3.3$ d, $c = 3.8 \times 10^{-7}$ M). To clarify our optimization analysis, our goal was to identify the most effective strategy for accelerating population extinction (minimizing long-term fitness) while maintaining clinically meaningful and accessible protocol parameters. We constrained our analysis to clinically relevant parameter ranges, specifically $[0, 7] \times [0, 7] \times [10^{-8}, 10^{-4}]$. The optimal protocol depends critically on the specific extinction objective. If the aim is to minimize the average total number of cycles to extinction, then the optimal protocol is $P1 = (\tau_t = 7$ d, $\tau_r = 3.3$ d, $c = 3.8 \times 10^{-7}$ M). Conversely, if the goal is to minimize the total absolute time to population extinction, then the optimal protocol becomes $P2 = (\tau_t = 7$ d, $\tau_r = 0$ d, $c = 2 \times 10^{-7}$ M). It is worth noting that while the observation of an optimal release duration is consistent across all clones when using the $\frac{1}{\text{cycle}}$ unit, the presence or absence of an optimal release duration in the $\frac{1}{\text{time}}$ unit depends on the specific population-dynamics parameters considered [22]. For instance, with the WiDr cl. B5 parameters, an optimal release time is found even in the $\frac{1}{\text{time}}$ unit (SM Fig. 21 [17]). The optimal protocol derived from Eq. (11) of the main text for this case is $P3 = (\tau_t = 5$ d, $\tau_r = 3.3$ d, $c = 4.8 \times 10^{-7}$ M). This result is not surprising, as intermittent treatments are crucial for eradicating immortal or very slowly dying persisters, as seen in the WiDr B5 clone.

3. Tumor size after a clinically relevant period

One can also define a different metric for treatment success, which considers the tumor size after a fixed treatment period (whether in cycles or in absolute time). To evaluate this quantity, we numerically integrated our dynamical model up to a fixed observation time, defined either in terms of absolute

time (we set it to 6 months after the start of treatment) or in terms of the total number of cycles (we set it to 30 cycles). We find that the choice of observation time whether in terms of absolute time or cycles does not alter our findings. This holds true both for success-failure diagrams (SM Figs. 15(a)–15(h) [17]) and protocol optimization (SM Figs. 15(i)–15(p) [17]). Moreover, we note that when the observation time is fixed in terms of the number of cycles, the final tumor size behaves similarly to long-term fitness expressed in units of 1/cycles (SM Figs. 15(a) and 15(b) [17], SM Figs. 15(e) and 15(f) [17], SM Figs. 15(i) and 15(j) [17], and SM Figs. 15(m) and 15(n) [17]), while when the observation time is fixed in absolute time, the final size corresponds to long-term fitness in units of 1/time (SM Figs. 15(c) and 15(d) [17], SM Figs. 15(g) and 15(h) [17], SM Figs. 15(k) and 15(l) [17], and SM Figs. 15(o) and 15(p) [17]). In other words, these analyses show that the average tumor size after months of treatment provides no additional information beyond the long-term fitness when evaluating the efficacy and optimality of a treatment-release protocol. This result arises because the observation time is far beyond the timescale of the transient dynamics, which generally last only a few cycles. As a result, the two variables become naturally related, with $N \sim \exp(\Lambda q)$. Therefore, the only difference between the two approaches is the effect of the transient phase. On the other hand, knowing the exact tumor reduction factor may be relevant for assessing the practical feasibility of a protocol within clinically relevant timescales. While tumor reduction may be negligible for protocols corresponding to very small values of long-term fitness, we note that this is precisely why our analysis of all possible protocols is valuable: It discriminates between truly effective protocols (in the $\Lambda < 0$ region) and those that are only theoretically effective. We also note that typical tumor reduction factor values in the $\Lambda < 0$ region of the protocol space vary in $[10^{-9}, 10^{-4}]$ (SM Fig. 15 [17]), which would ensure complete tumor extinction for the initial sizes considered in this study. For example, the optimal protocols minimizing the average tumor size after 30 cycles or 6 months from the start of treatment (red stars in SM Figs. 15(f) and 15(h) [17], respectively) induce tumor reduction factors of about 10^{-9} and 10^{-7} , respectively.

