



# Delaying cancer progression by integrating toxicity constraints in a model of adaptive therapy



Jana L. Gevertz<sup>1,12</sup>✉, Harsh Vardhan Jain<sup>2,12</sup>, Irina Kareva<sup>3,4,12</sup>, Kathleen P. Wilkie<sup>5,12</sup>, Joel Brown<sup>6</sup>, Yitong Pepper Huang<sup>7</sup>, Eduardo Sontag<sup>8</sup>, Vladimir Vinogradov<sup>9,10</sup> & Mark Davies<sup>11</sup>✉

Cancer therapies often fail when intolerable toxicity or drug-resistant cancer cells undermine otherwise effective treatment strategies. Over the past decade, adaptive therapy has emerged as a promising approach to postpone emergence of resistance by altering dose timing based on tumor burden thresholds. Despite encouraging results, these protocols often overlook the crucial role of toxicity-induced treatment breaks, which may permit tumor regrowth. Herein, we explore the following question: would incorporating toxicity feedback improve or hinder the efficacy of adaptive therapy? To address this question, we propose a mathematical framework for incorporating toxic feedback into treatment design. We find that the degree of competition between sensitive and resistant populations, along with the growth rate of resistant cells, critically modulates the impact of toxicity feedback on time to progression. Further, our conceptual model identifies circumstances where strategic treatment breaks, which may be based on either tumor size or toxicity, can mitigate overtreatment and extend time to progression, both at the baseline parameterization and across a heterogeneous virtual population. Taken together, these findings highlight the importance of integrating toxicity considerations into the design of adaptive therapy.

Systemic anti-cancer therapies, including chemotherapy, targeted therapies, immunotherapies, and hormone treatments, are essential components of modern oncology. In patients with metastatic disease, where cure is often elusive, the primary goals are to control tumor growth, extend survival, minimize treatment-related toxicities, and alleviate cancer-associated symptoms<sup>1</sup>. Typically, a maximum tolerated dose (MTD) protocol is used to achieve these goals. The goal of MTD is to eliminate as many cancer cells as possible by administering the drug at the highest dose deemed safe until toxicity limitations are reached.

Despite significant advances in the precision and effectiveness of systemic anti-cancer therapies, two persistent challenges remain. The first is the management of treatment-induced toxicity<sup>2–4</sup>. Toxicities associated with MTD treatment include fatigue, gastrointestinal side effects, and hematological complications<sup>5–9</sup>. These toxicities adversely affect patient

quality of life and often necessitate dose reductions, treatment interruptions, or even discontinuation, all of which ultimately undermine therapeutic effectiveness.

The second persistent challenge of MTD therapy is the emergence of therapeutic resistance. In clinical practice, a tumor is considered sensitive if it shrinks or remains stable with treatment, and resistant if it continues to grow. However, because tumors are inherently heterogeneous, they typically contain subpopulations of cells with varying treatment sensitivities<sup>10</sup>. Under the sustained treatment typical of MTD therapy, the elimination of sensitive cells can create an environment in which resistant clones proliferate freely. This phenomenon, known as competitive release, accelerates resistance-driven progression and diminishes treatment efficacy over time<sup>11,12</sup>.

Adaptive therapy (AT) has emerged as a promising alternative to traditional MTD approaches. Rather than aiming for complete tumor

<sup>1</sup>Department of Mathematics and Statistics, The College of New Jersey, Ewing, NJ, USA. <sup>2</sup>Department of Mathematics and Statistics, University of Minnesota Duluth, Duluth, MN, USA. <sup>3</sup>Department of Biomedical Engineering, Northeastern University, Boston, MA, USA. <sup>4</sup>Quercus Insights, MA Boston, USA. <sup>5</sup>Department of Mathematics, Toronto Metropolitan University, Toronto, ON, Canada. <sup>6</sup>Department of Integrated Mathematical Oncology, Moffitt Cancer Center, Tampa, FL, USA. <sup>7</sup>Department of Mathematical Sciences, Smith College, Northampton, MA, USA. <sup>8</sup>Departments of Electrical and Computer Engineering and of Bioengineering, Northeastern University, Boston, MA, USA. <sup>9</sup>Department of Mathematics, Ohio University, Athens, OH, USA. <sup>10</sup>Department of Computer and Mathematical Sciences, University of Toronto Scarborough, Toronto, ON, Canada. <sup>11</sup>Division of Cancer and Genetics, Cardiff University, Cardiff, UK. <sup>12</sup>These authors contributed equally: Jana L. Gevertz, Harsh Vardhan Jain, Irina Kareva, Kathleen P. Wilkie. ✉e-mail: [gevertz@tcnj.edu](mailto:gevertz@tcnj.edu); [DaviesDM15@cardiff.ac.uk](mailto:DaviesDM15@cardiff.ac.uk)

eradication, AT strategically modulates dosing schedules to preserve a population of sensitive cells that can suppress resistant clones through intra-tumor competition<sup>13,14</sup>. Its implementation is governed by clinical decision rules based on measurable indicators, such as tumor size or circulating biomarkers. For example, a dose-skipping regimen administers a high dose until a predetermined tumor response (e.g., a 50% reduction in size) is achieved and then pauses treatment until the tumor regrows to a defined threshold, often its initial size. Alternatively, a dose-modulation strategy adjusts treatment doses incrementally, increasing or decreasing them on the basis of the tumor's response over time<sup>15</sup>.

The effectiveness of AT can be enhanced if resistance comes with a fitness cost. This occurs when the mechanisms conferring resistance, such as the upregulation of energy-intensive drug efflux pumps, reduce a cell's overall fitness in the absence of therapy<sup>16–18</sup>. Indeed, adaptive therapies have already been tested experimentally<sup>19–21</sup> and are currently being applied across multiple clinical trials (NCT02415621; NCT03511196; NCT03543969; NCT03630120).

A crucial, yet underexplored, aspect of AT is the impact of treatment-induced toxicity on dosing decisions. Toxicity constraints can limit treatment frequency or intensity, potentially allowing both sensitive and resistant cell populations to proliferate during treatment interruptions. This could disrupt the competitive dynamics that are essential to AT's success. Equally, integrating toxicity feedback into dosing strategies could improve patient tolerability and enable prolonged treatment, ultimately enhancing outcomes.

Attempts to model toxicity in the literature mainly focus on white blood cell counts and neutropenia induced by chemotherapy<sup>22–26</sup>. In ref. 27, the authors model chemo-induced toxicity effects as a loss of muscle mass, and in ref. 28, the authors model toxicity effects via both weight loss and a comprehensive index that takes into account 57 side effects. Toxicity has also been included as an additional constraint in the design of optimal control dosing strategies<sup>29,30</sup>.

Approaches to model toxicity include an assumption that the effect is present and constant during the whole treatment phase<sup>22,23</sup>, that it is proportional to drug concentration<sup>25,29,30</sup>, or that it is proportional to drug and tumor burden<sup>28</sup>. In ref. 24, the toxic effect of the drug directly reduces the neutrophil proliferation rate via an inverse polynomial form of the drug concentration. In ref. 22, a stochastic model is introduced wherein toxicity is modeled as an increase in the death rate of white blood cells. Through a virtual clinical trial, the authors demonstrate that managing toxicity throughout treatment by modulating or pausing doses does not compromise overall treatment outcomes. In ref. 31, the authors use a modified logistic competition model with sensitive and resistant cells to show that in numerical simulations, continuously-dosed AT outperforms discretely-dosed AT in terms of both time to disease progression and controlling toxicity. Their results assume toxicity is a nonlinear function of drug dose satisfying specific conditions (such as being strictly increasing and concave up).

Previously considered control problems<sup>30</sup> introduce toxicity as a constraint, but do not monitor the emergence of drug resistance or the strategy of AT. Dose interruptions in a model of resistance is presented in ref. 22, but in that paper, toxicity is considered as a constant effect that reduces white blood cell growth, and is not dependent on drug concentration or accumulated exposure. In ref. 31, where the authors consider continuous or intermittent adaptive therapies to manage emergence of resistance, toxicity is modeled as a side effect that can be used to determine when too much drug has been administered but is not included in dosing decision making.

In this work, we explore the following critical question: would incorporating toxicity feedback improve or hinder the efficacy of AT? To address this question, we propose a novel mathematical framework for including toxicity both in the decision-making process of daily therapy, and in the formulation of adaptive regimes based on maintaining the tumor size within a pre-specified range. Our simulations, designed to mimic a theoretical murine study with an unspecified chemotherapeutic agent, compare the impact of toxicity constraints on time to progression (TTP) for daily

treatment versus AT schedules. Our results demonstrate that the degree of competition between sensitive and resistant populations, along with the growth rate of resistant cells, critically modulates how TTP is impacted by the incorporation of toxicity feedback. Notably, our conceptual model identifies circumstances where strategic treatment breaks, which may be based on either tumor size or toxicity, can mitigate overtreatment and extend TTP, both at the baseline parameterization and across a virtual population. We conclude with a discussion of the model's limitations and propose future directions aimed at optimizing cancer treatment strategies to simultaneously address resistance and toxicity, thereby improving patient quality of life.

