Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Adaptive and sequential cancer therapies emerge from treatment schedule optimization

Charles D. Kocher
Memorial Sloan Kettering Cancer Center
Joseph O. Deasy
Memorial Sloan Kettering Cancer Center
Damon R. Reed
Memorial Sloan Kettering Cancer Center
Larry Norton
Memorial Sloan Kettering Cancer Center
Corey Weistuch
weistucc@mskcc.org

Memorial Sloan Kettering Cancer Center

Article

Keywords:

Posted Date: May 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9695978/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. DRR discloses the following relationships and financial interests. Ad board: Recordati; Data Safety Monitoring Committees: Springworks, Eisai.

Adaptive and sequential cancer therapies emerge from treatment schedule optimization

Charles D. Kocher, Joseph O. Deasy, Damon R. Reed, Larry Norton, and Corey Weistuch*

Memorial Sloan Kettering Cancer Center, New York, NY, USA

(Dated: May 14, 2026)

Abstract

Most cancer chemotherapy is delivered at the maximum tolerated dose (MTD) as often as possible, switching to a new combination of drugs upon progression. Clinical trial results have demonstrated that this sequential MTD schedule is not optimal for every patient in every situation; however, there is no systematic program for discovering what the optimal schedule actually is. We developed a simulated-annealing-based algorithm for computing personalized dosing schedules that best achieve treatment objectives. Given a mathematical model of a patient's drug response, the algorithm returns the optimal schedule for either minimizing tumor burden or maximizing progression-free survival. We tested this algorithm on a simple mathematical model of patient treatment response; we found that there are only three optimal one-drug schedules: continuous therapy and two distinct adaptive therapies. For two drugs, each schedule used in the clinic was found to be optimal for some patient parameters, and the algorithm was able to quantitatively map out the boundaries between these schedules. The optimization algorithm developed here can play a key role in upcoming personalized cancer treatment pipelines.

I. INTRODUCTION: A BRIEF SURVEY OF CANCER DRUG SCHEDULING

The prevailing paradigm in cancer chemotherapy treatment is that of dose-dense treatment at the maximum tolerated dose (MTD) [1–7]. In the adjuvant setting, the patient is given a fixed number of treatment cycles, generally followed by observation to detect relapse. In the non-resectable and relapse settings, patients are given sequential cycles of (combinations of) drugs until either unacceptable toxicity or disease progression occurs, at which time the next drug combination in the sequence is used (sequential MTD).

Increasingly, there have been calls to re-examine the MTD paradigm as the default in every case of cancer treatment, both in terms of the dose used per cycle and the scheduling of those cycles. For example, MTD dosing as a guiding principle was carried over by default into the realm of targeted cancer therapies, but there is limited evidence that the maximum dose of a targeted agent is the one that best balances treatment efficacy and toxicity [2, 3, 8, 9]. On the scheduling side, MTD has also been challenged by proponents of evolution-informed

* To whom correspondence should be addressed: weistucc@mskcc.org

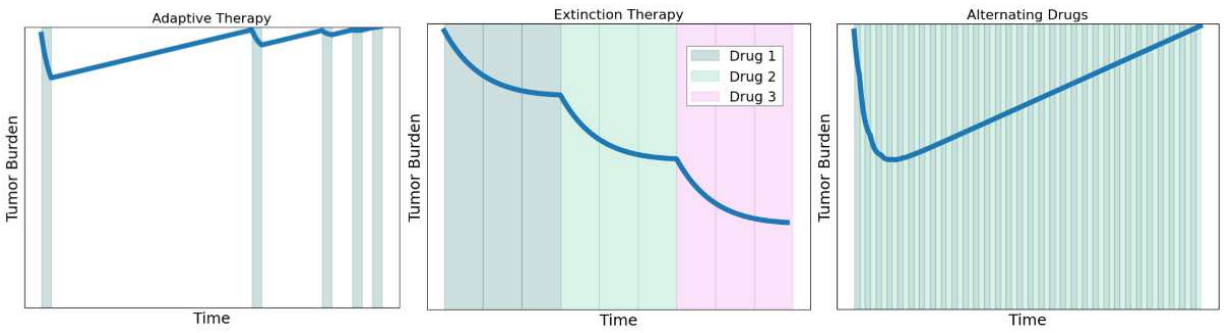


FIG. 1. **Some proposed alternatives to sequential MTD scheduling.** As discussed in the text, these schedules are motivated by evolutionary theory [14–16], and they have differing levels of clinical testing and implementation.

treatments [1, 10–12], particularly for relapsed or metastatic patients. The evolutionary argument against MTD is that maximum drug dosing applies the maximum selective pressure for resistance, so that MTD scheduling hastens the development of resistance and ultimately progressive disease. This perspective is bolstered by recent positive clinical experience with adaptive therapy [13], which showed that introducing response-informed treatment holidays increased progression-free survival (PFS) of metastatic castrate resistant prostate cancer patients beyond the PFS achieved by using MTD (note, however, that the drug used in this study was abiraterone—a hormone therapy—not a cytotoxic chemotherapy).

Using mathematical models of treatment response, many alternative treatment schedules have been developed that seek to replace MTD in different scenarios (see Fig 1); these schedules are being tested in the clinic. Above, we mentioned adaptive therapy [13, 14, 17], which aims to prolong PFS. There are currently clinical trials in basal cell carcinoma (NCT05651828), melanoma (NCT03543969), metastatic castrate resistant (NCT05393791, NCT06216249) [18] and sensitive (NCT03511196, NCT06734130) prostate cancer, and ovarian cancer (NCT05080556). Extinction therapy [1, 15, 19] aims to cure patients using a well-timed sequence of therapy changes (successive "strikes"). It is currently being evaluated in rhabdomyosarcoma (NCT04388839) [20], Ewing sarcoma (NCT07194044) [21, 22], metastatic castration sensitive prostate cancer (NCT05189457), and metastatic breast cancer (NCT06409390). In addition, the standard of care treatment in acute lymphoblastic leukemia is an extinction-like therapy, although the exact sequencing of drugs was discovered empirically, not through modeling [10, 23]; similarly for the chemotherapy regimen

AC-T (sequential administration of a taxane like paclitaxel and combined doxorubicin and cyclophosphamide) used in the adjuvant setting to treat localized breast cancer, where the treatment is switched without any sign of progression [24–26].

These alternative schedules hope to improve patient outcomes, decrease toxicities, and lower treatment costs. Caution is warranted, however. It has been demonstrated that intermittent treatment does not always mean higher quality of life, even though toxicity is lower [27]. Modeling incorporating evolutionary principles has tried to predict better treatment schedules before with limited success. For example, in [16], it was proposed that alternating cycles of non-cross resistant drugs would give the best outcome[1]. When this alternating strategy was tested in the clinic, a sequential schedule proved significantly better than the alternating schedule [28]; sequential MTD prevailed. In spite of this notable failure, alternating schedules are still the standard of care for Ewing sarcoma (VDC/IE) [29, 30], osteosarcoma (MAP) [31, 32], and Burkitt lymphoma (CODOX-M/IVAC, hyperCVAD + R) [33–35], as well as finding use in germ cell tumors (TI-CE) [36], pancreatic cancer (nab-P/Gem-mFOLFOX) [37], and Wilms tumor (e.g. DD4A and Regimen M, among others) [38–44].

In summary, improving on the sequential MTD paradigm is an achievable goal, but the difficulty of the problem, as evidenced by the paucity of clinical trial results that go against sequential MTD, should be appreciated. Breakthroughs have been made, but current results are limited. Moreover, the alternative schedules (adaptive therapy, extinction therapy, alternating therapies) that have been proposed and are currently being tested in or are already in use in the clinic are based on heuristics, leaving room for further optimization and personalization [17, 45–51]. In order to maximize the chances of success, it would be ideal to know that a proposed schedule is in some way optimal before undertaking a clinical trial.

To address this need, we have developed an algorithm based on simulated annealing for finding the optimal drug schedule—either for minimizing tumor burden or maximizing PFS—given a reasonable model of a patient's treatment response. We used this algorithm to probe the optimal schedules of a simple mathematical model of patient treatment response which was proposed previously in [52], creating "phase diagrams" of what schedules would be optimal for which types of patients. We found that, when only one drug was available,

[1] Interestingly, in the same paper, they model a version of extinction therapy, foreshadowing its development many years later.

there were only three classes of optimal schedules, including the adaptive therapy schedule we discussed above. For two drugs, the sequential MTD, alternating, and second-strike extinction therapy schedules were all found as optimal schedules in different parts of the patient parameter space. The algorithm we developed can play a key role in personalizing patient treatment decisions once the rest of the personalization pipeline (better modeling and prediction of patient treatment response) has been developed.

II. METHODS

A. Using simulated annealing as the basis for the schedule optimization algorithm

First, we formulate our schedule optimization problem. To make the problem tractable, we cut time up into segments of one week. During each week, patients can receive one of N treatments[2], or can receive no drug (denoted by 0). Therefore, a treatment schedule C is a sequence of numbers drawn from the set $\{0,1,2,\ldots,N\}$.

We assume that a mathematical model of the patient's response to all treatments is known. In our applications below, this assumption means that we know all of the parameters that go into an ODE model of tumor burden for a patient before beginning treatment. Estimation of these parameters a priori is beyond the scope of this manuscript and is the subject of ongoing research.

Given a model of treatment response and some initial conditions of that model, the drug schedule C can be used to simulate the patient's tumor burden trajectory, $T(t;C)$. The simulation should continue until either cure is reached (during the drug dosing period, by dropping the tumor burden below a set cure threshold T_{cure}), or progression is reached (by the tumor burden crossing a set progression threshold T_{prog}). If cure is not reached, then the treatment schedule can be padded with 0's until progression is reached. Therefore, each simulated trajectory ends at either T_{cure} or T_{prog}.

Next, we quantitatively defined the objectives we wanted to optimize. The first treatment objective is curing the cancer. The optimal curative schedule was assumed to be the one that reached the lowest tumor burden level at any point of its trajectory (i.e. the one that

[2] For simplicity, treatments are assumed to be dosed at the MTD, so that combinations are not allowed. However, treatments themselves can denote combinations of drugs. I.e., treatment 1 can be a three drug cocktail, where we assume those drugs are always used together.

was closest to T_{cure}). We introduced an energy landscape,

$$\mathcal{E}_{\text{cure}}(C) = \log\left[1 + \min_t T(t;C)\right] \ , \tag{1}$$

which gives the optimal curative schedule when minimized over all schedules C. The second treatment objective we considered was maximizing the PFS. We similarly introduced an energy landscape[3],

$$\mathcal{E}_{\text{PFS}}(C) = \frac{100}{\log\left[1 + \min_t\{t : T(t) \geq T_{\text{prog}}\}\right]} \ , \tag{2}$$

which gives the optimal PFS schedule when minimized over all schedules C.

Simulated annealing can be used to find the global minima of complicated energy landscapes [53, 54], like the ones in Eqs (1) and (2). We now sketch how we applied it to the schedule optimization problem. First, a schedule C was initialized and the energy $\mathcal{E}$ was calculated. Then, a change was made to the schedule (either switching one of the numbers in the sequence C to one of the N other numbers, or moving a non-zero number to the left or right if the target space was a zero) to produce a test schedule C' with corresponding energy $\mathcal{E}'$. If $\mathcal{E}' < \mathcal{E}$, the change is accepted. If $\mathcal{E}' > \mathcal{E}$, then the change is probabilistically accepted with probability

$$P(\text{accept}, n) = \exp\left[-\beta(n)(\mathcal{E}' - \mathcal{E})\right] \ . \tag{3}$$

where n is the number of iterations performed so far and $\beta(n)$ is the "cooling schedule" (an increasing function of n, since it is the inverse temperature) that ensures the algorithm converges to the minimum of $\mathcal{E}$. We used

$$\beta(n) = \begin{cases} 0.01 & n < n_{\max}/4 \\ 0.01\exp\left[10(n/n_{\max} - 1/4)\right] & n \geq n_{\max}/4 \end{cases} \ , \tag{4}$$

with $n_{\max} = 10,000$ being the number of iterations of the simulated annealing algorithm, throughout, for both minimizing the tumor burden and maximizing PFS.

For our specific application, $\mathcal{E}' = \mathcal{E}$ is a special case: Not all schedules that have the same energy are actually equal. For example, if cure is achieved in both schedules (say

[3] Note that the factor of 100 in the numerator of Eq (2) was used to roughly equalize the scales of $\mathcal{E}_{\text{cure}}$ and $\mathcal{E}_{\text{PFS}}$ in the simulations that we performed, so that the same cooling schedule could be used for both.