APPENDIX H: DERIVATION OF DABRAFENIB ABSORPTION AND ELIMINATION DYNAMICS OBSERVED IN PATIENTS

For the absorption and elimination dynamics of dabrafenib in patients we assume a double exponential trend of the

form

$$c(t) = \xi (e^{-k_e t} - e^{-k_a t}), \quad (\text{H1})$$

where k_a and k_e are the absorption and elimination constant rates, respectively, and ξ is a factor whose determination will be discussed below. Allowing for a lag time t_0 before actual absorption of dabrafenib, the previous equation reads

$$c(t) = \theta(t - t_0) \xi (e^{-k_e(t-t_0)} - e^{-k_a(t-t_0)}), \quad (\text{H2})$$

with $\theta(t)$ the standard Heaviside function. In the previous equation, ξ and k_e are unknowns that we now express in terms of measured experimental quantities ($c_{\max}$, the peak absorbed dose, $\text{AUC}(0, \tau)$, the area under the curve in the experimental time window and the measured absorbed rate k_a).

First, $c_{\max}$ can be determined as $c(t^*)$ where t^* is such that $c'(t^*) = 0$. After simple math, one finds $t^* = t_0 + \frac{\log k_a/k_e}{k_a - k_e}$ and $c_{\max} = \xi (\frac{k_a}{k_e} - 1) (\frac{k_a}{k_e})^{-\frac{k_a}{k_a - k_e}}$, so that

$$\xi = \frac{c_{\max}}{\left(1 - \frac{k_e}{k_a}\right) \left(\frac{k_e}{k_a}\right)^{\frac{-k_e/k_a}{k_e/k_a - 1}}} \quad (\text{H3})$$

Second, from the definition of $\text{AUC}(0, \tau) \equiv \int_0^\tau dt c(t)$ we find $\text{AUC}(0, \tau) = \xi \int_{t_0}^\tau (e^{-k_e(t-t_0)} - e^{-k_a(t-t_0)}) dt$ which we approximate to $\text{AUC}(0, \tau) \approx \xi (\frac{1}{k_e} - \frac{1}{k_a})$, so that

$$\xi = \frac{\text{AUC}}{\frac{1}{k_a} \left(\frac{k_a}{k_e} - 1\right)}. \quad (\text{H4})$$

Using Eqs. (H3) and (H4) we can therefore write

$$c(t) = \theta(t - t_0) \frac{\left(\frac{k_e}{k_a}\right) k_a \text{AUC}}{\left(\frac{k_e}{k_a} - 1\right)} (e^{-k_a(t-t_0)} - e^{-k_a(t-t_0) \frac{k_e}{k_a}}), \quad (\text{H5})$$

where k_e/k_a , solving the equation $\frac{k_e}{k_a}^{1/(\frac{k_e}{k_a} - 1)} = \frac{k_a \text{AUC}}{c_{\max}}$, can be written as

$$\frac{k_e}{k_a} = - \frac{W\left(-\frac{c_{\max}}{k_a \text{AUC}} \log \frac{k_a \text{AUC}}{c_{\max}}\right)}{\log \frac{k_a \text{AUC}}{c_{\max}}} \quad (\text{H6})$$

with W the Lambert (product logarithm) function.

Substituting the experimental values derived from the clinical trial [19], $c_{\max} = 1512 \text{ ng ml}^{-1}$, $\text{AUC} = 8820 \text{ ng ml}^{-1} \text{ h}$, $k_a = 1.88 \text{ h}^{-1}$, and $t_0 = 0.48 \text{ h}$, we can finally write

$$c(t) = \theta(t - 0.48) 2301 [e^{-0.23(t-0.48)} - e^{-1.88(t-0.48)}], \quad (\text{H7})$$

which is Eq. (18) in the main text and the function plotted in Fig. 5.

- [1] J. Bigger, Treatment of staphylococcal infections with penicillin by intermittent sterilisation, *The Lancet* **244**, 497 (1944).
- [2] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler, Bacterial persistence as a phenotypic switch, *Science* **305**, 1622 (2004).