## Results

To investigate the relationship between AT and toxicity, we analyze the dynamics captured by a mathematical model that integrates drug pharmacokinetics and toxicity dynamics into an existing framework<sup>16</sup> of competitive interactions between treatment-sensitive ( $S(t)$ ) and treatment-resistant ( $R(t)$ ) cancer cells. This conceptual model, described in detail in the Methods, assumes that both subpopulations grow logistically to a shared carrying capacity and that only sensitive cells are killed by the drug. It further assumes that drug pharmacokinetics are described by a one-compartment model and that toxicity increases at a rate proportional to drug concentration and resolves at a linear rate.

In this section, we evaluate the impact of the four proposed protocols (Fig. 1) on treatment dynamics and time-to-progression (TTP): (i) daily protocol, (ii) adaptive protocol, (iii) daily protocol with toxicity feedback, and (iv) adaptive protocol with toxicity feedback. More details on the implementation of these protocols are provided in the Methods section.

In all simulations, parameters are fixed at their baseline values defined in Table 1 of the Methods, except for those being explicitly varied in a parameter sweep, with the same normalized dose administered across all treatment protocols. We monitor tumor composition (fractions of sensitive and resistant cells), drug concentration, and accumulated toxicity. Protocol effectiveness is assessed based on TTP—the longer a tumor takes to fail treatment, the more effective the protocol is deemed to be.

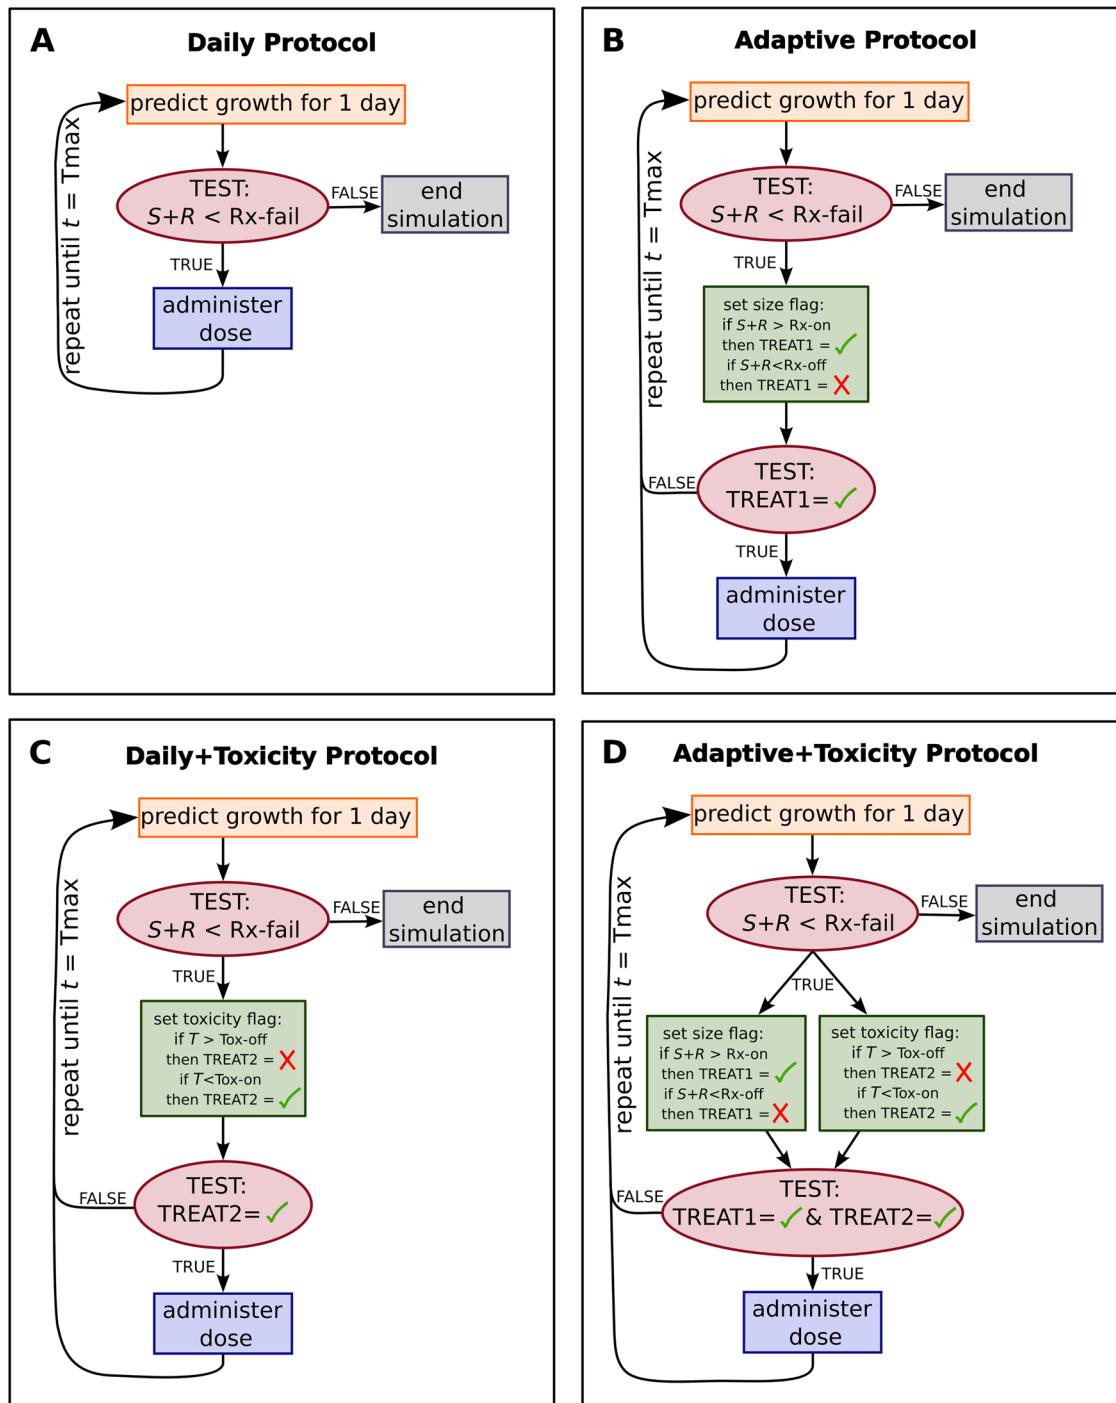
### Addition of Toxicity Constraints Can Extend TTP

In Fig. 2, we show representative model dynamics from a simulated experiment comparing the four treatment protocols using parameter values from Table 1. Each row corresponds to a protocol, with tumor cell time-courses plotted in the left panels, and drug concentration and toxicity time-courses in the right panels. Periods of drug administration are shaded in gray, and simulations run until the disease progresses. For each protocol, representative examples of model dynamics at alternative parameterizations are shown in Supplementary Fig. 1.

In the daily protocol (Fig. 2, top row), the drug rapidly depletes sensitive cells (blue curve), resulting in a large initial reduction in total tumor burden. However, this depletion removes intra-tumor competition, allowing resistant cells (red curve) to proliferate unchecked, leading to swift treatment failure. High and sustained drug exposure results in persistently elevated toxicity throughout treatment. This protocol yields the shortest TTP (19.2 days), creating a “lose-lose” scenario where resistance emerges quickly, and toxicity remains critically high.

In the adaptive protocol, tumor size-based feedback modulates treatment cycles, alternating between drug administration and off-treatment periods (Fig. 2, second row). This strategy allows sensitive cells to recover during treatment pauses, enabling them to compete with and temporarily suppress resistant cells, thereby delaying treatment failure. Periodic drug holidays also reduce toxicity compared to the daily protocol, though overall toxicity continues to increase with time. AT extends TTP to 44 days, more than doubling the duration achieved under the daily protocol. However, as resistant cells gradually escape suppression by the sensitive cells, the adaptive protocol ultimately loses effectiveness.

An additional observation in the adaptive protocol (Fig. 2C) is that the tumor volume sometimes overshoots the  $R_{x_0}$  threshold by a large margin



**Fig. 1 | Summary of the four proposed treatment protocols.** **A** daily protocol, **B** adaptive protocol with tumor size-driven treatment cycles, **C** daily protocol with toxicity-driven treatment cycles, and **D** adaptive protocol with both tumor size- and toxicity-driven treatment cycles.

before therapy resumes, while in other cycles the overshoot is much smaller. This irregularity is not seen in the protocols that include toxicity feedback. The explanation lies in the interaction between continuous tumor growth and the discrete, once-daily treatment decision schedule. In the adaptive protocol, the tumor may cross the  $Rx_{on}$  threshold at any time between daily checks; if this occurs early in the off-treatment interval, the tumor has nearly a full day to continue growing before therapy resumes, leading to a larger overshoot. If the threshold is crossed later in the interval, the overshoot is smaller. In contrast, when toxicity feedback is incorporated, therapy cannot restart immediately after tumor regrowth, and one must wait for toxicity to decline below  $Tox_{on}$ . This constraint effectively prolongs and regularizes the

off-treatment intervals. Similar effects of overshooting thresholds have been observed in recent work by Gallagher et al.<sup>32</sup>, where AT protocols with discrete monitoring intervals often allow the tumor to exceed the upper size threshold before treatment is resumed, due to the inherent delay between checks.