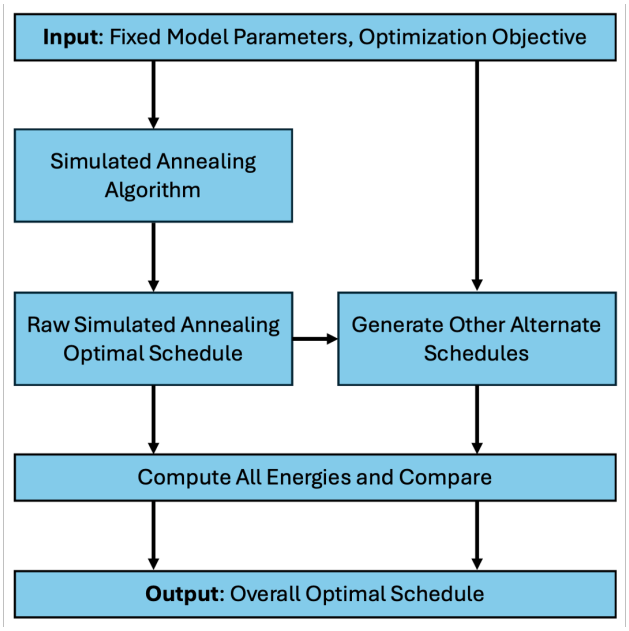


FIG. 2. **Schematic of the schedule optimization algorithm.** Given a mathematical model of drug response, a patient's parameters in that model, and the treatment objective (minimize tumor burden or maximize PFS), the algorithm outputs the schedule that optimizes that treatment objective for that patient in that model.

at t_{cure} and t'_{cure}, the schedule which cures first should be preferred. In this case, the acceptance probability was set to $P(\text{accept}) = \min\{t_{\text{cure}}/t'_{\text{cure}}, 1\}$ (if only one schedule cured, that schedule was definitely accepted). Similarly, if two schedules had the same $\mathcal{E}$, but used different amounts of drug, the schedule that minimizes the amount of drug used was preferred. In this case,

$$P(\text{accept}) = \min\left\{\frac{1+|\{c \in C : c > 0\}|}{1+|\{c \in C' : c > 0\}|}, 1\right\} , \tag{5}$$

where $|\cdot|$ denotes the cardinality of a set.

As discussed in SI Appendix B, we tested the simulated annealing algorithm by comparing its output to a few cases where the global optimal schedule could be found by brute-force enumerating all possible schedules. We observed some noise and variability around the optimal schedule when the total required treatment time was long. While this variability had little impact on the performance of the algorithm (i.e. the PFS times or minimum tumor burdens achieved by the raw simulated annealing output schedule were always very close to optimal), the theoretical interpretation of the model's output, e.g. in our phase

diagrams below, could be distorted by it. To reduce this noise, we added a comparison of the simulated annealing algorithm's raw output to a few other well-motivated schedules, then picked the best of these as the optimum. For example, for one drug, we additionally tested a continuous dosing schedule, an adaptive therapy schedule, and a "standardized" version of the simulated annealing output. More details can be found in SI Appendix B. A full schematic of our optimization algorithm can be seen in Fig 2.

B. The GDRS model is a simple but realistic description of patient treatment response

Our optimization algorithm requires a model that predicts patient treatment response and an individual patient's parameter values in that model. To test the algorithm, we used the GDRS model of [52]:

$$\frac{dT}{dt} = \left(\gamma - \sum_i \delta_i D_i(t) E_i(t)\right) T(t) \ , \tag{6a}$$

$$\frac{dE_i}{dt} = \left[s_i(1 - E_i(t))(1 - D_i(t)) - r_i D_i(t)\right] E_i(t) \ , \tag{6b}$$

where $T(t)$ is the tumor volume (or proxy for it) at time t and $D_i(t) \in [0, 1]$ is the dose of drug i (scaled by the MTD). The GDRS model is a minimal model describing tumor Growth (rate γ), Death due to drug (rate δ_i), Resistance to drug (rate r_i), and re-Sensitization to drug (rate s_i). It is a phenomenological model, via the use of the "drug efficacy" parameter $E_i(t) \in [0, 1]$ to quantify the degree of resistance in the tumor population.

A similar model was derived from first principles in [55] that connects the terms in GDRS-like models to the mechanisms from which they arise and clarifies the GDRS model's implicit assumptions. GDRS assumes that tumors grow exponentially, with no difference in growth rate between the sensitive and resistant cells when there is no drug. Drug action/resistance mechanisms are assumed to be independent, so that the population fractions of sensitive and resistant cells act as independent probabilities. For example, if p_k is the fraction of cells sensitive to drug k, then $p_{1\&2} = p_1 p_2$. This type of independence has been used before to accurately model treatment response [56].

For our simulations, we used the GDRS model with the initial condition $T(0) = 1 \times 10^{12}$ cells (this starting point was also used in [56]) and $E_i(0) = 1$ for all drugs i. The progression

threshold was set to $T_{\text{prog}} = 2T(0) = 2\times10^{12}$ cells so that an untreated tumor would progress in one doubling time, and the cure threshold was set to $T_{\text{cure}} = 1$ cell.

To demonstrate the utility of this model beyond the examples shown in [52], we have given an example in SI Appendix A of how it fits clinical data from MSK-CHORD [57]. Given the limited resolution of cancer patient data, we conclude that, while simple, the GDRS model should give a reasonable representation of the space of patient drug responses, so that the optimal schedules we compute using it will give a reasonable representation of the space of optimal strategies.

III. RESULTS AND INTERPRETATION

A. There are only three types of optimal schedule for one drug.

After formulating the algorithm, we explored the GDRS parameter space for one drug. Through this non-systematic searching, we noticed that the three types of schedules[4] in Fig 3(A-C) were the optimal ones. One was continuous treatment, where drug was used every time possible. The other two were types of adaptive therapy, which we have called T-type and E-type after the parameters of the GDRS model that mainly inform the strategy.

In E-type adaptive therapy, one window of treatment is followed by a treatment break that lasts until the drug efficacy E reaches some acceptable level (via re-sensitization), at which point drug is used again. These schedules are largely characterized by cycles in which E returns to its original value, but T ends up lower. This mechanism and the trajectories produced by it are reminiscent of radiotherapy, in which treatment breaks via fractionation allow for cell cycle redistribution and other mechanisms of re-sensitization [58].

In T-type adaptive therapy, a window of treatment is followed by a treatment break that lasts until the tumor burden T approaches the progression threshold, at which point drug is used again to prolong PFS. The idea of this schedule is to use drug only when necessary so that sensitivity can build up in between doses. This type of adaptive therapy is close to, but not exactly the same as, the adaptive therapy introduced in [13, 14], because here we do not set a response threshold for stopping treatment. We only use the progression threshold

[4] In these results and the ones that follow, we enforced a rest week after each treatment, so that drug must be followed a 0 in the sequence. This rest week can be turned on or off. See SI Appendix D for a discussion on the guiding principles we used when choosing which parts of parameter space to explore.

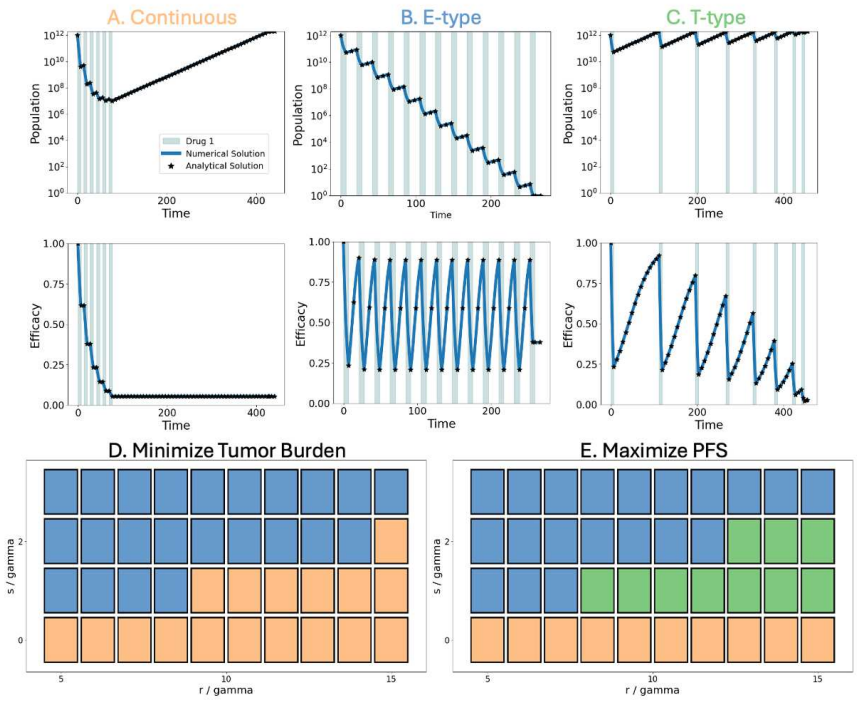


FIG. 3. **Phase diagrams for optimal schedules with one drug.** There are three types of optimal schedule for one drug, shown here with their effects on cancer population (first row) and drug efficacy (second row): (A) Continuous therapy, where drug is used every possible window. (B) E-type adaptive therapy, where drug is used when an efficacy threshold is crossed. (C) T-type adaptive therapy, where drug is used just before progression. (D and E): Results of sweeping through r-s phase space at constant γ and δ and finding the optimal schedule. Color on the phase diagram corresponds to the color of the titles for (A-C).

for turning the drug on.

Our heuristic understanding from anecdotal examples was that these three schedules were the only ones that could be optimal for one drug. We fixed the GDRS parameters γ (growth) and δ (drug kill) and swept over resistance-sensitivity r-s parameter space, finding the optimal schedule at each point. We were then able to sort the obtained optimal schedules into the three types we identified above using the schedules we tested for the optimum (see SI Appendix B): The continuous schedule was identified with the continuous strategy, the adaptive therapy schedule with the T-type adaptive therapy strategy, and both the raw and standardized simulated annealing output were mapped to the E-type strategy.

The resulting phase diagrams for both objectives, minimizing tumor burden and maxi-

mizing PFS, are shown in Fig 3(D-E). Except for three points on the phase boundary, the PFS phase diagram is the same as the tumor burden phase diagram but with the $s > 0$ continuous strategies turned into T-type adaptive therapy. The three exception points turned E-type adaptive therapy into T-type adaptive therapy. In Appendix C, we investigate these points, showing that the anomalous behavior happens due to the discretization of time into cycles. Accordingly, the phase boundary here acts like a second-order phase transition, where properties of both the T-type and E-type schedules can emerge near the boundary, for both optimization objectives.

B. For two drugs, the clinically used schedules are optimal for some parameters

We next tested various parameters for two drugs. As in the single drug case, we first explored parameter space by choosing different parameter values in an ad hoc manner. We were able to confirm that clinically used schedules (sequential MTD, alternating, and second-strike) were the optimal schedule for some parameter values. Some of the optimal schedules we obtained are shown in Fig 4(A-E).

The full phase space for two drugs is vast; instead of doing a systematic analysis as for one drug, we chose a few different ranges of parameter values for the first drug and created phase diagrams; see Fig 4(F-G). We were able to capture phase boundaries between pairs of the alternating, second-strike, sequential MTD, and adaptive therapy strategies. The optimal schedules that we identified as two drug adaptive therapies could again be sorted into two different buckets, curative and non-curative, in direct analogy to the E- and T-type adaptive therapies of Fig 3. Further discussion of this comparison is given in SI Appendix C, as the "second-order-like" phase transition boundary becomes more pronounced in the two drug case.