- [3] P. Patra and S. Klumpp, Emergence of phenotype switching through continuous and discontinuous evolutionary transitions, *Phys. Biol.* **12**, 046004 (2015).
- [4] E. Kussell and S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments, *Science* **309**, 2075 (2005).

- [5] E. Kussell, R. Kishony, N. Q. Balaban, and S. Leibler, Bacterial persistence, *Genetics* **169**, 1807 (2005).
- [6] E. Levien, J. Min, J. Kondev, and A. Amir, Non-genetic variability in microbial populations: survival strategy or nuisance? *Rep. Prog. Phys.* **84**, 116601 (2021).
- [7] 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. Krishnamurthy 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 *et al.*, Tumor cells can follow distinct evolutionary paths to become resistant to epidermal growth factor receptor inhibition, *Nat. Med.* **22**, 262 (2016).
- [8] S. Shen, S. Vagner, and C. Robert, Persistent cancer cells: The deadly survivors, *Cell* **183**, 860 (2020).
- [9] K. Kochanowski, L. Morinishi, S. J. Altschuler, and L. F. Wu, Drug persistence—From antibiotics to cancer therapies, *Curr. Opin. Syst. Biol.* **10**, 1 (2018).
- [10] M. Decollogny and S. Rottenberg, Persisting cancer cells are different from bacterial persisters, *Trends Cancer* **10**, 393 (2024).
- [11] M. Russo, G. Crisafulli, A. Sogari, N. M. Reilly, S. Arena, S. Lamba, A. Bartolini, V. Amodio, A. Magrì, L. Novara, I. Sarotto, Z. D. Nagel, C. G. Piatt, A. Amatu, A. Sartore-Bianchi, S. Siena, A. Bertotti, L. Trusolino, M. Corigliano, M. Gherardi *et al.*, Adaptive mutability of colorectal cancers in response to targeted therapies, *Science* **366**, 1473 (2019).
- [12] 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, A modified fluctuation-test framework characterizes the population dynamics and mutation rate of colorectal cancer persister cells, *Nat. Genet.* **54**, 976 (2022).
- [13] N. Q. Balaban, S. Helaine, K. Lewis, M. Ackermann, B. Aldridge, D. I. Andersson, M. P. Brynildsen, D. Bumann, A. Camilli, J. J. Collins *et al.*, Definitions and guidelines for research on antibiotic persistence, *Nat. Rev. Microbiol.* **17**, 441 (2019).
- [14] M. Acar, J. T. Mettetal, and A. van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments, *Nat. Genet.* **40**, 471 (2008).
- [15] P. Patra and S. Klumpp, Population dynamics of bacterial persistence, *PLoS One* **8**, e62814 (2013).
- [16] Y. Pu, L. Li, H. Peng, L. Liu, D. Heymann, C. Robert, F. Vallette, and S. Shen, Drug-tolerant persister cells in cancer: The cutting edges and future directions, *Nat. Rev. Clin. Oncol.* **20**, 799 (2023).
- [17] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/mh5v-lcyw> for additional figures and analyses, including results on other cell lines and investigations of resistance and stochastic effects.
- [18] J. Tabernero, F. Ciardiello, F. Rivera, E. Rodriguez-Braun, F. Ramos, E. Martinelli, M. Vega-Villegas, S. Roselló, S. Liebscher, O. Kisker *et al.*, Cetuximab administered once every second week to patients with metastatic colorectal cancer: A two-part pharmacokinetic/pharmacodynamic phase I dose-escalation study, *Ann. Oncol.* **21**, 1537 (2010).