Incorporating toxicity-based dose adjustments into the daily protocol introduces periodic drug holidays (Fig. 2, third row). These treatment interruptions, designed to allow recovery from drug toxicity, also enable partial regrowth of sensitive cells, creating an unintended adaptative therapy protocol. Interestingly, for our baseline parameter values, this approach achieves a longer TTP (59.8 days) than even the adaptive protocol. This is

**Table 1 | Baseline parameter values used in Eqs. (1)–(4)**

Parameter	Description	Value	Dimensions
$\alpha_S$	Growth rate of sensitive cancer cells	1	per day
$\alpha_R$	Growth rate of resistant cancer cells	$\epsilon\alpha_S$	per day
$\epsilon$	Dimensionless constant between 0 and 1	0.4	dimensionless
$K$	Shared carrying capacity	100	dimensionless
$\delta$	Rate of drug-induced cancer cell death	1	per day
$\beta$	Competition parameter between sensitive and resistant cancer	2.4	dimensionless
$\lambda$	Drug clearance rate	0.693	per day
$\mu$	Rate of toxicity accumulation	1	per day
$\gamma$	Toxicity recovery rate	0.4	per day
$Rx_{off}$	Efficacy threshold triggers therapy pause if total tumor size falls below this value	40% of initial tumor size	dimensionless
$Rx_{on}$	Efficacy threshold triggers therapy restart if total tumor size surpasses this value	Initial tumor size	dimensionless
$Tox_{off}$	Toxicity threshold triggers therapy pause if toxicity surpasses this value	2	dimensionless
$Tox_{on}$	Toxicity threshold triggers therapy restart if toxicity falls below this value	1	dimensionless
Treatment Failure	Tumor size defining treatment failure and disease progression	75	dimensionless

because the toxicity feedback effectively mimics an adaptive strategy with higher  $Rx_{on}$  thresholds; that is, treatment resumes only after the tumor has grown beyond a threshold that would have triggered re-administration in the adaptive protocol.

The adaptive protocol with toxicity modulation integrates both tumor size-based and toxicity-based treatment decisions (Fig. 2, bottom row). This doubly adaptive approach results in sustained on/off treatment cycles. Like the daily protocol with toxicity feedback, tumor sizes overshoot the  $Rx_{on}$  threshold because treatment cannot resume until toxicity has declined sufficiently. This has the dual effect of reducing drug exposure, and consequently, toxicity, while preserving the sensitive cell population during early cycles. As a result, this protocol achieves the longest TTP (79.1 days).

### Sensitivity analysis

Next, we perform a global sensitivity analysis using eFAST to quantify the impact of parameter variations on TTP across the four treatment protocols. The results are summarized in Fig. 3. It should be noted that eFAST assigns a minimal sensitivity value to all parameters, including dummy parameters, making the  $p$ -values (indicated by asterisks in each plot) crucial for interpreting the significance of each sensitivity index.

The sensitivity indices plotted in Fig. 3 indicate that no single parameter strongly influences TTP in all four protocols. In the adaptive protocols (Fig. 3C, D), the competition parameter  $\beta$  most significantly influences TTP, as indicated by having the largest first order (blue bars), and higher order (orange bars), sensitivity indices. In the daily protocols (Fig. 3A, B), the parameter with the most influence TTP is the relative growth rate of resistant cells relative to sensitive cells ( $\epsilon$ ,  $\alpha_R = \epsilon\alpha_S$ ). Counterintuitively, the rate of toxicity recovery,  $\gamma$ , is either the least or second least sensitive parameter in any of the protocols. We explore this further in ‘Effect of Parameters on TTP’ subsection.

### Effect of Parameters on TTP

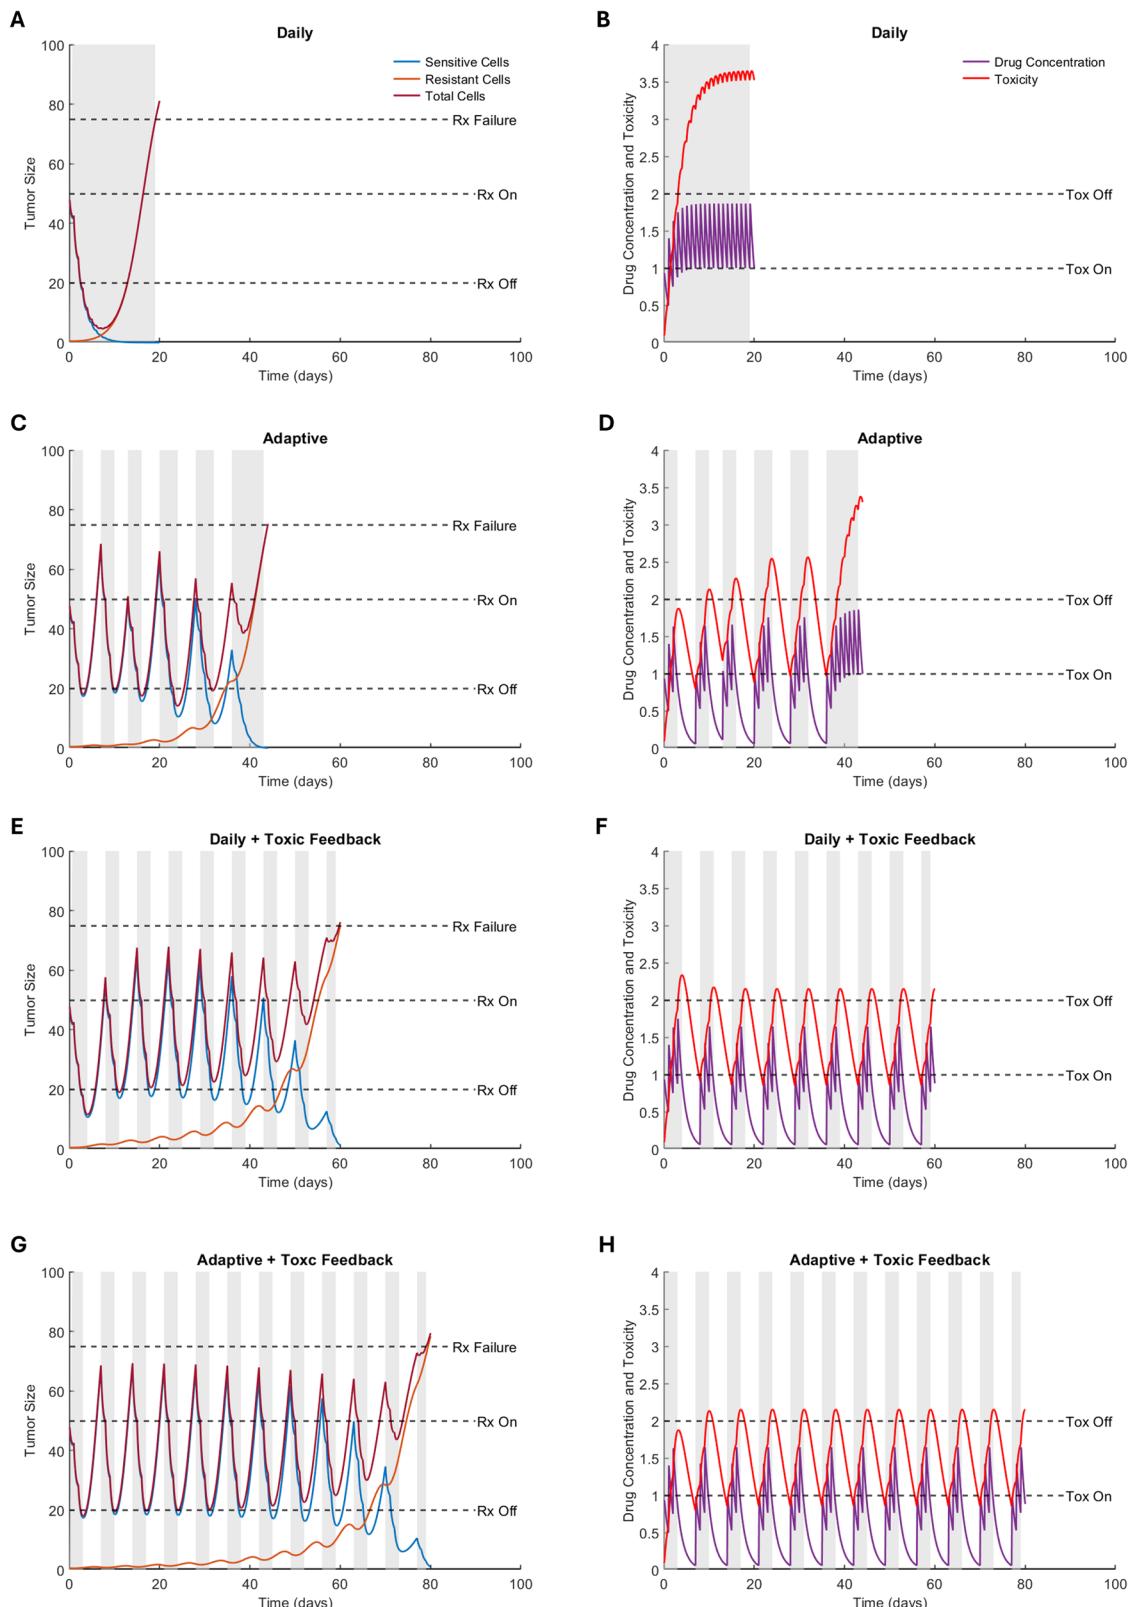
Parameters  $\beta$  and  $\epsilon$  (or  $\alpha_R$  as  $\alpha_R = \epsilon\alpha_S$ ) were identified as significant for most protocols in Fig. 3, so we start our analysis with them. We first investigate the effect on TTP of varying  $\beta$  and  $\alpha_R$  over the ranges defined in the global sensitivity analysis, with all other parameters fixed at their baseline values from Table 1. For each pair of values ( $\beta$ ,  $\alpha_R$ ) we compute the TTP for each protocol (Supplementary Fig. 2). The results of these four parameter sweeps are summarized in Fig. 4. In particular, Fig. 4A indicates which protocol(s) result in the longest TTP, and Fig. 4B indicates the actual time to progression for the protocol(s) that maximized TTP. Recall that the simulation is run for 100 days. We assign a value of 150 to any simulation that does not reach treatment failure within this simulation time.