IV. DISCUSSION

In this paper, we have reported two main results. First, we have created a simulated-annealing-based algorithm for finding a patient's personalized optimal treatment schedule given that a way to model the patient's response to each possible treatment is known. This algorithm was validated in simple cases, and it can be easily adapted to use any mathematical

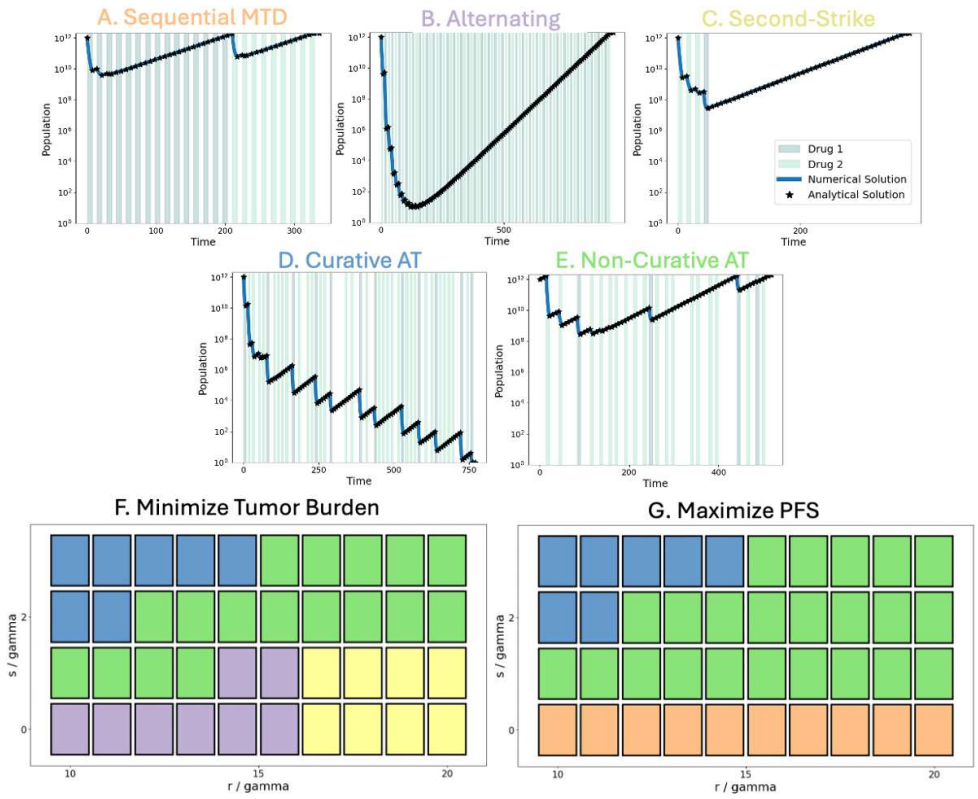


FIG. 4. **Example phase diagrams for the two drug case.** The simulated annealing algorithm returns the schedules of Fig 1 as optimal for different parts of parameter space: (A) sequential MTD, (B) alternating drugs, (C) second-strike. Additionally, (D) curative and (E) non-curative adaptive therapies analogous to the E- and T-type, respectively, adaptive therapies for one drug were found to be optimal. (F and G) Sweeping through resistance-sensitivity phase space of the first drug. Coloring corresponds to the titles for (A-E). For more details on the adaptive therapy schedules, see SI Appendix C.

model of patient treatment response. Second, we used the algorithm to explore the space of optimal schedules in the simple GDRS model of cancer patient treatment response. For one drug, we mapped out the phase diagram of optimal schedules and found only three different types of schedules could be optimal: continuous therapy, T-type adaptive therapy, and E-type adaptive therapy. We found features similar to those of a second-order phase transition at the boundary. For two drugs, we found that each of the clinically used schedules was optimal for some parameters and gave examples of the phase diagrams, illustrating the boundaries between these schedules.

These latter results all depend on the GDRS model. In principle, if we were to use a different treatment response model in our optimization algorithm, then they could change. However, as we argued in Section II B and SI Appendix A, the GDRS model is a reasonable description of actual patient treatment responses. Therefore, the results that we report here should have a similar degree of generality: optimal schedules using other models with independently-acting drugs and U-shaped treatment response trajectories (ones that reproduce patient dynamics such as in Fig S1) should have similar features.

The optimization algorithm was able to resolve the precise phase boundaries between these strategies, which had not previously been achieved. Our phase diagram approach leads to an important observation: Different patients in the same clinical context could fall on different sides of these phase boundaries (because of the intrinsic characteristics of their disease, which manifest as different GDRS parameters). For example, in a hypothetical clinical trial where this situation arises, the best results could be achieved only by personalizing the patients' strategies. A "one-schedule-fits-all" approach would always be sub-optimal. Predicting where patients fall on the phase diagram ahead of time using baseline demographic, clinical, genomic, or other data (e.g. generalizing successes in predicting radiotherapy response [59–61]), or from early treatment response [62, 63], is an open problem; the combination of machine learning and large-scale clinical data aggregation [57] is a promising potential route toward a solution.

We also found that sequential MTD was optimal for maximizing PFS when the resensitization parameter of each drug was low. This result should serve as another reminder that improving upon the sequential MTD paradigm could be an incredibly difficult problem, as we highlighted in Section I. It very well could be the case that many of our patients and treatments are in the region where sequential MTD is actually optimal, so that improvement is not possible (a similar argument could explain why alternating therapy failed in a breast cancer trial [7, 28]). On the other hand, for situations where sequential MTD is not the clinically used schedule, it is possible that these patients, tumors, and drugs are in a different region of parameter space. However, we must also consider explanations drawing on effects beyond those of the simple GDRS model, such as collateral sensitivities [64–66] or pharmacodynamics.

V. CONCLUSION

So far in clinical practice, testing alternatives to the sequential MTD scheduling paradigm has been a patchwork that has produced notable successes and failures. We have introduced an algorithm for finding the optimal cancer treatment schedule given a mathematical model of a patient's response to each possible treatment and demonstrated the types of optimal schedules that it can find using the simple GDRS model. Once reliable methods for predicting (and updating in real-time) patient treatment response are implemented, such an algorithm can provide a framework for systematically optimizing the use of existing drugs through personalized scheduling.

CODE AVAILABILITY

Code used to create all the plots in this paper, which includes the simulated annealing optimization algorithm, can be found at `https://github.com/cdkocher/optimal-cancer-treatment-scheduling`.

AUTHOR CONTRIBUTIONS

CDK conceptualized the project, wrote the code, performed the analysis, interpreted results, wrote the manuscript, and revised the manuscript. JOD conceptualized the project, interpreted results, and revised the manuscript. DRR conceptualized the project, interpreted results, and revised the manuscript. LN conceptualized the project, interpreted results, and revised the manuscript. CW conceptualized the project, interpreted results, revised the manuscript, and supervised research.

COMPETING INTERESTS

DRR discloses the following relationships and financial interests. Ad board: Recordati; Data Safety Monitoring Committees: Springworks, Eisai.

ACKNOWLEDGMENTS

This research was funded in part by the Simons Foundation, the Breast Cancer Research Foundation (grant MATH-23-001 and BCRF-17-193), and an NIH/NCI Cancer Center Support Grant (P30CA008748). CW is supported by an MIB Agents Outsmarting Osteosarcoma Hero Award. DRR support includes from Brandon's Defense and the Tsigutkin Family Fund.

This manuscript is the result of funding in whole or in part by the National Institutes of Health (NIH). It is subject to the NIH Public Access Policy. Through acceptance of this federal funding, NIH has been given a right to make this manuscript publicly available in PubMed Central upon the Official Date of Publication, as defined by NIH.

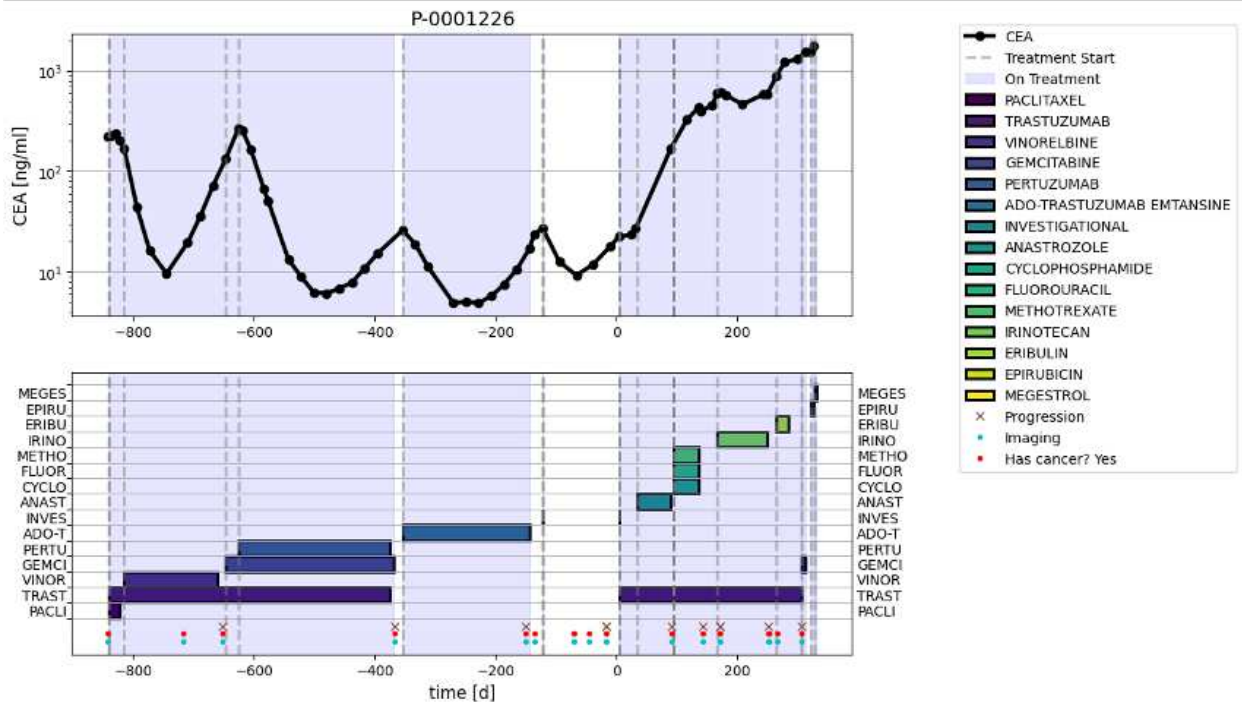


FIG. S1. **An example patient trajectory from MSK-CHORD.** This patient had metastatic breast cancer.

Appendix A: GDRS model fits data from MSK-CHORD

MSK-CHORD [57] is a dataset that has longitudinal blood biomarker data for a few cancer types, as well as treatment data and genomics. This data set can be viewed on cBioPortal. Here, we just use the clinical information, not the genomics. As an example, look at Fig S1. Shown is patient P-0001226, diagnosed at time $t = -3665$ days (t is the difference from when the first genomic sequencing was finished) with AJCC stage I (HR+, HER2+) breast invasive ductal carcinoma which ultimately metastasized, leading to the patient's death at $t = 334$ days. MSK-CHORD reports the carcinoembryonic antigen (CEA, a proxy for tumor burden in some advanced breast cancer cases) level over time for this patient, as well as all treatments, imaging results, and progression, as extracted from clinical and radiological notes.

Fig S1 shows a clear U-shaped response of the patient's CEA level to the different treatments, the characteristic sign of evolutionary rescue by drug resistance [1, 10, 67, 68]. That is, the population of cancer cells that is treatment sensitive—the bulk of the initial population—declines in number while the drug is applied, but a small resistant population—which is

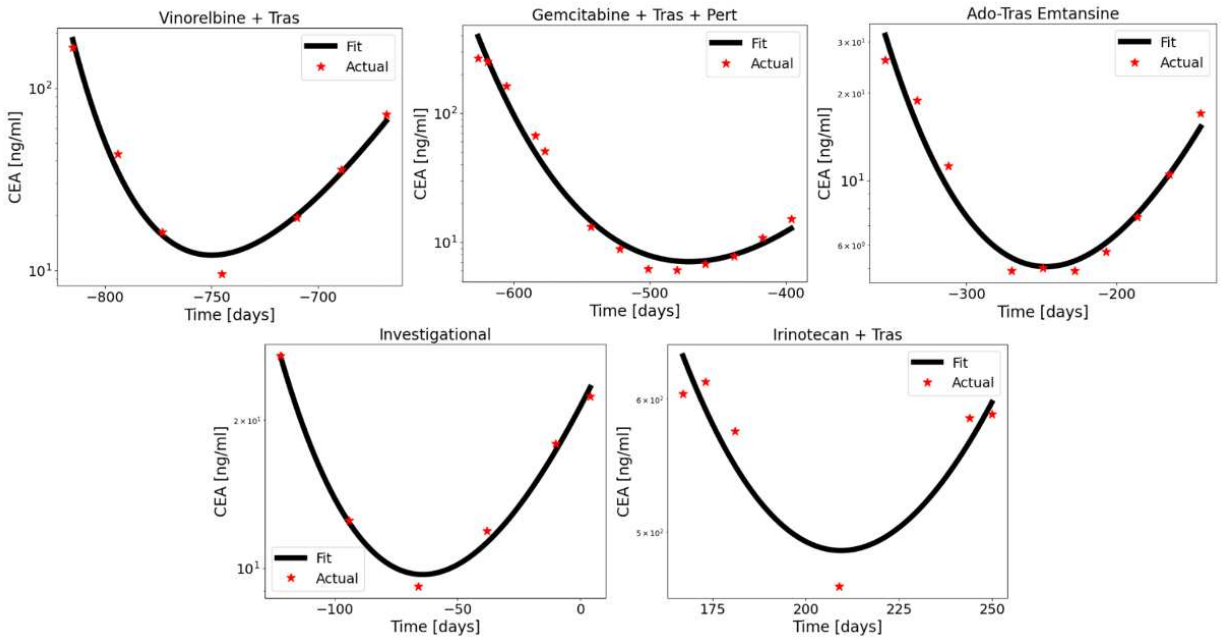


FIG. S2. **Fitting the GDRS model to MSK-CHORD patient data.** The GDRS model, Eqs (6), was fit to the five U-shaped regions of the patient CEA trajectory shown in Fig S1. Agreement between the fit and data are reasonable: The fit is able to roughly capture the timing and depth of the CEA minimum, suggesting the GDR model is capturing the essence of the CEA dynamics.

either pre-existing or discovered *de novo* through genetic or epigenetic changes—is able to continue growing. Eventually, the total population of sensitive plus resistant cells reaches a minimum, after which the population is mostly resistant cells, and treatment begins to fail. Ultimately, imaging-confirmed progression is reached (brown x on the bottom plot), prompting a switch in treatment.