- [19] D. Ouellet, E. Gibiansky, C. Leonowens, A. O'Hagan, P. Haney, J. Switzky, and V. L. Goodman, Population pharmacokinetics of dabrafenib, a braf inhibitor: effect of dose, time, covariates, and relationship with its metabolites, *J. Clin. Pharmacol.* **54**, 696 (2014).
- [20] A. Dayimu, A. Gupta, R. N. Matin, J. Nobes, R. Board, M. Payne, A. Rao, A. Fusi, S. Danson, B. Eccles *et al.*, A randomised phase 2 study of intermittent versus continuous dosing of dabrafenib plus trametinib in patients with brafv600 mutant advanced melanoma (interim), *Eur. J. Cancer* **196**, 113455 (2024).
- [21] M. M. Fischer and N. Blüthgen, On minimising tumoural growth under treatment resistance, *J. Theor. Biol.* **579**, 111716 (2024).
- [22] E. B. Gunnarsson, B. V. Magnússon, and J. Foo, Optimal dosing of anti-cancer treatment under drug-induced plasticity, [arXiv:2412.16391](https://arxiv.org/abs/2412.16391) [q-bio.PE]
- [23] L. H. Biller and D. Schrag, Diagnosis and treatment of metastatic colorectal cancer: A review, *J. Am. Med. Assoc.* **325**, 669 (2021).
- [24] P. K. Dhanyamraju, T. D. Schell, S. Amin, and G. P. Robertson, Drug-tolerant persister cells in cancer therapy resistance, *Cancer Res.* **82**, 2503 (2022).
- [25] T. Kuosmanen, J. Cairns, R. Noble, N. Beerenwinkel, T. Mononen, and V. Mustonen, Drug-induced resistance evolution necessitates less aggressive treatment, *PLoS Comput. Biol.* **17**, e1009418 (2021).
- [26] J. Wen, J. P. P. M. Smelt, N. O. E. Vischer, A. L. de Vos, P. Setlow, and S. Brul, Heat activation and inactivation of bacterial spores: Is there an overlap? *Appl. Environ. Microbiol.* **88**, e02324-21 (2022).
- [27] N. G. Cogan, J. Brown, K. Darres, and K. Petty, Optimal control strategies for disinfection of bacterial populations with persister and susceptible dynamics, *Antimicrob. Agents Chemother.* **56**, 4816 (2012).
- [28] N. G. Cogan, H. Rath, N. Kommerein, S. N. Stumpp, and M. Stiesch, Theoretical and experimental evidence for eliminating persister bacteria by manipulating killing timing, *FEMS Microbiol. Lett.* **363**, fnw264 (2016).
- [29] K. J. Meyer, H. B. Taylor, J. Seidel, M. F. Gates, and K. Lewis, Pulse dosing of antibiotic enhances killing of a staphylococcus aureus biofilm, *Front. Microbiol.* **11**, 596227 (2020).
- [30] G. Singh, M. A. Orman, J. C. Conrad, and M. Nikolaou, Systematic design of pulse dosing to eradicate persister bacteria, *PLoS Comput. Biol.* **19**, e1010243 (2023).
- [31] B. Xue and S. Leibler, Evolutionary learning of adaptation to varying environments through a transgenerational feedback, *Proc. Natl. Acad. Sci. USA* **113**, 11266 (2016).
- [32] M. Raatz, S. Shah, G. Chitadze, M. Brüggemann, and A. Traulsen, The impact of phenotypic heterogeneity of tumour cells on treatment and relapse dynamics, *PLoS Comput. Biol.* **17**, e1008702 (2021).
- [33] Y. Viossat and R. Noble, A theoretical analysis of tumour containment, *Nat. Ecol. Evol.* **5**, 826 (2021).
- [34] M. Wang, J. G. Scott, and A. Vladimirovsky, Threshold-awareness in adaptive cancer therapy, *PLoS Comput. Biol.* **20**, e1012165 (2024).
- [35] J. T. Park and H. Levine, Mathematical characterization of drug-induced persistence in cancer, *bioRxiv* (2025).
- [36] V. Heinemann, L. F. von Weikersthal, T. Decker, A. Kiani, F. Kaiser, S.-E. Al-Batran, T. Heintges, C. Lerchenmüller, C. Kahl, G. Seipelt, F. Kullmann, M. Moehler, W. Scheithauer, S.