As shown in Fig. 4A, one of the two adaptive protocols is overwhelmingly likely to maximize TTP when the resistant cells have a competitive advantage over the sensitive cells ( $\beta < 1$ ). The only exception to this is when resistant cells grow at nearly the same rate as the sensitive cells ( $\alpha_R$  near 1). In this scenario, the competitive advantage of the resistant cells is so extreme that none of the protocols can contain the tumor, and any stopping of treatment decreases TTP. Figure 4B shows that as the competitive advantage of the resistant cells increases (smaller  $\beta$ , larger  $\alpha_R$ ), the optimal protocol becomes less effective. As resistant cells lose their direct competitive advantage over sensitive cells (in particular for  $1 \leq \beta \leq 2$ ), the incorporation of toxicity feedback into the adaptive protocol yields the optimal result, though treatment still fails within the simulation period. Finally, when sensitive cells have a very strong competitive advantage over resistant cells ( $\beta \geq 2.75$ ) or a sufficiently strong competitive advantage ( $2.25 \leq \beta \leq 2.5$ ) coupled with significantly weaker growth of resistant cells relative to sensitive cells (small  $\epsilon$ ), TTP extends beyond the simulation window. As we move towards the parameter regime where the tumor does not progress (higher  $\beta$ ) from our baseline parameterization (red star), the adaptive protocol with toxicity modulation is the only protocol that appears as optimal (protocol 4 in Fig. 4A). As we move further to the right, we interestingly find that it is the daily protocol with toxicity modulation that can also extend TTP beyond the simulation period (protocol 3 in Fig. 4A). Finally, increasing  $\beta$  further, moving to the extreme end of parameter space for which sensitive cells have a competitive advantage over resistant cells, the adaptive protocol joins the two toxicity protocols as all three prevent progression within the simulation period.

### Rates of drug-induced death and of toxicity recovery affect the nature of treatment failure

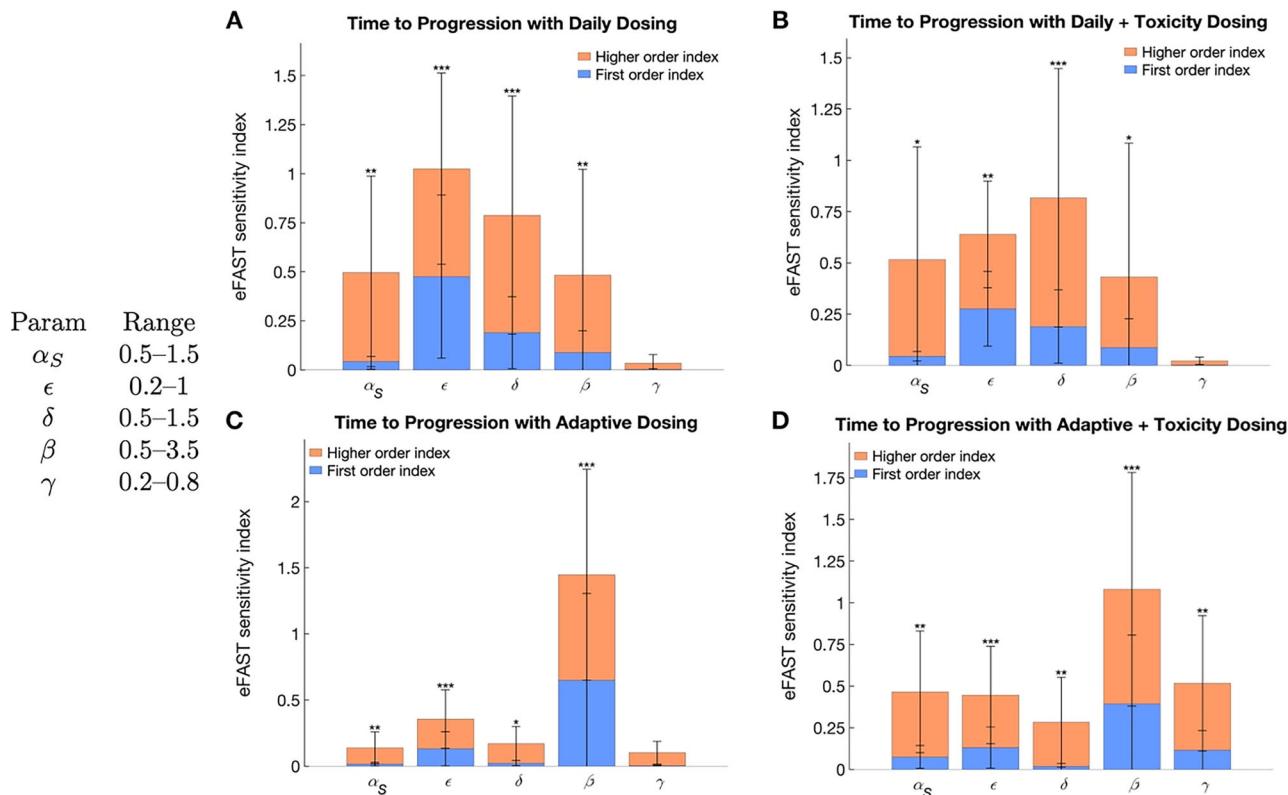
Next, we want to assess whether the *cause* of treatment failure is due to resistant or sensitive cells. Surprisingly, simulations revealed that while  $\alpha_R$  (varied by changing  $\epsilon$ ) strongly influences time to failure, it does not play an important role in determining the cause of failure (data not shown). This is because treatment failure by resistant cells or sensitive cells depends primarily on whether toxicity constrains the ability to control sensitive cells which is relatively independent of the intrinsic growth rate of the resistance cells.

To explore the underlying cause of progression, we performed a parameter sweep in  $(\beta, \delta)$ -space, where  $\delta$  is the drug-induced death rate and  $\beta$  is the competitive effect of sensitive cells on resistant cells. For each protocol and each  $(\beta, \delta)$  pair, keeping all other parameters fixed at their baseline values, we assess tumor composition (fraction of resistant cells) at the time of treatment failure. In the protocols without toxicity (Fig. 5A, B)

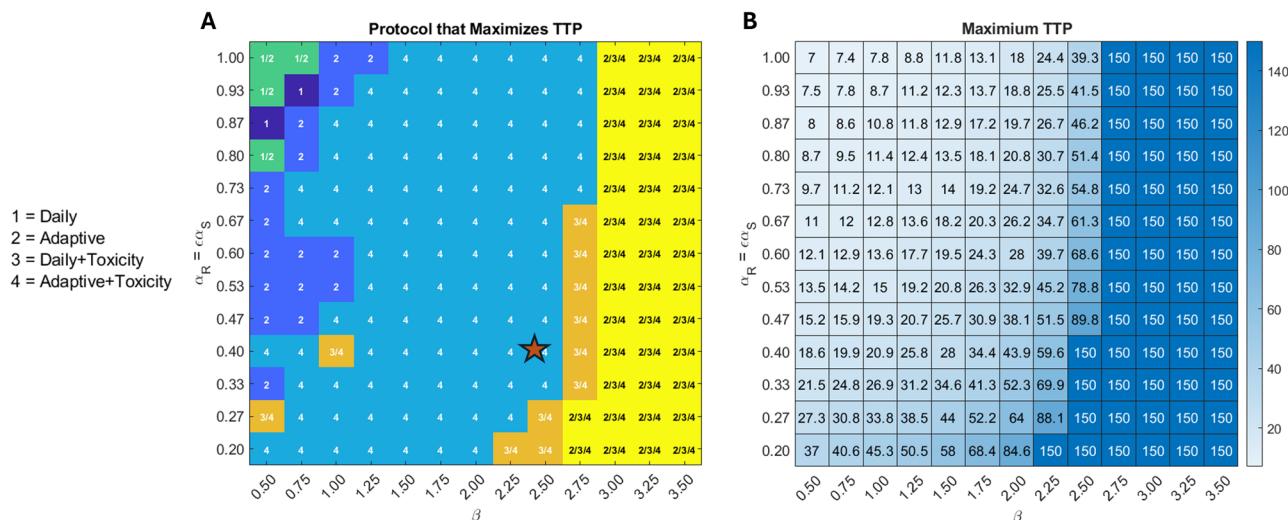


**Fig. 2 | The four treatment protocols and resulting model outputs using parameter values from Table 1.** The left column shows the time course of the total tumor size, along with the sensitive and resistant subpopulations. The right column shows the dynamics of drug concentration and accumulated toxicity. Compared protocols

are daily (A,B), adaptive (C,D), daily with toxicity feedback (E,F), and adaptive with toxicity feedback (G,H). Time periods shaded gray represent when the treatment is on.