We can fit the U-shaped curves from this patient using the GDR part of the GDRS model (that is, no re-sensitization because the drug is never turned back on after it is turned off). Results are shown in Fig S2. To do so, we manually picked out the treatment windows that this patient had U-shaped responses. We then estimated this patient's CEA growth rate as $\gamma = 0.042$ per day by taking the maximum growth rate implied by the difference between consecutive CEA measurements. We then used this value of γ to find δ_i and r_i for each drug by using SciPy's curve_fit function on the log of the CEA data with the analytical solution to the GDR model. The full method can be found in the code provided on GitHub (see code

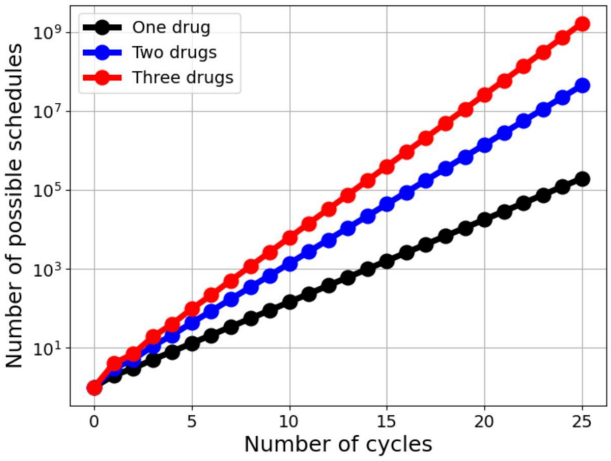


FIG. S3. **Number of schedules as a function of the number of cycles (weeks), for one to three drugs.** The scaling is exponential in the number of cycles, but searching approximately 10^4 schedules was found to run reasonably quickly on the author's laptop.

availability statement in the main text).

Appendix B: Validating the simulated annealing algorithm using brute-force enumeration and fine-tuning by testing additional schedules

We wanted to test the simulated annealing algorithm in cases where the optimal schedule was known exactly. To do so, we wrote a separate optimization algorithm that enumerated all schedules of a specified length, created their tumor population trajectories, then picked out the optimal one according to the rules used for the simulated annealing algorithm. This brute force method is tractable if the number of weeks is not too large and a rest week is enforced: see Fig S3, which shows the total number of schedules as a function of the total number of weeks for one to three drugs.

To test minimizing the tumor burden, we used $\gamma = \log(2)/20$, $\delta = 50\gamma$, $r = 6\gamma$, and $s = 45\gamma/4$. To test maximizing the PFS, we used $\gamma = \log(2)/20$, $\delta = 25\gamma$, $r = 15\gamma$, and $s = \gamma/2$. We enumerated all schedules over a 16 week period, with mandatory rest weeks after using drug. The results of the enumeration are shown as the top row in Fig S4; the simulated annealing output is shown as the bottom row.

The only difference between the figures is that we have shown the full 16 weeks in the brute force case to illustrate our next point. We now prove that the brute force enumeration

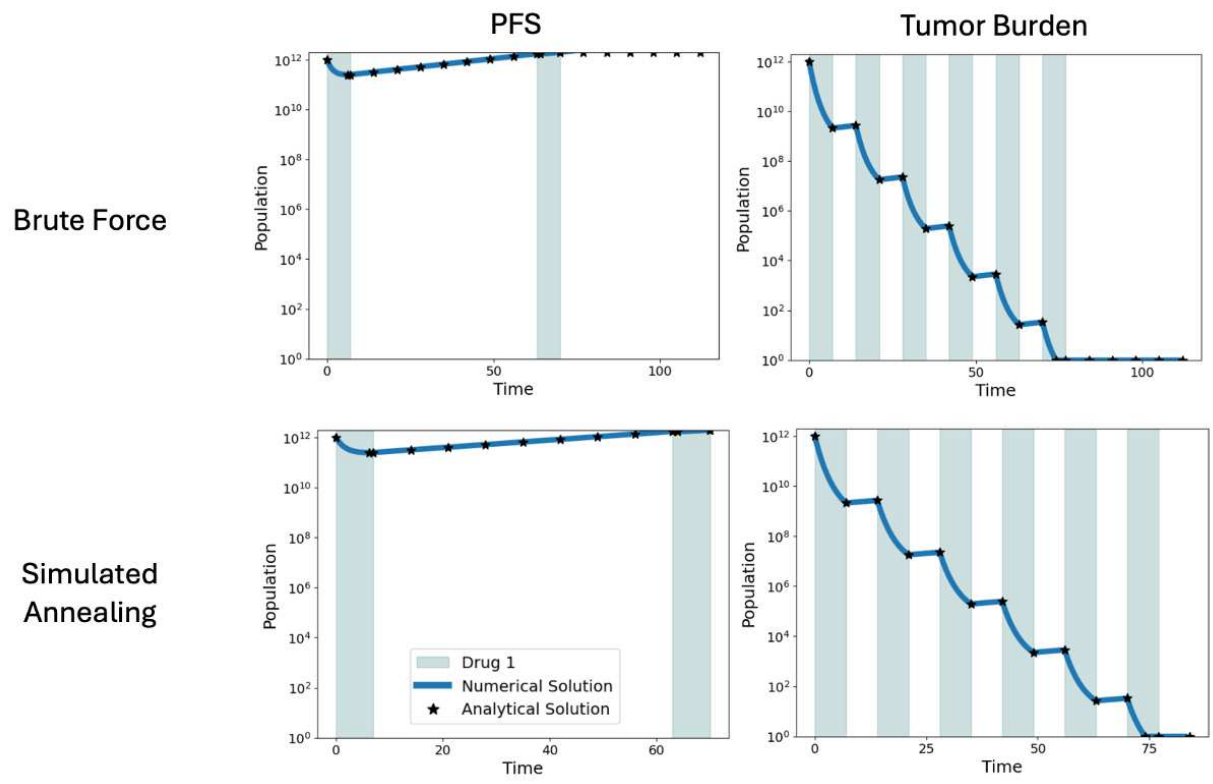


FIG. S4. **Comparing the simulated annealing algorithm to a brute force approach.** Simulated annealing gives the same optimal schedule. These schedules are their respective global optima because adding more cycles cannot impact the optimum PFS or cure time.

is finding the global optimum in each case. First, in the PFS case, all enumerated schedules progressed by the 16th cycle (the optimal schedule shown being the one that took the longest to do so). Therefore, adding cycles cannot change the optimal schedule: any treatments after the last one shown cannot impact the trajectory because the patient has already progressed. We have enumerated all potentially optimal schedules, meaning we have found the global optimum. Similarly, in the cure case, adding cycles cannot impact the optimum. The optimal schedule is the one that cures first, so only changes before the shown cure time would impact the outcome. Since all such schedules were already in the enumeration, we have found the global optimum.

However, the simulated annealing algorithm is stochastic. Accordingly, it does not always find the same optimal schedule. In the simple cases of Fig S4, there is no issue: we ran the simulated annealing algorithm 100 times on each and observed that it found the globally

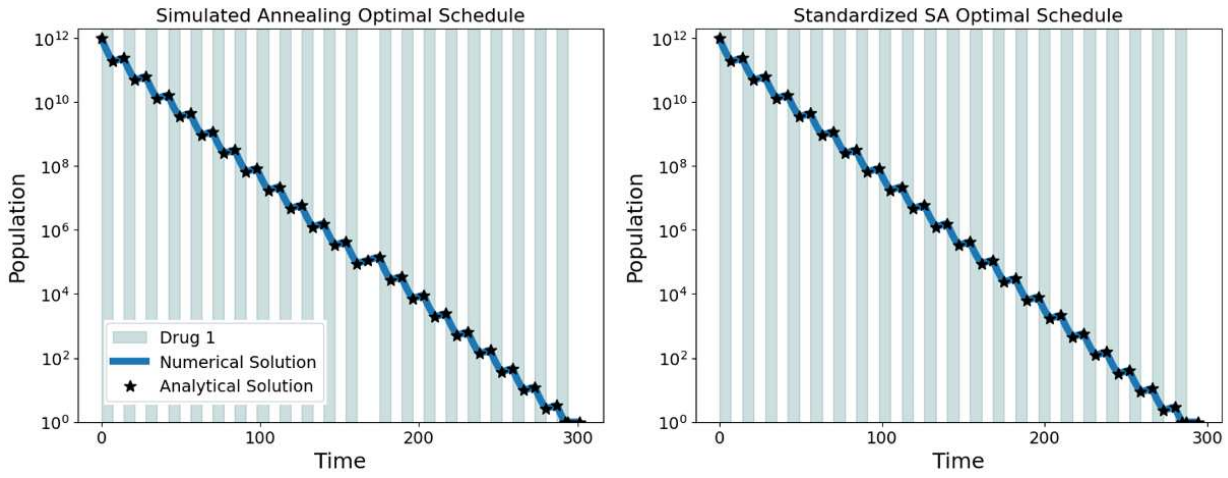


FIG. S5. **Testing a standardized version of the simulated annealing output increases the accuracy.** Simulated annealing struggles to find the global optimum (right) when the search space gets large, outputting incorrect answers like the one shown here (left) 69/100 times for these parameters. Testing a schedule that standardizes the time between drug doses increases the accuracy to 99/100 in this case.

optimal schedule all 100 times (see the code).

When the GDRS parameters are such that the optimal schedule is much longer (i.e. the space of potentially optimal schedules is much larger than the 10,000 iterations we are doing for the simulated annealing), then the stochasticity can cause variability in the output. For example, consider Fig S5. The optimal strategy for these parameters is to use drug every time available, leading to cure in 41 cycles. However, there are over 430 million schedules of this length, much more than the 10,000 iterations of the simulated annealing algorithm. Accordingly, the algorithm does not always sample the exact global optimum: for this case, it only found the correct schedule 31/100 times (see the code). As shown in the figure, the difference in cure time between the global optimum and the returned schedule was often very small; however, the interpretation of the optimal schedule is different for the left and right side of Fig S5. On the right, the true optimum is a continuous dosing schedule. On the left, the nearly-optimal schedule is not a continuous dosing schedule. Such differences can cause problems when trying to draw theoretical conclusions about optimal scheduling.