- Held, L. Miller-Phillips, D. P. Modest, A. Jung, T. Kirchner, and S. Stintzing, Folfiri plus cetuximab or bevacizumab for advanced colorectal cancer: Final survival and per-protocol analysis of fire-3, a randomised clinical trial, *Br. J. Cancer* **124**, 587 (2021).
- [37] D.-S. Wang, C. Ren, S.-S. Li, W. P. Fong, X.-J. Wu, J. Xiao, B.-K. Li, Y. Zheng, P.-R. Ding, G. Chen, M.-Z. Qiu, Z.-Q. Wang, F.-H. Wang, H.-Y. Luo, F. Wang, X.-Z. Wang, L.-Y. Wang, D.-J. Xie, T. Chen, L.-R. Li *et al.*, Cetuximab plus folfoxiri versus cetuximab plus folfox as conversion regimen in ras/braf wild-type patients with initially unresectable colorectal liver metastases (trice trial): A randomized controlled trial, *PLOS Med.* **21**, e1004389 (2024).
- [38] S. Qin, J. Li, L. Wang, J. Xu, Y. Cheng, Y. Bai, W. Li, N. Xu, L.-z. Lin, Q. Wu, Y. Li, J. Yang, H. Pan, X. Ouyang, W. Qiu, K. Wu, J. Xiong, G. Dai, H. Liang, C. Hu *et al.*, Efficacy and tolerability of first-line cetuximab plus leucovorin, fluorouracil, and oxaliplatin (folfox-4) versus folfox-4 in patients with ras wild-type metastatic colorectal cancer: The open-label, randomized, phase iii tailor trial, *J. Clin. Oncol.* **36**, 3031 (2018).
- [39] A. Avallone, F. Giuliani, A. De Stefano, G. Santabarbara, G. Nasti, V. Montesarchio, G. Rosati, A. Cassata, S. Leo, C. Romano, E. Tamburini, L. Silvestro, C. Lotesoriere, A. Nappi, D. Santini, A. Petrillo, A. Colombo, A. Febbraro, A. Leone, F. Mannavola *et al.*, Intermittent or continuous panitumumab plus fluorouracil, leucovorin, and irinotecan for first-line treatment of *RAS* and *BRAF* wild-type metastatic colorectal cancer: The IMPROVE trial, *JCO* **43**, 829 (2024).
- [40] V. Amodio, R. Yaeger, P. Arcella, C. Cancelliere, S. Lamba, A. Lorenzato, S. Arena, M. Montone, B. Mussolin, Y. Bian, A. Whaley, M. Pinnelli, Y. R. Murciano-Goroff, E. Vakiani, N. Valeri, W.-L. Liao, A. Bhalkikar, S. Thyparambil, H.-Y. Zhao, E. de Stanchina *et al.*, EGFR blockade reverts resistance to KRAS^{G12C} inhibition in colorectal cancer, *Cancer Discov.* **10**, 1129 (2020).
- [41] G. Rizzo, A. Bertotti, S. M. Leto, and S. Vetrano, Patient-derived tumor models: A more suitable tool for pre-clinical studies in colorectal cancer, *J. Exp. Clin. Cancer Res.* **40**, 178 (2021).
- [42] H. F. Cabanos and A. N. Hata, Emerging insights into targeted therapy-tolerant persister cells in cancer, *Cancers* **13**, 2666 (2021).
- [43] GitHub repository for M. Corigliano *et al.*, Optimal treatment for drug-induced cancer persists involves release periods and intermediate drug doses (2025), <https://doi.org/10.5281/zenodo.15647096>.
- [44] A. Dewanji, E. G. Luebeck, and S. H. Moolgavkar, A generalized Luria–Delbrück model, *Math. Biosci.* **197**, 140 (2005).
- [45] I. Bozic, J. G. Reiter, B. Allen, T. Antal, K. Chatterjee, P. Shah, Y. S. Moon, A. Yaqubie, N. Kelly, D. T. Le, E. J. Lipson, P. B. Chapman, J. Diaz, Luis A. B. Vogelstein, and M. A. Nowak, Evolutionary dynamics of cancer in response to targeted combination therapy, *eLife* **2**, e00747 (2013).
- [46] Y. Iwasa, M. A. Nowak, and F. Michor, Evolution of resistance during clonal expansion, *Genetics* **172**, 2557 (2006).
- [47] D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions, *J. Phys. Chem.* **81**, 2340 (1977).
- [48] E. B. Gunnarsson, S. De, K. Leder, and J. Foo, Understanding the role of phenotypic switching in cancer drug resistance, *J. Theor. Biol.* **490**, 110162 (2020).