**Fig. 3 | eFAST sensitivity indices as quantified by TTP.** Result for A daily protocol, B daily protocol with toxicity, C adaptive protocol, and D adaptive protocol with toxicity. Asterisks indicate statistical significance: \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Parameter ranges used are indicated in the graph. Note that  $\alpha_R = \epsilon \alpha_S$ .



**Fig. 4 | Parameter sweeps over  $\beta$ - $\alpha_R$  space.** A Protocol that maximizes time to progression for the corresponding pair of parameter values. B Maximum time to progression for the optimal protocol(s) identified in (A). A value of 150 indicates that the predicted tumor did not progress within the simulation time of 100 days and

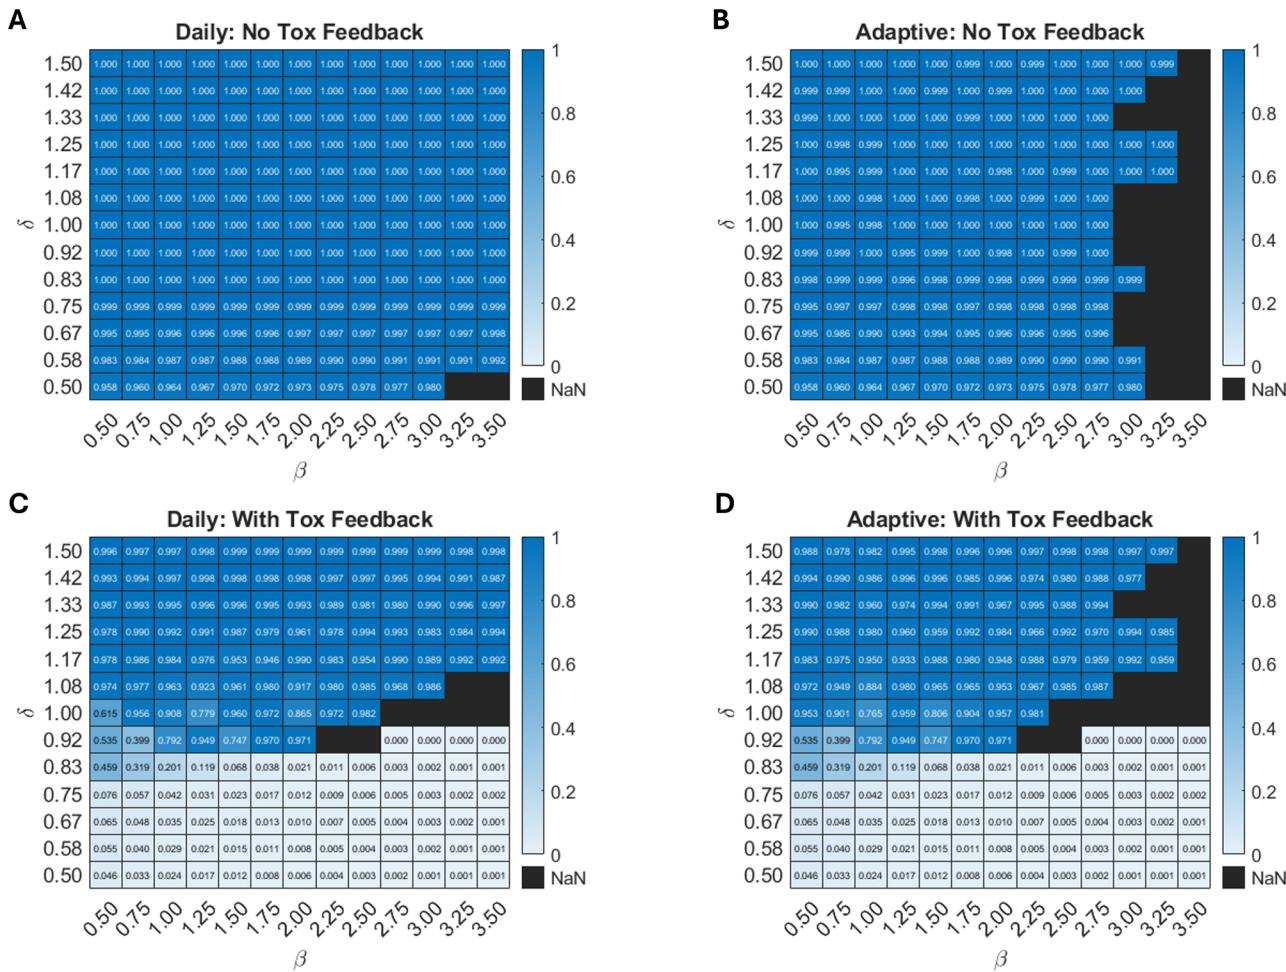
thus, that the treatment protocol did not (yet) fail. The baseline parameterization is indicated with a red star. Parameters that are not varied are fixed at their baseline value in Table 1 with  $\alpha_R$  varied from 0.2 to 1 by varying  $\epsilon$ .

treatment failure results from the eventual dominance of resistant cells, with sensitive cells nearly eradicated.

With toxicity, drug doses may be constrained to the point where sensitive cells drive progression regardless of treatment strategy (daily versus adaptive in Fig. 5C, D, respectively). When drug efficacy is low (small  $\delta$  values), toxicity prevents sufficient drug dosing to control the sensitive cells and they drive progression. The switch from resistant to sensitive cells

causing progression is relatively abrupt, occurring close to  $\delta = 1$ . Despite the importance of  $\beta$  in determining TTP (larger values prolong TTP), the role of  $\beta$  in the cause of treatment failure is minimal. However, the transition from resistant to sensitive cells causing progression as  $\delta$  declines becomes less abrupt at small values of  $\beta$ .

The global sensitivity analysis in Fig. 3 identified the toxicity recovery rate  $\gamma$  as either the least sensitive, or second least sensitive, parameter in all



**Fig. 5 | Fraction of resistant cells at TTP for different combinations of  $\beta$  and  $\delta$  tested over the four protocols.** Results for A daily, B adaptive, C daily with toxicity, and D adaptive with toxicity feedback. All other parameters are fixed at their baseline values reported in Table 1. Black regions indicate that treatment did not fail within

the simulation time of 100 days. As expected, progression when it occurs is always from resistant cells when there are no toxicity constraints. With toxicity constraints, as the rate of drug-induced cell mortality ( $\delta$ ) declines, the cause of progression switches from resistant to sensitive cells.

four protocols. This is surprising, as one would think that the rate of toxicity recovery would affect the outcomes of protocols with toxicity feedback. If the rate of toxicity recovery does not significantly influence TTP, perhaps it affects the cause of treatment failure? We explore this question next.

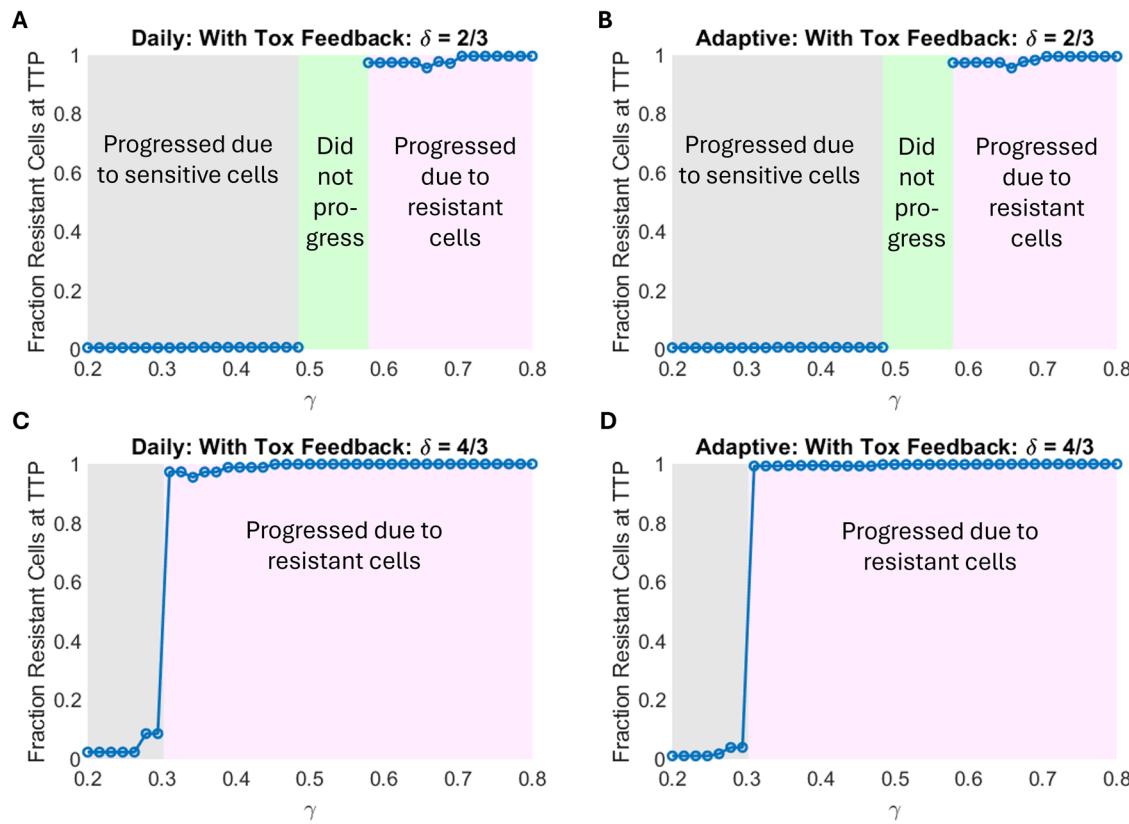
Figure 6 assesses the effect of  $\gamma$  in two different parameter regimes identified from Fig. 5C, D: when sensitive cells drive progression ( $\delta = 2/3$ ) and when resistant cells drive progression ( $\delta = 4/3$ ). In both parameter regimes, we observe a bifurcation in the reason for treatment failure. When  $\gamma$  is small toxicity accumulates, forcing treatment breaks that allow sensitive cells to regrow and outcompete resistant cells, ultimately driving treatment failure. Conversely, when  $\gamma$  is large, toxicity resolves quickly, permitting more frequent treatment cycles that, in turn, eliminate sensitive cells leaving the resistant cells to drive treatment failure. The bifurcation in behavior occurs at the same value of  $\gamma$  for each value of  $\delta$  considered (compare Fig. 6A, B with  $\delta = 2/3$ , then compare Fig. 6C, D with  $\delta = 4/3$ ).