To address this issue and therefore decrease some of the stochasticity around the globally optimal schedule of the raw simulated annealing algorithm's output, we tested a small

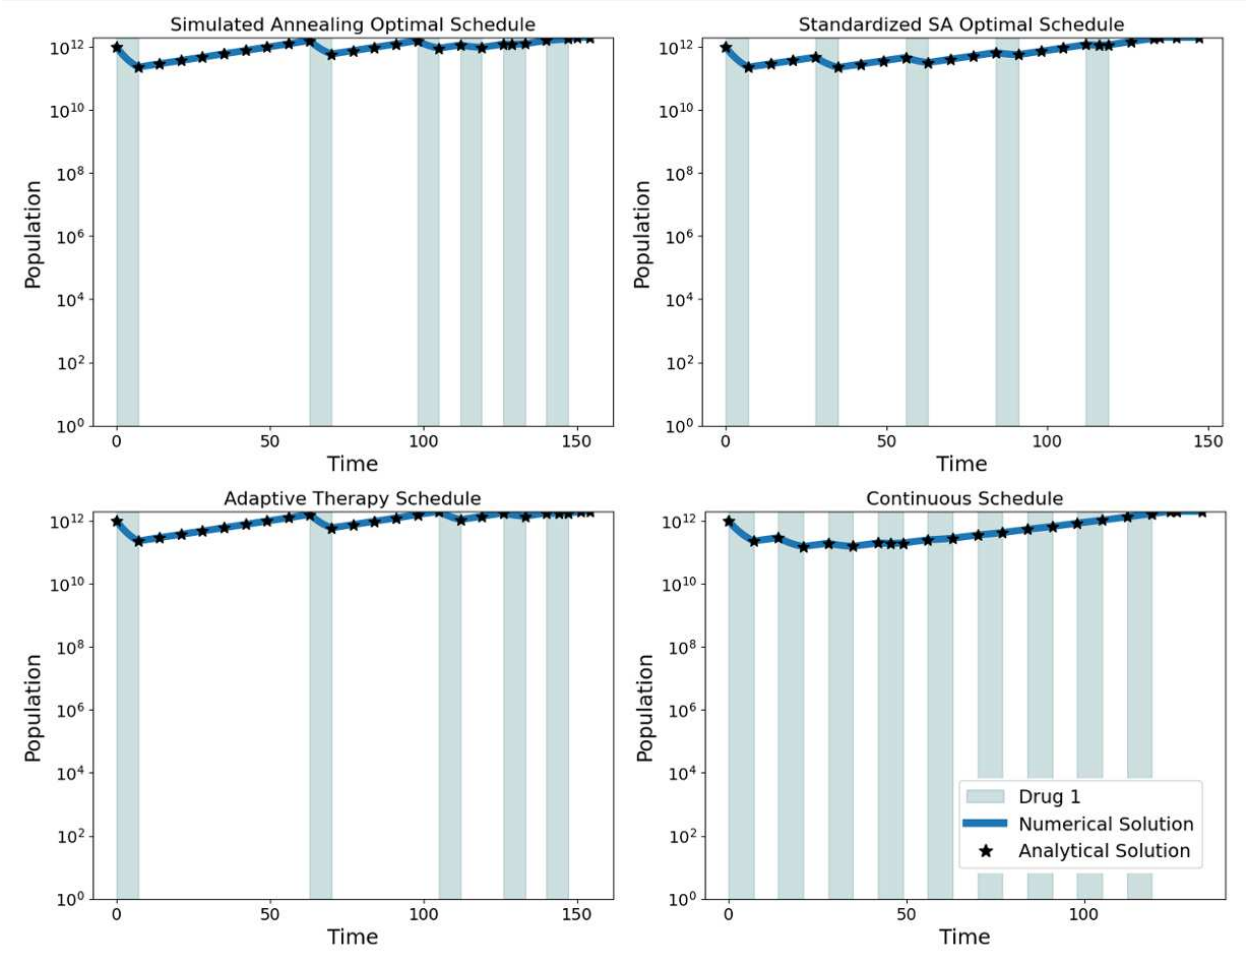


FIG. S6. **The additional tested schedules and the simulated annealing optimum for one drug.** We test the simulated annealing output against a standardized version of itself, an adaptive therapy schedule, and a continuous dosing schedule. The optimal schedule, which in this case (trying to maximize PFS) was the adaptive therapy schedule, is the one that has the lowest energy (including tiebreakers). PFS values were 252 days for the simulated annealing schedule, 230 days for the standardized SA schedule, 255 days for the adaptive therapy schedule, and 184 days for the continuous therapy schedule.

number of other schedules in addition to those sampled during the simulated annealing, as shown in Fig S6. These schedules were ones we expected to be optimal for various parameters. For one drug, they included continuous therapy (dose every available time), a form of adaptive therapy (only dose when the patient would progress the next cycle without it), and a "standardized" version of the simulated annealing output that used a constant period between cycles that matched the rounded average period between cycles of the simulated

annealing algorithm's returned optimal schedule. The schedule with the lowest energy (including tiebreakers) was then reported as the optimal schedule. To further illustrate the utility of this additional testing, for the parameters used to create Fig S5, adding just the standardization resulted in the correct optimal schedule being found 99/100 times (see the code). Explicitly testing the continuous dosing schedule would increase this accuracy to 100/100, since continuous dosing is the exact global optimum in this case.

For two drugs, there are far more heuristic schedules that could be optimal and could be separately tested for each run of the algorithm. As shown in Fig 4(A-C), we chose to test the sequential, alternating, and second-strike schedules, of both drug orders, since these schedules were the ones we are clinically interested in.

Appendix C: Further information on the phase boundaries in the one drug case

One drug. To investigate the anomalous points in Fig 3(E) that we pointed out in the main text, we needed to create analytical metrics for predicting which strategy should be optimal on the PFS phase diagram. For example, we can analytically calculate when "curative cycles" that keep E unchanged but decrease T exist for GDRS parameters (γ, δ, r, s), giving E-type trajectories that lead to ever-decreasing tumor burden. Setting $D(t) \in \{0, 1\}$, the GDRS equations for one drug are

$$\frac{dT}{dt} = \begin{cases} \gamma T \,, & \text{drug off} \\ (\gamma - \delta E)T \,, & \text{drug on} \end{cases} , \tag{S1}$$

$$\frac{dE}{dt} = \begin{cases} sE(1-E) \,, & \text{drug off} \\ -rE \,, & \text{drug on} \end{cases} . \tag{S2}$$

We want to find a cycle that starts at (T_0, E_0), applies drug for time dt_1, leaves drug off for time dt_2, and then ends up at $(T_0 + dT < T_0, E_0)$. During the first leg of the cycle, E changes by an amount $dE_1 = -rE_0 dt_1$. During the second leg, E changes by an amount $dE_2 = sE_0(1 - E_0)dt_2 + O(dt^2)$. For E to return to the same value E_0 at the end of the two-step cycle, we must have $dE_1 = -dE_2$, or

$$dt_2 = \frac{r dt_1}{s(1 - E_0)} \,. \tag{S3}$$

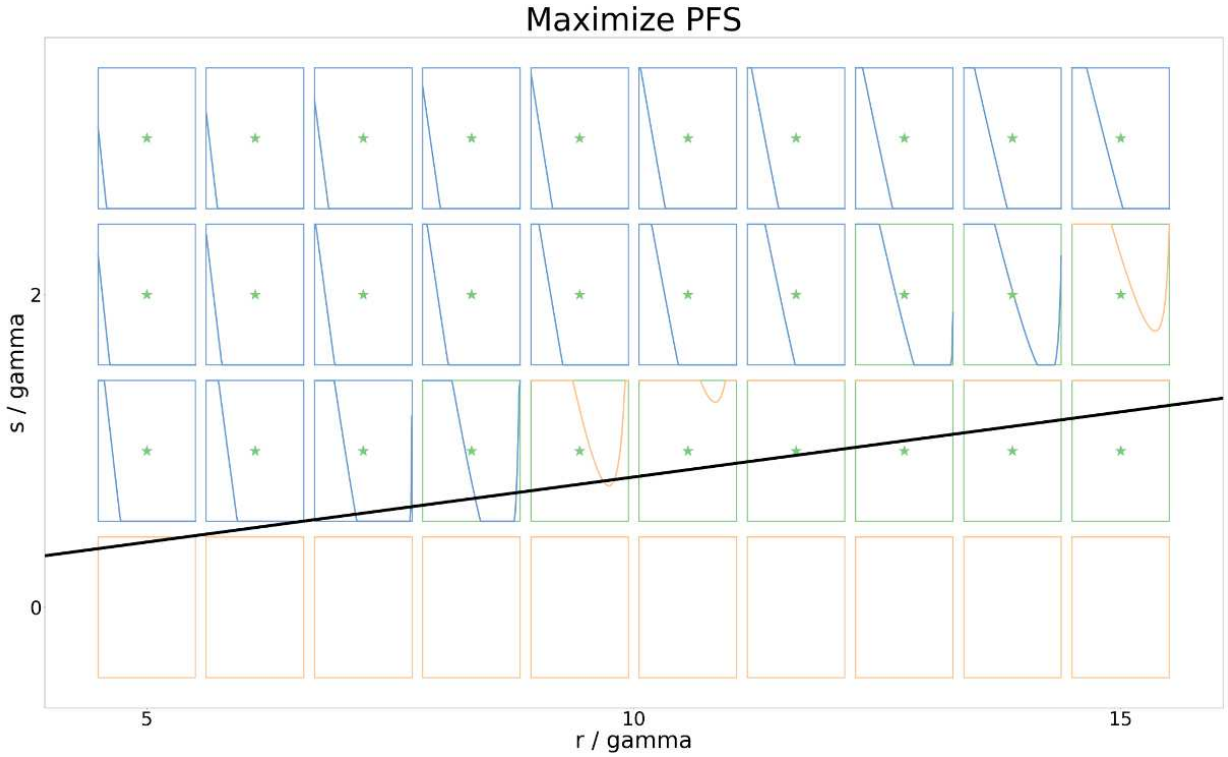


FIG. S7. **Examining the E-type/T-type phase boundary using different curative cycle conditions.** Each box (color coded as in the PFS phase diagram of Fig 3(E)) shows Eq (S13) plotted on the restricted range $[0, 1]$. When it has a root, a curative cycle is possible (blue); otherwise, a curative cycle is not possible (orange). Plots with a green star favor adaptive therapy (delaying a dose gives a PFS benefit). Points above the black line have a curative cycle when infinitesimal switches are allowed, as predicted by Eq (S6); however, some of these schedules become inaccessible due to the discrete dosing cycle time of 7 days.

The change in T over the first leg of the cycle is $dT_1 = (\gamma - \delta E_0) T_0 dt_1$. Over the second leg, it is

$$dT_2 = \gamma T_0 dt_2 = \frac{\gamma T_0 r dt_1}{s(1 - E_0)} \ . \tag{S4}$$

Combining, we find the total change in T is

$$dT = dT_1 + dT_2 = \left[\gamma - E_0 \delta + \frac{\gamma r}{s(1 - E_0)} \right] T_0 dt_1 \ . \tag{S5}$$

The sign of dT is determined by the term in brackets in Equation (S5), which is equivalent to a quadratic in E_0. If the quadratic has a root, then the curative cycle is possible. The

corresponding condition is

$$\frac{\delta}{4\gamma}\left(1+\frac{\gamma}{\delta}\right)^2 - 1 > \frac{r}{s}\,, \tag{S6}$$

which gives a linear boundary for the phase diagram, shown as the black line in Fig S7.

However, this curative schedule is not always accessible due to the discretization of time. Infinitesimal switching times are not allowed. The drug must be active for at least 7 days, and then it must be off for at least 7 more days. Without the discretization, every schedule above the black phase boundary of Fig S7 would be exactly E-type and curative, making for a sharp phase boundary (at least on that side; this argument does not comment on the optimal schedule on the other side of the boundary). The chunking of time into weekly cycles blurs the boundary, creating behavior like a second-order phase transition.

Because of this difficulty, we sought another metric to characterize the phase boundary in the case of weekly cycles. This metric is the result of the same calculation as above, but now $dt_i \to \Delta t_i$ becomes a finite time, with $\Delta t_1 = 7$ days. Instead of being able to use Eqs (S1) and (S2), we need the solutions evaluated at Δt_1 and $\Delta t_1 + \Delta t_2$:

$$E(\Delta t_1) = E_0 e^{-r\Delta t_1}\,, \tag{S7}$$

$$T(\Delta t_1) = T_0 e^{\gamma \Delta t_1} \exp\left[\frac{E_0\delta}{r}\left(e^{-r\Delta t_1}-1\right)\right]\,, \tag{S8}$$

$$E(\Delta t_1 + \Delta t_2) = \frac{E(\Delta t_1)}{E(\Delta t_1) + (1-E(\Delta t_1))e^{-s\Delta t_2}}\,, \tag{S9}$$

$$T(\Delta t_1 + \Delta t_2) = T(\Delta t_1) e^{\gamma \Delta t_2}\,. \tag{S10}$$

Setting $E(\Delta t_1 + \Delta t_2) = E_0$ allows us to find

$$\Delta t_2 = \frac{1}{s}\ln\left[\frac{1/E(\Delta t_1) - 1}{1/E_0 - 1}\right]\,. \tag{S11}$$

Using this expression for Δt_2 and setting $T(\Delta t_1 + \Delta t_2) < T_0$ gives a condition that must be satisfied for a curative cycle with $\Delta t_1 = 7$ days to exist:

$$T_0 > T_0 \exp\left[\gamma\Delta t_1 + \frac{E_0\delta}{r}\left(e^{-r\Delta t_1}-1\right) + \frac{\gamma}{s}\ln\left(\frac{e^{r\Delta t_1}/E_0 - 1}{1/E_0 - 1}\right)\right]\,. \tag{S12}$$

Upon dividing by T_0 and taking the natural log, we get

$$0 = \gamma \Delta t_1 + \frac{E_0 \delta}{r}\left(e^{-r\Delta t_1} - 1\right) + \frac{\gamma}{s} \ln\left[\frac{1 - E_0 e^{-r\Delta t_1}}{(1 - E_0)e^{-r\Delta t_1}}\right] . \tag{S13}$$

If there exists an E_0 that satisfies this equation, then a curative cycle with $\Delta t_1 = 7$ days is theoretically possible. However, there is no reason Δt_2 should be a multiple of 7 days, and there is no reason why E_0 should be exactly attainable from the simulation's starting value of $E(0) = 1$ (e.g. treating for one week may overshoot E_0).