Interestingly, there is a range of intermediate  $\gamma$  values in the  $\delta = 2/3$  case, where treatment does not fail within the simulation time frame. These cases result in prolonged on-off cycling of therapy. With this moderate drug efficacy, toxicity can be controlled and the resulting feedback-modified treatments work well. However, at higher values of treatment efficacy ( $\delta$ ), this ‘Goldilocks’ window vanishes, and treatment invariably fails during the simulation period. Prolonged cycling of treatment is no longer a possibility as resistant cells will lead to treatment failure. These results demonstrate that even though  $\gamma$  has a minimal effect on TTP (sensitivity analysis in Fig. 3), it can play a crucial role in determining the cause of treatment failure.

### Assessing therapy outcome robustness to protocol thresholds

Next, we assess whether treatment outcomes can be improved by adjusting the therapy on/off thresholds in adaptive and toxicity-tracking protocols. We previously showed that at baseline tumor size and toxicity thresholds for starting or stopping therapy, TTP varies significantly with  $\beta$  and  $\alpha_R$ . Here, we fix all model parameters to the baseline values from Table 1 and instead vary the tumor size thresholds  $Rx_{on}$  and  $Rx_{off}$  in the adaptive protocols, or the toxicity  $Tox_{on}$  and  $Tox_{off}$  in the protocols with toxicity feedback. We report predicted TTP for each of the following cases: the effect of varying  $Tox_{on}$  and  $Tox_{off}$  in daily and adaptive protocols with toxicity feedback (Fig. 7A, B, respectively) and the effect of varying  $Rx_{on}$  and  $Rx_{off}$  in adaptive protocols without or with toxicity feedback (Fig. 7C, D, respectively).

Figure 7A, B show TTP as a function of  $Tox_{on}$  (the toxicity level at which treatment may resume, horizontal axis) and  $Tox_{off}$  (the toxicity level at which therapy is paused, vertical axis). The constraint  $Tox_{on} < Tox_{off}$  must be satisfied; protocols violating this condition are blacked out in the bottom right corners. When  $Tox_{off}$  is too high, that is, when we are allowing a high level of toxicity in the simulation, toxicity-induced treatment pauses are infrequent. As a result, TTP in toxicity-feedback protocols is similar to that of protocols without toxicity feedback in both daily (TTP of 19.2 days) and adaptive (TTP of 44 days) dosing. Low values of  $Tox_{on}$ , meaning treatment can only resume after toxicity has mostly resolved, result in sharp declines in TTP. In these cases, if toxicity triggers a treatment pause (as seen with moderate  $Tox_{off}$  values in the daily protocol, and all values in the adaptive protocol), the toxicity level never falls low



**Fig. 6 | Effect of recovery from toxicity,  $\gamma$ , on the nature of treatment failure.** As  $\gamma$  increases, the reason for treatment failure switches from sensitive to resistant cells. **A** Daily with toxicity and **B** adaptive with toxicity protocols when drug effectiveness,

$\delta$ , is small ( $\delta = 2/3$ ). **C** Daily with toxicity and **D** adaptive with toxicity protocols when  $\delta$  is large ( $\delta = 4/3$ ). The remaining parameters are fixed at their baseline value in Table 1.

enough for the treatment to restart. The result is that cancer cells grow unchecked and TTP is very short. Conversely, high values of  $\text{Tox}_{\text{on}}$  allow for treatment to be resumed sooner, minimizing the effect of toxic feedback on the protocols, and resulting in similar TTPs to that of protocols without toxicity feedback.

Notably, the longest TTP for daily and adaptive protocols with toxicity feedback occurs at intermediate toxicity thresholds that are fairly close together, forcing toxicity to remain relatively constant and moderate. Within this sweet spot in protocol space, the model predicts that the tumor will not progress during the 100-day simulation window. These protocols with an optimal TTP occur at the boundary of the blacked-out region (Fig. 7A, B). Interestingly, this tumor control behavior occurs over a larger range of protocol thresholds for the adaptive with toxicity protocol, as compared to the daily with toxicity protocol.

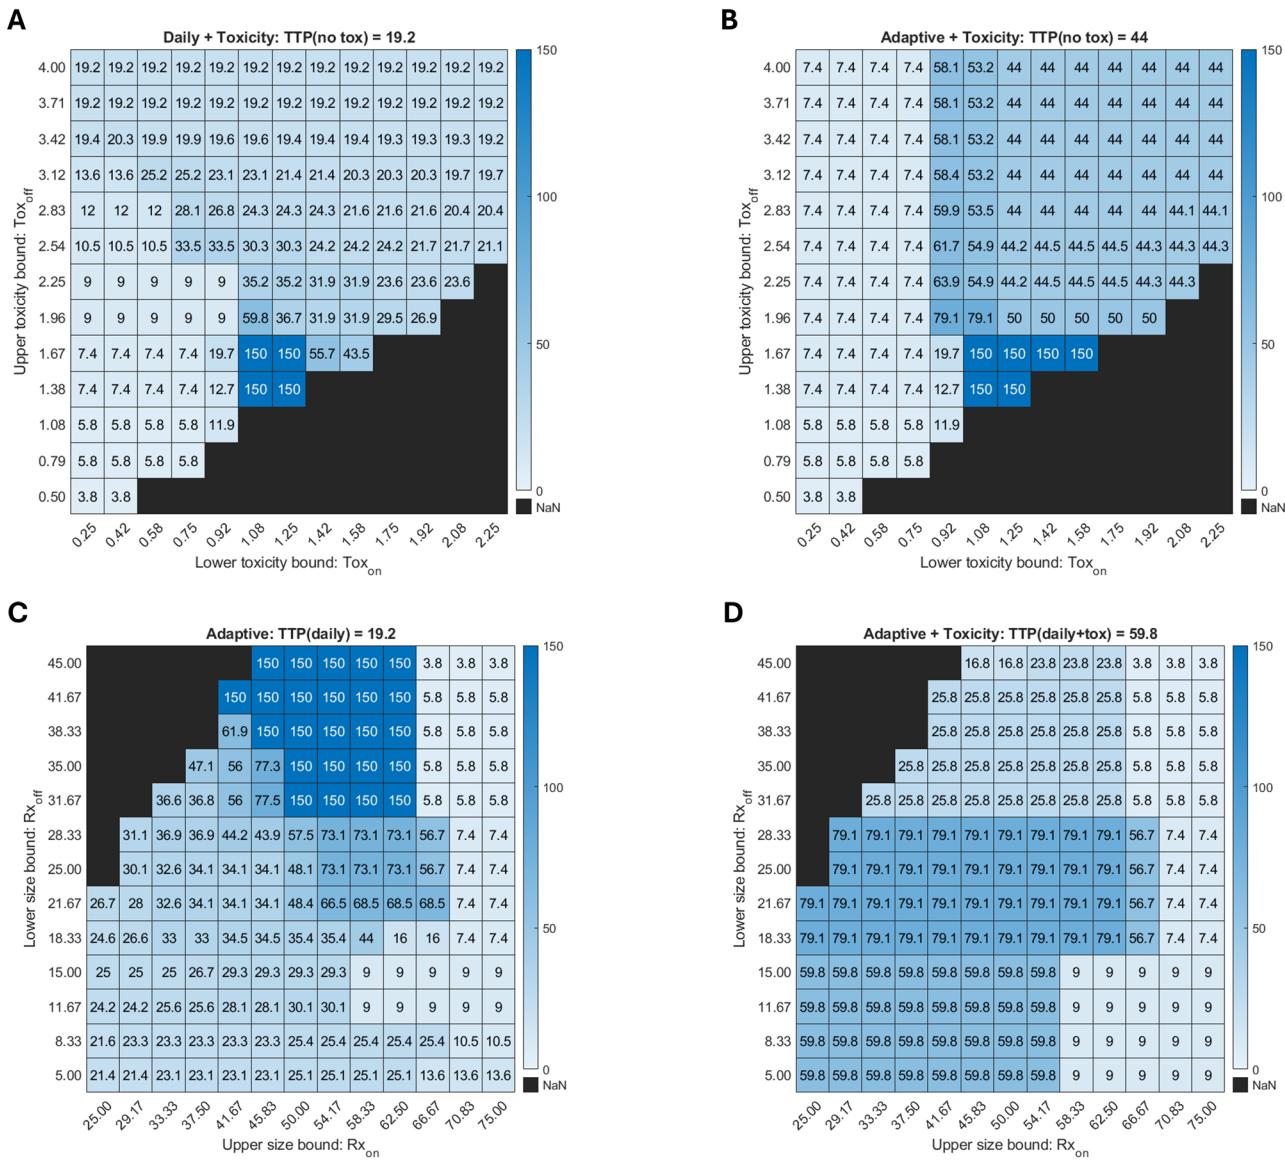
Figure 7C, D show TTP as a function of  $\text{Rx}_{\text{on}}$  (the tumor size at which treatment resumes, horizontal axis) and  $\text{Rx}_{\text{off}}$  (the tumor size at which treatment pauses, vertical axis). The constraint  $\text{Rx}_{\text{on}} > \text{Rx}_{\text{off}}$  must be satisfied; parameter sets violating this condition are blacked out in the top left corners. Looking at the rightmost columns in Fig. 7C, D, when  $\text{Rx}_{\text{on}}$  is too high, the tumor is allowed to grow to a sufficiently large size before a new treatment cycle begins. Under both adaptive protocols, TTP is noticeably lower than in the corresponding daily protocols (reported in the title of each subplot), regardless of  $\text{Rx}_{\text{off}}$ . When  $\text{Rx}_{\text{on}}$  is small, both adaptive protocols (with or without toxicity feedback) are comparable to, or modestly outperform, the corresponding daily protocol.

As we saw when varying the toxicity thresholds, there is a sweet spot in the adaptive protocol space for which the tumor is controlled over the 100-day window. This optimal window is found at high  $\text{Rx}_{\text{off}}$  and intermediate  $\text{Rx}_{\text{on}}$  values (indicated by the boxes labeled '150' in Fig. 7C). This region suggests that the best outcome is obtained by pausing treatment before the tumor shrinks too much, thereby

allowing the sensitive cell population to recover and suppress resistant cells through competition, and by restarting treatment before the tumor becomes too large.

Interestingly, we do not observe the same sweet spot in protocol space for the AT with toxicity feedback (Fig. 7D). At high  $\text{Rx}_{\text{off}}$  and intermediate  $\text{Rx}_{\text{on}}$  values, TTP is smaller than what we observe with the baseline threshold values (59.8 days). Instead, an optimal treatment window emerges at intermediate  $\text{Rx}_{\text{off}}$  values, where tumor burden is maintained between 36.6% and 56.6% of its initial size. Although adjusting size-based thresholds can improve TTP, none of the tested protocols produce a durable response – defined here as TTP exceeding the 100-day simulation period – once toxicity feedback is included in the adaptive protocol. While we did not identify tumor size thresholds for which the adaptive protocol with toxicity can control the tumor in Fig. 7D, Fig. 7B shows that toxicity feedback can produce the desired outcome. Both Fig. 7B, D are looking at two-dimensional slices of a four-dimensional protocol space, so it is important not to interpret this as concluding that the adaptive protocol with toxicity feedback is inferior to the others, as these are snapshots of a larger space.

The threshold sweeps in Fig. 7 also highlight the personalization level required to optimize the toxicity feedback protocol. We see that, depending on the chosen toxicity thresholds, the efficacy of a toxicity feedback protocol can improve or worsen relative to the daily protocol's TTP of 19.2 days. As an example, consider Fig. 7A and fix the maximum tolerated toxicity at  $\text{Tox}_{\text{off}} = 1.67$ . Moving the treatment restart threshold  $\text{Tox}_{\text{on}}$  from 0.92 to 1.08 results in a jump in TTP from 12.7 days (worse than the daily protocol) to over 100 days (much better than the daily protocol). However, a further increase of  $\text{Tox}_{\text{on}}$  reduces the predicted TTP, revealing the existence of a toxicity threshold "sweet spot" for this fixed parametrization. Such "sweet spots" exist for all the treatment cycling protocols, and the landscapes will depend on the model parameters.



**Fig. 7 | TTP as a function of protocol thresholds at baseline parameter values.** Top row: toxicity thresholds varied in the daily with toxicity (A) and adaptive with toxicity (B) protocols. Bottom row: tumor size thresholds varied in the adaptive (C) and adaptive with toxicity (D) protocols. Parameter sets for which the constraint

$Rx_{on} > Rx_{off}$  is not satisfied are blacked out. The results of the same protocol sweep for smaller and larger values of  $\beta$  than the baseline value of  $\beta = 2.4$  used here are shown in Supplementary Fig. 3.

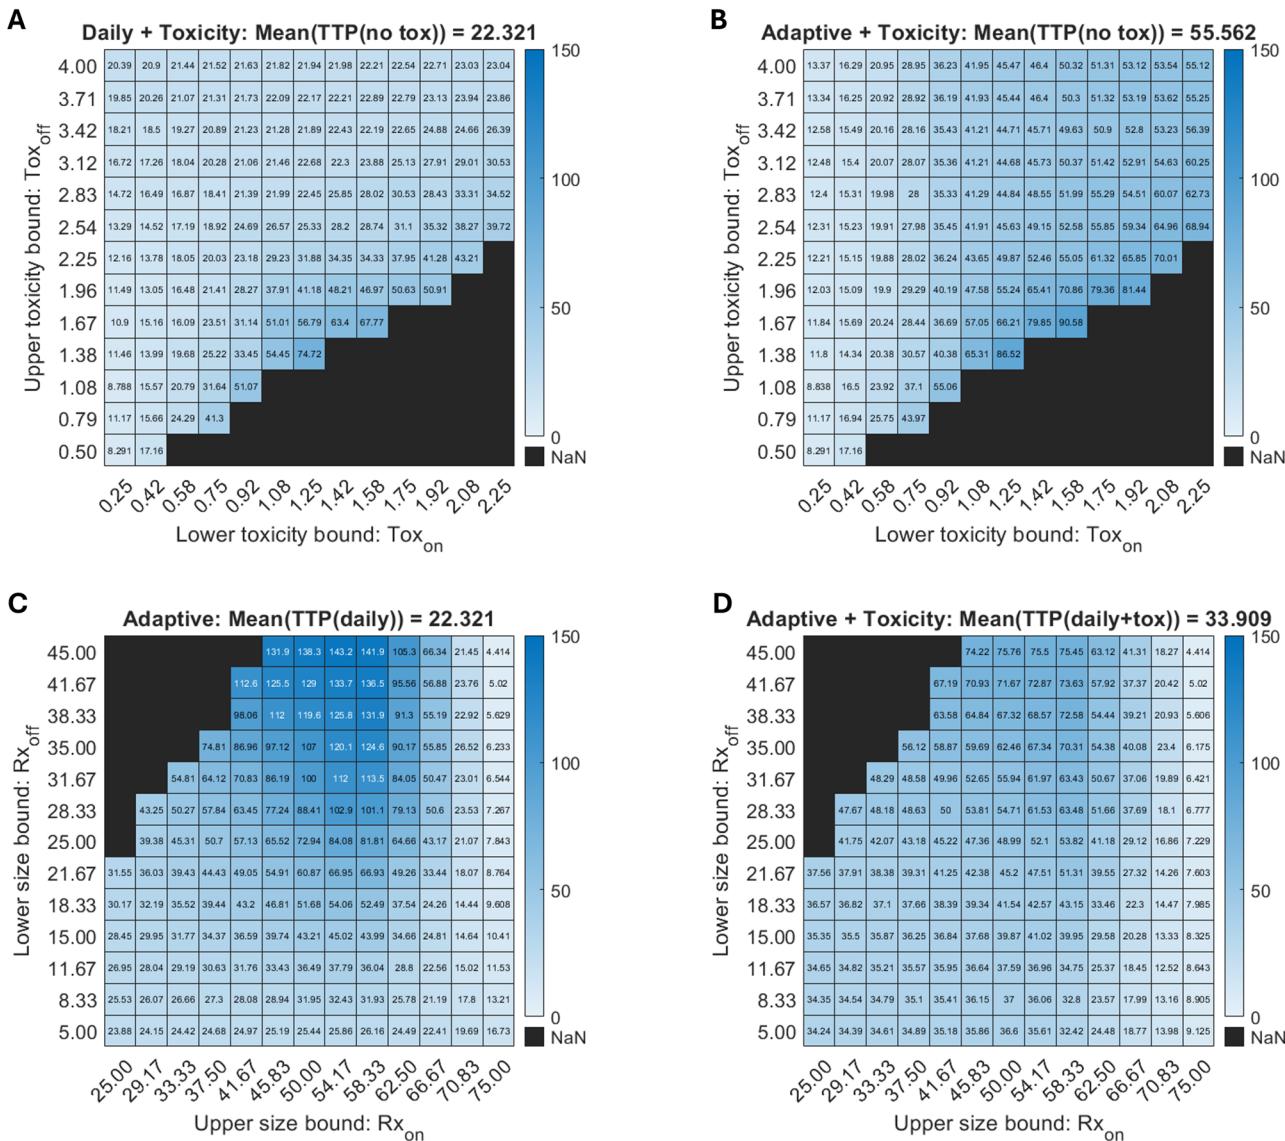
### Optimizing threshold “Sweet Spots” across a virtual population