In Fig S7, we plot the right-hand side of Eq (S13) as E_0 varies from 0 to 1 for each point we tested in phase space. For clarity, we have restricted its range to $[0, 1]$. The color of the curve denotes if a root exists (E-type favored, blue) or not (E-type not favored, orange). The color of the outline box corresponds to the strategy sorting we did in Fig 3(E). Finally, we tested if the PFS time was greater for a schedule that started with two consecutive doses or for a schedule with the second dose delayed until just before progression. Such a set of parameter values would favor an adaptive therapy; these are denoted with green stars. Note that the curative phase boundary line defined by the roots of Eq (S13) is not the same as the one corresponding to Eq (S6) (black line) where infinitesimal cycling is allowed; as expected, the infinitesimal cycle boundary is more permissive, saying that some parameters where we found only T-type adaptive therapy as the optimal schedule actually permit a curative schedule.

In the three anomalous cases (blue curves but green outline boxes), we can see that Eq (S13) does have a root, so that in theory there is a curative E-type schedule. However, this schedule cannot be found because of the discretization of time into 7 day cycles. Accordingly, the best that the algorithm can do in the PFS case is the T-type schedule. In this sense, the phase transition here acts like a second-order phase transition, where properties of both the T-type and E-type schedules can emerge near the boundary, for both optimization objectives.

Our results have largely used fixed values for δ and γ. Eq (S6) tells us directly how the curative cycle boundary (and thus the E-type/T-type phase boundary) shifts as δ and γ change. Unsurprisingly, as shown in Fig S8, decreasing δ or increasing γ makes it more difficult to have a curative cycle, so the region of parameter space where curative cycles are possible gets smaller (the phase boundary shifts upward, toward the s-axis).

Two drugs. In the main text, we analyzed a somewhat special case of the two drug optimization problem where the re-sensitization of the second drug was zero, $s_2 = 0$, and

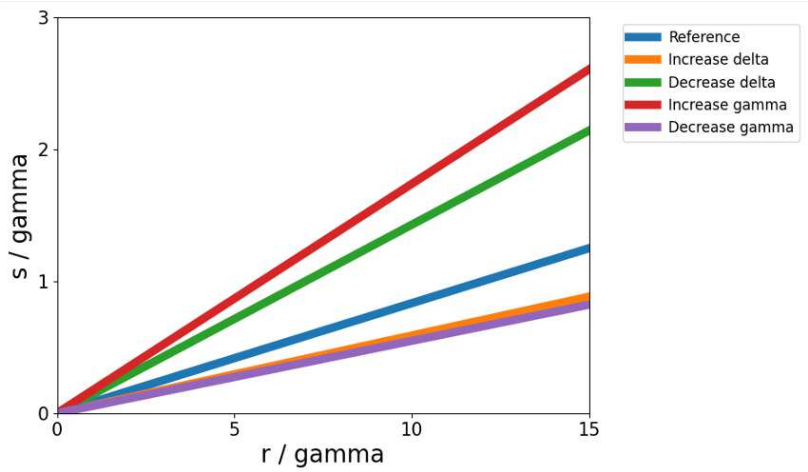


FIG. S8. **Shifting the curative cycle boundary of Eq** (S6) **by varying** δ **and** γ**.** Increasing γ or decreasing δ makes it harder to have curative cycles, decreasing the amount of parameter space for which these cycles are possible. Conversely, increasing δ or decreasing γ makes curative cycles more attainable, increasing the size of parameter space over which these cycles are possible.

the efficacy of the second drug is essentially entirely used up in one treatment window. Because of these characteristics, the scheduling optimization problem collapses to the one drug scenario after the use of the second drug (which typically happens very early in the trajectory, during the first few treatment windows). Therefore, it should be possible to use the categories of the one drug solution to categorize the adaptive therapy optimal schedules in the case of maximizing PFS. But now, the analogous schedules would use the second drug as much as possible in the treatment breaks.

The result of this analysis is shown in Fig S9, where the green denotes the T-type analog and blue denotes the rest, including the E-type analog. The phases seem to be quite mixed in this case, with no apparent phase boundary. However, upon inspecting the trajectories (e.g. see the code), this phenomenon can be explained by the second-order-like phase transition we discussed above. The black line marks the (lower) edge of parameter space where curative infinitesimal cycles of drug one are possible, as predicted by Eq (S6). Some of the schedules above have curative cycles in the discrete cycle time case, as shown in Fig 4. The others here either do not, or could not be run long enough in the simulated annealing algorithm to find the curative schedule. In particular, for the green squares that appear above the black line, the T-type adaptive therapies tested were able to achieve much longer times than

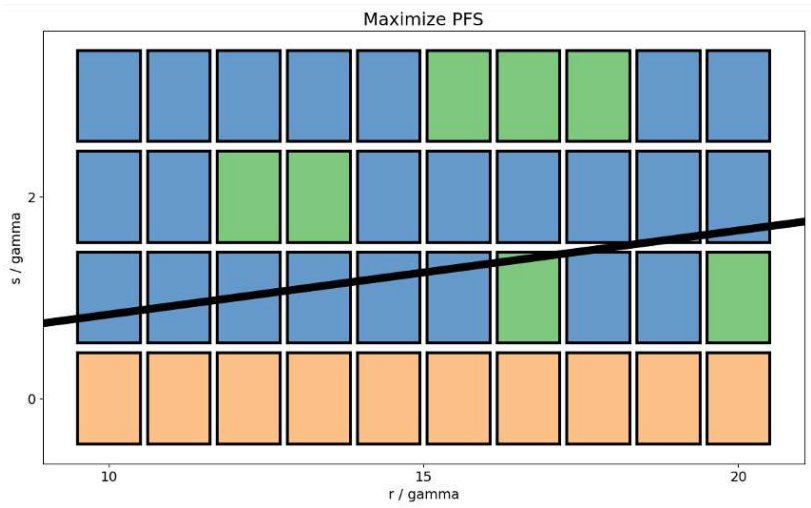


FIG. S9. **Alternate optimal schedule sorting of the two drug phase boundary from Fig 4G.** Because of the special parameter values of the second drug, we were able to map the adaptive therapy schedules onto the one drug case (orange is still the sequential schedule): the T-type analog schedule is shown in green, while all others, including the E-type schedule, are shown as blue. Infinitesimal curative cycles of the first drug exist above the black line, which was calculated using Eq (S6). As before, the second-order nature of the phase boundary causes anomalies and fuzziness, although it is much more pronounced in this case than in Fig S7.

the simulated annealing algorithm could sample. We expect E-type-like schedules to again emerge as the optimum for these cases if the simulated annealing algorithm is permitted to run for much longer.

Below the infinitesimally curative boundary, there are many blue boxes. These points largely correspond to schedules that blend together the E- and T-type strategies, as is the behavior of the second-order phase transition. In the two drug case, the effect is now more pronounced because the use of the second drug creates such a large gap between the tumor population minimum and the progression threshold. So once again, the discretization of time into cycles and the nature of the resulting phase transition explain the anomalous behavior of the phase diagram.

Appendix D: Specific parameters used to create the figures in this paper

We wanted to focus the model on patient-realistic regions of GDRS parameter space. Since the standard of care for Ewing sarcoma is the alternating VDC/IE chemotherapy regimen with two week cycles [30], we used it as the basis for our parameter selection. We set our cycle time smaller, to one week, so we had more flexibility in timing (that is, so doses could be moved around with one week resolution instead of two week resolution). Lung metastases from bone tumors have a doubling time of around 23 days [69], so we generally chose our growth rate $\gamma = \log(2)/20$ to be near this value. We then picked some baseline values for δ, r, and s that created "reasonable" trajectories of the tumor burden, i.e. ones that could create a multitude of responses (non-responders, stable disease, partial response, and complete response, as determined by a tumor detection threshold of $\sim 10^9$ cells). These explorations informed our later choices of parameters that went into creating the figures, as detailed below.

Our simulations were greatly sped up by using the analytical solutions of Eqs (S1) and (S2) in the simulated annealing algorithm instead of using the numerical solution for the computations. We only use the numerical solution for visualization at the end, when we plot both it and the analytical solution; e.g. see Fig S5, where the analytical solution is denoted by the black stars. We will now explain how this solution was utilized in the algorithm.

The analytical solutions to Eqs (S1) and (S2) are

$$T(t) = \begin{cases} T_0 e^{\gamma t} \,, & \text{drug off} \\ T_0 e^{\gamma t} \exp\left[\frac{E_0 \delta}{r}\left(e^{-rt} - 1\right)\right] \,, & \text{drug on} \end{cases}, \tag{S1}$$

$$E(t) = \begin{cases} \frac{E_0}{E_0 + (1 - E_0)e^{-st}} \,, & \text{drug off} \\ E_0 e^{-rt} \,, & \text{drug on} \end{cases}, \tag{S2}$$

where, for the purposes of the simulation, drug is on or off for the entire cycle period and T_0 and E_0 are known from the previous cycle (that is, $t = 0$ is the current cycle start time). So as not to clutter the equations with indices, these equations assume only one drug. For more drugs, each $E_i(t)$ would need to be tracked for each cycle, using its own E_{i0}, δ_i, r_i, and s_i, as well as the drug on or off solution corresponding to whether that drug was on or off during that cycle. The full solution for the entire duration of the treatment is a piecewise

combination of these individual cycle solutions arranged and solved according to the drug schedule.

These analytical functional forms allow for easy evaluation of the values of T and E at the end of the cycle. However, we also need to know the minimum and maximum values that T attains during the period to check for progression or cure. The maximum value of T will always occur at the cycle start or end, so no additional calculations were needed to find the maximum. The minimum, however, can occur at an intermediate point when drug is on (due to resistance). There is an intermediate minimum if the sign of $\delta - \gamma E(t)$ changes from negative at the start of the cycle to positive at the end of the cycle. If this condition is satisfied, then the time at which the minimum occurred was

$$t_{\text{minimum}} = \frac{1}{r} \log\left(\frac{\delta E_0}{\gamma}\right) , \tag{S3}$$

where again $t = 0$ is set to be the cycle start time. Then, the intermediate values of T and E could be solved for using Eqs (S1) and (S2).

In addition to checking for and returning the values of the intermediate minima, we also stopped the dynamics (just returned the previous value every subsequent cycle) if the patient progressed or was cured. The patient was checked for progression or cure that occurred at intermediate times as well to achieve this goal. This set of rules explains the dynamics of the analytical solution black stars for every plot in which they appear. For full details of the implementation, see the function "GDRSanalytical" in the "schedule_utils.py" file of the code provided on GitHub (see code availability statement in the main text).

Finally, before we give a detailed account of each figure, we revisit the cooling schedule used throughout. The cooling schedule of Eq (4) is compared to the energy of the simulation in Fig S10. This cooling schedule allowed for a period of exploration (temperature higher than energy) and a period of exploitation (temperature lower than energy), giving generally good convergence near the global optimum for every parameter set tested.

Figure 1. For these figures, we used $\gamma = \log(2)/20$, $\delta_1 = \delta_2 = \delta_3 = 30\gamma$, $r_1 = r_2 = r_3 = 4\gamma$, and $s_1 = s_2 = s_3 = 0$. We then numerically simulated the GDRS model for the shown schedules of drug on/off.

Figure 3. (A) We found this continuous schedule to be optimal for $\gamma = \log(2)/20$, $\delta = 30\gamma$, $r = 2\gamma$, and $s = 0$ when trying to minimize the tumor burden. (B) We found this E-type adaptive therapy schedule to be optimal for $\gamma = \log(2)/20$, $\delta = 25\gamma$, $r = 6\gamma$,

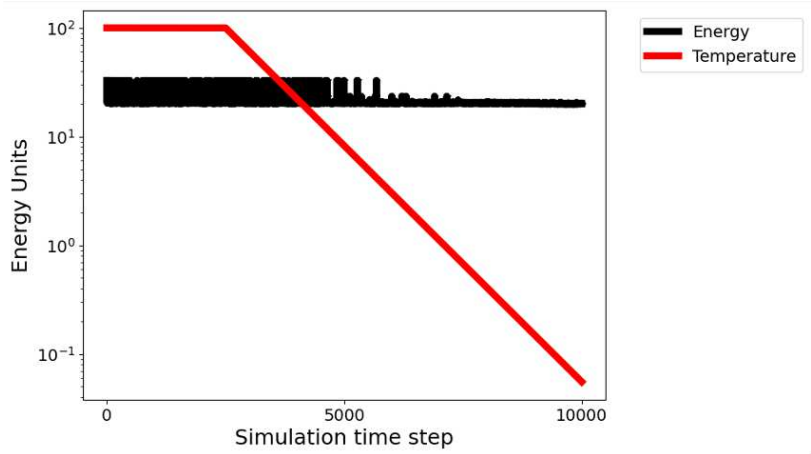


FIG. S10. **The cooling schedule used for the simulated annealing algorithm, Eq** (4)**,** demonstrated here using the simulation results from Fig S6.

and $s = 7\gamma$ when trying to minimize the tumor burden. (C) We found this T-type adaptive therapy to be optimal for $\gamma = \log(2)/20$, $\delta = 25\gamma$, $r = 6\gamma$, and $s = \gamma$ when trying to maximize PFS. (D and E) With $\gamma = \log(2)/20$ and $\delta = 50\gamma$, we sampled 40 points in $r - s$ space, evenly spaced between $r = 5\gamma$ to 15γ and $s = 0$ to 3γ for each objective.

Figure 4. (A) We found this sequential MTD schedule to be optimal for $\gamma = \log(2)/20$, $\delta_1 = \delta_2 = 40\gamma$, $r_1 = 6\gamma$, $r_2 = 10\gamma$, and $s_1 = s_2 = 0$ when trying to maximize PFS. (B) We found this alternating schedule to be optimal for $\gamma = \log(2)/20$, $\delta_1 = 50\gamma$, $\delta_2 = 30\gamma$, $r_1 = 3\gamma$, $r_2 = 2\gamma$, and $s_1 = s_2 = 0$ when trying to maximize the PFS. (C) We found this second-strike therapy to be optimal for $\gamma = \log(2)/20$, $\delta_1 = 50\gamma$, $\delta_2 = 40\gamma$, $r_1 = 18\gamma$, $r_2 = 4\gamma$, $s_1 = \gamma$, and $s_2 = 0$ when trying to minimize the tumor burden. (D) We found this curative adaptive therapy to be optimal for $\gamma = \log(2)/20$, $\delta_1 = 50\gamma$, $\delta_2 = 40\gamma$, $r_1 = 10\gamma$, $r_2 = 4\gamma$, $s_1 = 2\gamma$, and $s_2 = 0$ when trying to minimize the tumor burden. (E) We found this non-curative adaptive therapy to be optimal for $\gamma = \log(2)/20$, $\delta_1 = 50\gamma$, $\delta_2 = 40\gamma$, $r_1 = 18\gamma$, $r_2 = 4\gamma$, $s_1 = \gamma$, and $s_2 = 0$ when trying to maximize the PFS. (F and G) With $\gamma = \log(2)/20$, $\delta_1 = 50\gamma$, $\delta_2 = 40\gamma$, $r_2 = 7\gamma$, and $s_2 = 0$, we sampled 40 points in $r_1 - s_2$ phase space, evenly spaced between $r_1 = 10\gamma$ to 20γ and $s_1 = 0$ to 3γ for each objective.

Figure S4. For the example schedule maximizing the PFS, we used $\gamma = \log(2)/20$, $\delta = 25\gamma$, $r = 15\gamma$, and $s = \gamma/2$. For the example schedule minimizing the tumor burden, we used $\gamma = \log(2)/20$, $\delta = 50\gamma$, $r = 6\gamma$, and $s = 45\gamma/4$.

Figure S5. Here we minimized the tumor burden using $\gamma = \log(2)/20$, $\delta = 10\gamma$, $r = 2\gamma$,

and $s = 10\gamma$.

Figure S6. Here we maximized the PFS using $\gamma = \log(2)/20$, $\delta = 10\gamma$, $r = 3\gamma$, and $s = \gamma/2$.

Figure S7. We used the same values as for Figure 3(E), plotting the additional equations as described above.

Figure S8. Here we plotted $r(s)$ for $s = 0$ to $3\log(2)/20$ using Eq (S6) for the following various parameter combinations. Reference: $\gamma = \log(2)/20$, $\delta = 5\log(2)/2$. Increase delta: $\gamma = \log(2)/20$, $\delta = 7\log(2)/2$. Decrease delta: $\gamma = \log(2)/20$, $\delta = 3\log(2)/2$. Increase gamma: $\gamma = \log(2)/10$, $\delta = 5\log(2)/2$. Decrease gamma: $\gamma = \log(2)/30$, $\delta = 5\log(2)/2$.

Figure S9. We used the same parameters as Fig 4(G).

[1] M. Strobl, J. Gallaher, M. Robertson-Tessi, J. West, and A. Anderson, Treatment of evolving cancers will require dynamic decision support, Annals of Oncology **34**, 867 (2023).

[2] C. H. Takimoto, Maximum tolerated dose: clinical endpoint for a bygone era?, Targeted oncology **4**, 143 (2009).

[3] J. Fourie Zirkelbach, M. Shah, J. Vallejo, J. Cheng, A. Ayyoub, J. Liu, R. Hudson, R. Sridhara, G. Ison, L. Amiri-Kordestani, *et al.*, Improving dose-optimization processes used in oncology drug development to minimize toxicity and maximize benefit to patients, Journal of Clinical Oncology **40**, 3489 (2022).

[4] W. Hryniuk, More is better., Journal of clinical oncology **6**, 1365 (1988).

[5] V. T. DeVita Jr and E. Chu, A history of cancer chemotherapy, Cancer research **68**, 8643 (2008).

[6] R. Simon and L. Norton, The norton–simon hypothesis: designing more effective and less toxic chemotherapeutic regimens, Nature Clinical Practice Oncology **3**, 406 (2006).

[7] M. Perry, D. Doll, and C. Freter, eds., *Perry's The Chemotherapy Source Book*, 5th ed. (Wolters Kluwer Health, 2012).

[8] L. A. Levit, M. Shah, M. J. Ratain, E. Garrett-Mayer, A. Rahman, M. Theoret, and R. D. Harvey, Totality of the evidence: optimizing dosage selection strategies in oncology, Journal of Clinical Oncology **43**, 2827 (2025).

[9] E. L. Korn and B. Freidlin, 18 - clinical trial designs in oncology, in *Abeloff's Clinical Oncology (Sixth Edition)*, edited by J. E. Niederhuber, J. O. Armitage, M. B. Kastan, J. H. Doroshow, and J. E. Tepper (Elsevier, Philadelphia, 2020) sixth edition ed., pp. 296–307.e2.

[10] R. A. Gatenby and J. S. Brown, Integrating evolutionary dynamics into cancer therapy, Nature reviews Clinical oncology **17**, 675 (2020).

[11] P. M. Enriquez-Navas, J. W. Wojtkowiak, and R. A. Gatenby, Application of evolutionary principles to cancer therapy, Cancer research **75**, 4675 (2015).

[12] A. F. Read, T. Day, and S. Huijben, The evolution of drug resistance and the curious orthodoxy of aggressive chemotherapy, Proceedings of the National Academy of Sciences **108**, 10871 (2011).

[13] J. Zhang, J. J. Cunningham, J. S. Brown, and R. A. Gatenby, Integrating evolutionary dynam-

ics into treatment of metastatic castrate-resistant prostate cancer, Nature communications **8**, 1816 (2017).

[14] R. A. Gatenby, A. S. Silva, R. J. Gillies, and B. R. Frieden, Adaptive therapy, Cancer research **69**, 4894 (2009).

[15] R. A. Gatenby, J. Zhang, and J. S. Brown, First strike–second strike strategies in metastatic cancer: lessons from the evolutionary dynamics of extinction, Cancer research **79**, 3174 (2019).

[16] J. Goldie, A. Coldman, and G. Gudauskas, Rationale for the use of alternating non-cross-resistant chemotherapy, Cancer Treat Rep **66**, 439 (1982).

[17] K. Gallagher, M. A. R. Strobl, A. R. A. Anderson, and P. K. Maini, Deriving optimal treatment timing for adaptive therapy: Matching the model to the tumor dynamics, Bulletin of Mathematical Biology **87**, 146 (2025).

[18] A. Holzgreve, A. Delker, Z. Ells, J. Brosch-Lenz, L. M. Unterrainer, J. Nikitas, S. Zhu, M. M. Contreras, H. Alam, R. M. Nabong, *et al.*, Randomized phase 2 trial of an extended and flexible dosing schedule of 177lu-psma molecular radiotherapy in patients with metastatic castration-resistant prostate cancer (flex-mrt): Study protocol, Journal of Nuclear Medicine **66**, 1639 (2025).

[19] R. A. Gatenby, Y. Artzy-Randrup, T. Epstein, D. R. Reed, and J. S. Brown, Eradicating metastatic cancer and the eco-evolutionary dynamics of anthropocene extinctions, Cancer research **80**, 613 (2020).

[20] D. R. Reed, J. Metts, M. Pressley, B. L. Fridley, M. Hayashi, M. S. Isakoff, D. M. Loeb, R. Makanji, R. D. Roberts, M. Trucco, *et al.*, An evolutionary framework for treating pediatric sarcomas, Cancer **126**, 2577 (2020).

[21] J. Metts, T. Russell, D. Reed, and M. Trucco, A proposed trial design for the treatment of widely metastatic ewing sarcoma inspired by evolutionary dynamics, Cancers **14**, 736 (2022).

[22] D. R. Reed, A. Tulpule, J. Metts, M. Trucco, M. Robertson-Tessi, T. J. O'Donohue, F. Iglesias-Cardenas, M. S. Isakoff, A. Mauguen, N. Shukla, *et al.*, Pediatric leukemia roadmaps are a guide for positive metastatic bone sarcoma trials, Journal of Clinical Oncology **42**, 2955 (2024).

[23] E. B. Heikamp and C.-H. Pui, Next generation evaluation and treatment of pediatric acute lymphoblastic leukemia, The Journal of pediatrics **203**, 14 (2018).

[24] A. G. Waks and E. P. Winer, Breast cancer treatment: a review, Jama **321**, 288 (2019).

[25] J. L. Blum, P. J. Flynn, G. Yothers, L. Asmar, C. E. Geyer Jr, S. A. Jacobs, N. J. Robert,

J. O. Hopkins, J. A. O'Shaughnessy, C. T. Dang, *et al.*, Anthracyclines in early breast cancer: The abc trials—usor 06-090, nsabp b-46-i/usor 07132, and nsabp b-49 (nrg oncology), Journal of clinical oncology **35**, 2647 (2017).

[26] I. C. Henderson, D. A. Berry, G. D. Demetri, C. T. Cirrincione, L. J. Goldstein, S. Martino, J. N. Ingle, M. R. Cooper, D. F. Hayes, K. H. Tkaczuk, *et al.*, Improved outcomes from adding sequential paclitaxel but not from escalating doxorubicin dose in an adjuvant chemotherapy regimen for patients with node-positive primary breast cancer, Journal of clinical oncology **21**, 976 (2003).

[27] A. Coates, V. Gebski, J. F. Bishop, P. N. Jeal, R. L. Woods, R. Snyder, M. H. Tattersall, M. Byrne, V. Harvey, G. Gill, *et al.*, Improving the quality of life during chemotherapy for advanced breast cancer, New England Journal of Medicine **317**, 1490 (1987).

[28] G. Bonadonna, M. Zambetti, A. Moliterni, L. Gianni, and P. Valagussa, Clinical relevance of different sequencing of doxorubicin and cyclophosphamide, methotrexate, and fluorouracil in operable breast cancer, Journal of clinical oncology **22**, 1614 (2004).

[29] H. E. Grier, M. D. Krailo, N. J. Tarbell, M. P. Link, C. J. Fryer, D. J. Pritchard, M. C. Gebhardt, P. S. Dickman, E. J. Perlman, P. A. Meyers, *et al.*, Addition of ifosfamide and etoposide to standard chemotherapy for ewing's sarcoma and primitive neuroectodermal tumor of bone, New England Journal of Medicine **348**, 694 (2003).

[30] T. Cash, M. D. Krailo, A. B. Buxton, B. R. Pawel, J. H. Healey, O. Binitie, K. J. Marcus, H. E. Grier, P. J. Grohar, D. R. Reed, *et al.*, Long-term outcomes in patients with localized ewing sarcoma treated with interval-compressed chemotherapy on children's oncology group study aews0031, Journal of Clinical Oncology **41**, 4724 (2023).

[31] J. K. Anninga, H. Gelderblom, M. Fiocco, J. R. Kroep, A. H. Taminiau, P. C. Hogendoorn, and R. M. Egeler, Chemotherapeutic adjuvant treatment for osteosarcoma: where do we stand?, European journal of cancer **47**, 2431 (2011).

[32] J. Gill and R. Gorlick, Advancing therapy for osteosarcoma, Nature reviews Clinical oncology **18**, 609 (2021).

[33] I. Magrath, M. Adde, A. Shad, D. Venzon, N. Seibel, J. Gootenberg, J. Neely, C. Arndt, M. Nieder, E. Jaffe, *et al.*, Adults and children with small non-cleaved-cell lymphoma have a similar excellent outcome when treated with the same chemotherapy regimen., Journal of Clinical Oncology **14**, 925 (1996).

[34] D. A. Thomas, J. Cortes, S. O'Brien, S. Pierce, S. Faderl, M. Albitar, F. B. Hagemeister, F. F. Cabanillas, S. Murphy, M. J. Keating, *et al.*, Hyper-cvad program in burkitt's-type adult acute lymphoblastic leukemia, Journal of Clinical Oncology **17**, 2461 (1999).

[35] L. Fayad, D. Thomas, and J. Romaguera, Update of the md anderson cancer center experience with hyper-cvad and rituximab for the treatment of mantle cell and burkitt-type lymphomas, Clinical Lymphoma and Myeloma **8**, S57 (2007).

[36] D. R. Feldman, J. Sheinfeld, D. F. Bajorin, P. Fischer, S. Turkula, N. Ishill, S. Patil, M. Bains, L. M. Reich, G. J. Bosl, *et al.*, Ti-ce high-dose chemotherapy for patients with previously treated germ cell tumors: results and prognostic factor analysis, Journal of Clinical Oncology **28**, 1706 (2010).

[37] A. Carrato, R. Pazo-Cid, T. Macarulla, J. Gallego, P. Jiménez-Fonseca, F. Rivera, M. T. Cano, M. Rodriguez-Garrote, C. Pericay, I. Alés, *et al.*, Nab-paclitaxel plus gemcitabine and folfox in metastatic pancreatic cancer, NEJM evidence **3**, EVIDoa2300144 (2024).

[38] C. V. Fernandez, E. A. Mullen, Y.-Y. Chi, P. F. Ehrlich, E. J. Perlman, J. A. Kalapurakal, G. Khanna, A. C. Paulino, T. E. Hamilton, K. W. Gow, *et al.*, Outcome and prognostic factors in stage iii favorable-histology wilms tumor: a report from the children's oncology group study aren0532, Journal of Clinical Oncology **36**, 254 (2018).

[39] F. Spreafico, C. V. Fernandez, J. Brok, K. Nakata, G. Vujanic, J. I. Geller, M. Gessler, M. Maschietto, S. Behjati, A. Polanco, *et al.*, Wilms tumour, Nature reviews Disease primers **7**, 75 (2021).

[40] D. B. Dix, N. L. Seibel, Y.-Y. Chi, G. Khanna, E. Gratias, J. R. Anderson, E. A. Mullen, J. I. Geller, J. A. Kalapurakal, A. C. Paulino, *et al.*, Treatment of stage iv favorable histology wilms tumor with lung metastases: a report from the children's oncology group aren0533 study, Journal of Clinical Oncology **36**, 1564 (2018).

[41] D. B. Dix, C. V. Fernandez, Y.-Y. Chi, E. A. Mullen, J. I. Geller, E. J. Gratias, G. Khanna, J. A. Kalapurakal, E. J. Perlman, N. L. Seibel, *et al.*, Augmentation of therapy for combined loss of heterozygosity 1p and 16q in favorable histology wilms tumor: a children's oncology group aren0532 and aren0533 study report, Journal of Clinical Oncology **37**, 2769 (2019).

[42] P. Ehrlich, Y. Y. Chi, M. M. Chintagumpala, F. A. Hoffer, E. J. Perlman, J. A. Kalapurakal, A. Warwick, R. C. Shamberger, G. Khanna, T. E. Hamilton, *et al.*, Results of the first prospective multi-institutional treatment study in children with bilateral wilms tumor (aren0534): a

report from the children's oncology group, Annals of surgery **266**, 470 (2017).

[43] N. L. Seibel, Y.-Y. Chi, E. J. Perlman, J. Tian, J. Sun, J. R. Anderson, M. L. Ritchey, P. R. Thomas, J. Miser, J. A. Kalapurakal, *et al.*, Impact of cyclophosphamide and etoposide on outcome of clear cell sarcoma of the kidney treated on the national wilms tumor study-5 (nwts-5), Pediatric blood & cancer **66**, e27450 (2019).

[44] J. S. Dome, C. A. Cotton, E. J. Perlman, N. E. Breslow, J. A. Kalapurakal, M. L. Ritchey, P. E. Grundy, M. Malogolowkin, J. B. Beckwith, R. C. Shamberger, *et al.*, Treatment of anaplastic histology wilms' tumor: results from the fifth national wilms' tumor study, Journal of Clinical Oncology **24**, 2352 (2006).

[45] M. Gluzman, J. G. Scott, and A. Vladimirsky, Optimizing adaptive cancer therapy: dynamic programming and evolutionary game theory, Proceedings of the Royal Society B **287**, 20192454 (2020).

[46] K. Gallagher, M. Strobl, D. Park, F. Spendlin, R. Gatenby, P. Maini, and A. Anderson, Learning to adapt: Rational personalization of adaptive therapy using deep reinforcement learning, Cancer Research **84** (2024).

[47] D. Mathur, E. Barnett, H. I. Scher, and J. B. Xavier, Optimizing the future: how mathematical models inform treatment schedules for cancer, Trends in cancer **8**, 506 (2022).

[48] A. M. Jarrett, D. Faghihi, D. A. Hormuth, E. A. Lima, J. Virostko, G. Biros, D. Patt, and T. E. Yankeelov, Optimal control theory for personalized therapeutic regimens in oncology: Background, history, challenges, and opportunities, Journal of clinical medicine **9**, 1314 (2020).

[49] U. Ledzewicz and H. Schättler, On the role of the objective in the optimization of compartmental models for biomedical therapies, Journal of optimization theory and applications **187**, 305 (2020).

[50] K. Leder, K. Pitter, Q. LaPlant, D. Hambardzumyan, B. D. Ross, T. A. Chan, E. C. Holland, and F. Michor, Mathematical modeling of pdgf-driven glioblastoma reveals optimized radiation dosing schedules, Cell **156**, 603 (2014).

[51] A. Fu, Z. A. Gouw, J. Jeong, and J. O. Deasy, Optimal radiotherapy dose scheduling with variable fraction sizes and breaks via sequential mixed-integer convex programming, SSRN Pre-Print (2025).

[52] M. Robertson-Tessi, J. S. Brown, M. I. Poole, M. Johnson, A. Marusyk, J. A. Gallaher, K. A. Luddy, C. J. Whelan, J. West, M. Strobl, *et al.*, Feasibility of an evolutionary tumor board

for generating novel personalized therapeutic strategies, MedRxiv , 2023 (2023).

[53] M. Pincus, A monte carlo method for the approximate solution of certain types of constrained optimization problems, Operations research **18**, 1225 (1970).

[54] D. Bertsimas and J. Tsitsiklis, Simulated annealing, Statistical science **8**, 10 (1993).

[55] Y.-J. Yang, C. D. Kocher, and K. A. Dill, Nonequilibrium theory for adaptive systems in varying environments, arXiv preprint arXiv:2506.19018 (2025).

[56] A. E. Pomeroy, E. V. Schmidt, P. K. Sorger, and A. C. Palmer, Drug independence and the curability of cancer by combination chemotherapy, Trends in cancer **8**, 915 (2022).

[57] J. Jee, C. Fong, K. Pichotta, T. N. Tran, A. Luthra, M. Waters, C. Fu, M. Altoe, S.-Y. Liu, S. B. Maron, *et al.*, Automated real-world data integration improves cancer outcome prediction, Nature **636**, 1 (2024).

[58] H. R. Withers, The four r's of radiotherapy, in *Advances in radiation biology*, Vol. 5 (Elsevier, 1975) pp. 241–271.

[59] S. Eschrich, H. Zhang, H. Zhao, D. Boulware, J.-H. Lee, G. Bloom, and J. F. Torres-Roca, Systems biology modeling of the radiation sensitivity network: a biomarker discovery platform, International Journal of Radiation Oncology* Biology* Physics **75**, 497 (2009).

[60] J. G. Scott, A. Berglund, M. J. Schell, I. Mihaylov, W. J. Fulp, B. Yue, E. Welsh, J. J. Caudell, K. Ahmed, T. S. Strom, *et al.*, A genome-based model for adjusting radiotherapy dose (gard): a retrospective, cohort-based study, The lancet oncology **18**, 202 (2017).

[61] G. D. Grass, J. C. Alfonso, E. Welsh, K. A. Ahmed, J. K. Teer, S. Pilon-Thomas, L. B. Harrison, J. L. Cleveland, J. J. Mulé, S. A. Eschrich, *et al.*, The radiosensitivity index gene signature identifies distinct tumor immune microenvironment characteristics associated with susceptibility to radiation therapy, International Journal of Radiation Oncology* Biology* Physics **113**, 635 (2022).

[62] C. Mauz-Körholz, J. Landman-Parker, A. Fernández-Teijeiro, A. Attarbaschi, W. Balwierz, J. M. Bartelt, A. Beishuizen, S. Boudjemaa, M. Cepelova, F. Ceppi, *et al.*, Response-adapted omission of radiotherapy in children and adolescents with early-stage classical hodgkin lymphoma and an adequate response to vincristine, etoposide, prednisone, and doxorubicin (euronet-phl-c1): a titration study, The lancet oncology **24**, 252 (2023).

[63] M. K. Graeser, O. Gluz, C. zu Eulenburg, S. Kuemmel, M. Christgen, R. Wuerstlein, H. H. Kreipe, P. Schmid, M. Thill, M. W. Braun, *et al.*, Prediction of survival after de-escalated

neoadjuvant therapy in her2+ early breast cancer: A pooled analysis of three wsg trials. (2025).

[64] J. A. Scarborough, E. McClure, P. Anderson, A. Dhawan, A. Durmaz, S. L. Lessnick, M. Hitomi, and J. G. Scott, Identifying states of collateral sensitivity during the evolution of therapeutic resistance in ewing's sarcoma, Iscience **23** (2020).

[65] N. Yoon, R. Vander Velde, A. Marusyk, and J. G. Scott, Optimal therapy scheduling based on a pair of collaterally sensitive drugs, Bulletin of mathematical biology **80**, 1776 (2018).

[66] D. Nichol, J. Rutter, C. Bryant, A. M. Hujer, S. Lek, M. D. Adams, P. Jeavons, A. R. Anderson, R. A. Bonomo, and J. G. Scott, Antibiotic collateral sensitivity is contingent on the repeatability of evolution, Nature communications **10**, 334 (2019).

[67] R. Gomulkiewicz and R. D. Holt, When does evolution by natural selection prevent extinction?, Evolution **49**, 201 (1995).

[68] H. A. Orr and R. L. Unckless, Population extinction and the genetics of adaptation, The American Naturalist **172**, 160 (2008).

[69] S. B. Balmukhanov, I. Turdugulov, Z. Karibjanova, and L. Révész, The growth rate of bone sarcomas and survival after radiotherapy with tourniquet-induced hypoxia: A clinical study, Cancer **49**, 1597 (1982).