Finally, to systematically determine whether adaptive and toxicity-based thresholds can be optimized for maximal TTP across various parameter values, we conduct a virtual patient analysis. In this approach, each “virtual patient” is assigned a unique set of values for five parameters,  $\alpha_s$ ,  $\epsilon$ ,  $\beta$ ,  $\delta$ , and  $\gamma$ , drawn from independent lognormal distributions (Supplementary Fig. 4). The means of these distributions match the corresponding baseline values in Table 1, and the standard deviations are chosen to approximate the parameter ranges used in the sensitivity analysis of Fig. 3. We then repeat the protocol threshold analysis from Fig. 7 for each virtual patient and calculate the mean TTP across the entire virtual population, for every pair of threshold values. As before, VPs that do not reach treatment failure in 100 days are assigned a TTP of 150. The results for 100 simulated VPs are reported in Fig. 8.

Figure 8A, B show mean TTP as a function of  $Tox_{on}$  (the toxicity level at which treatment resumes, horizontal axis) and  $Tox_{off}$  (the toxicity level at which treatment pauses, vertical axis), with the tested protocol specified in

the panel title, along with the corresponding mean TTP for the protocol without toxicity feedback. Again, the blacked out bottom right corners correspond to pairs of values that violate the constraint that  $Tox_{on} < Tox_{off}$ . Figure 8C, D show mean TTP as a function of  $Rx_{on}$  (the tumor size at which treatment resumes, horizontal axis) and  $Rx_{off}$  (the tumor size at which treatment pauses, vertical axis), with the tested protocol specified in the panel title, along with the corresponding mean TTP for the protocol without tumor size-based feedback. The upper left corner is blacked out as this region violates the constraint that  $Rx_{on} > Rx_{off}$ .

It is instructive to compare the heatmaps in Fig. 7 for the baseline parameterization to the heatmaps in Fig. 8, where we are looking at average behavior across a virtual population. When the toxicity thresholds are varied, we generally see that the sweet spot in protocol space is comparable between the baseline parameterization (Fig. 7A, B) and the virtual population (Fig. 8A, B). The same is true when the tumor size thresholds are varied in the adaptive protocol (Fig. 7C versus Fig. 8C). Interestingly, the analogous consistency is not observed in the adaptive protocol with toxic



**Fig. 8 | Mean TTP for 100 VPs.** Top row: toxicity thresholds varied in the daily with toxicity (A) and adaptive with toxicity (B) protocols. Bottom row: tumor size thresholds varied in the adaptive (C) and adaptive with toxicity (D) protocols.

feedback (Fig. 7D versus Fig. 8D). At the baseline parameter values, intermediate  $Rx_{on}$  values coupled with large  $Rx_{off}$  values corresponded to a suboptimal treatment protocol. However, this same region of protocol space is optimal for the average of the virtual population. This demonstrates that the effect of tumor size thresholds in the adaptive with toxicity protocol depends, to a large extent, on the model parameters.

To demonstrate the variability across VPs, as opposed to the average behavior as in Fig. 8, we examine the distribution of patient outcomes by plotting Kaplan-Meier curves (Fig. 9). These curves illustrate the proportion of virtual patients who have not progressed over time. The Kaplan-Meier curve for each protocol uses the thresholds that were identified to optimize the mean TTP for the virtual population in Fig. 8. For instance, the adaptive protocol in Fig. 8 has the longest TTP when  $Rx_{on} = 58.33$  and  $Rx_{off} = 45$ , so these are the thresholds set for the adaptive protocol in Fig. 9. For reference, we also include the Kaplan-Meier curve for the daily protocol with no treatment pausing (blue curve). The adaptive with toxicity protocol (purple curve) in Fig. 9 uses the optimal thresholds found by searching over the four-dimensional space defined by the two toxicity thresholds and the two tumor size thresholds.

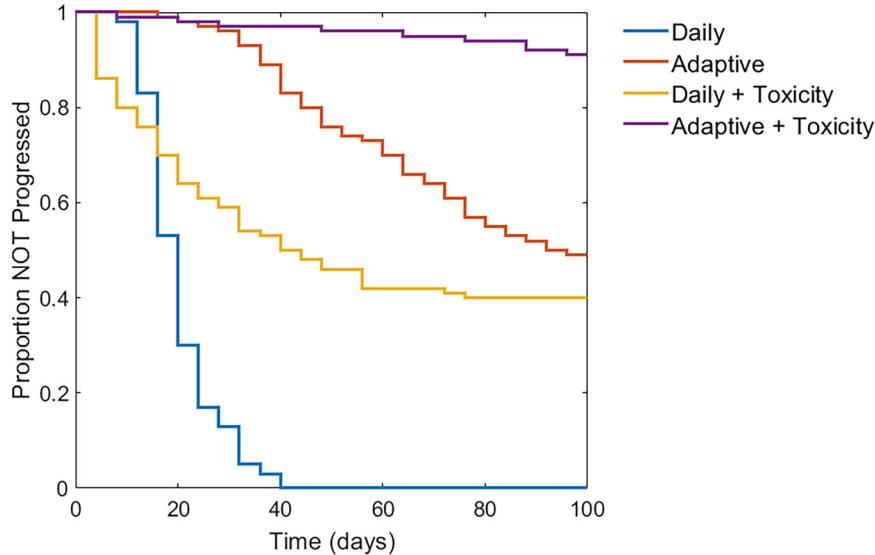
By adding toxicity considerations to the daily protocol (yellow curve in Fig. 9) 40% of VPs experience a TTP greater than 100 days, but 16%

experience a TTP shorter than the daily protocol. The adaptive protocol (red curve) significantly improves TTP for the majority of VPs compared to the daily protocol and leaves almost 50% of VPs with no disease progression by 100 days. Finally, the adaptive with toxicity protocol (purple curve), demonstrates the potential advantage of dual feedback: 91% experience no progression by 100 days.

## Discussion

Cancer treatments can fail for many reasons, including the emergence of resistance or unacceptable toxicity. In practice, toxicity often necessitates dose reductions, delays, or discontinuations, undermining the durability of systemic therapy. AT has been proposed as an alternative to conventional MTD approaches, with the aim of prolonging tumor control by maintaining a population of drug-sensitive cells that suppress resistant clones. In settings where cure is not possible due to the presence of resistant subpopulations, previous mathematical models and several clinical trials suggest that AT (e.g., pausing and restarting treatment based on tumor burden) can improve outcomes by delaying progression. However, most mathematical models of AT have focused on tumor burden alone and have not explicitly incorporated toxicity into treatment decisions. Our study addresses this gap by introducing a

**Fig. 9 | Variability across a virtual population for each protocol with optimized treatment thresholds.** Mean TTP of 100 VPs was used to determine the optimal thresholds for each protocol (except the daily protocol which does not have any thresholds).



conceptual framework in which toxicity constraints, alongside efficacy thresholds, could govern treatment administration.

We demonstrate that toxicity constraints can enhance or diminish outcomes depending on parameter values and protocol thresholds. In our analysis, we compared TTP of four protocols: daily; adaptive, where treatment is controlled based on tumor size criteria; daily with toxicity feedback, where treatment is controlled based on accumulated toxicity criteria; and adaptive with toxicity feedback, where treatment is controlled by both tumor size and toxicity criteria. We were able to identify scenarios in which each of the four protocols could optimize TTP. The optimal protocol depended on the growth rates of, and competition between, the sensitive and resistant cells, and on the defined protocol thresholds. Protocols leading to the longest TTPs typically incorporate either or both efficacy and toxicity feedback rules.

Additionally, by assessing tumor composition at TTP, we found that there are two paths to treatment failure. The tumor may progress due to significant growth of either sensitive or resistant cells, with progression due to sensitive cells occurring only in protocols with toxicity feedback. As such, drug dose timing becomes critically important: waiting too long to resume treatment (for example, because of slow toxicity recovery), may result in sensitive cell growth causing treatment failure.

Our results demonstrate that varying toxicity and efficacy thresholds can significantly affect treatment response, with relatively small changes in threshold values potentially having significant effects on TTP. We also observed that TTP is sensitive to several model parameters. For these reasons, we conducted a virtual population analysis to better understand the relationship between protocol thresholds and TTP across a heterogeneous population. In particular, we varied the protocol thresholds for both tumor size and/or toxicity (as relevant to the protocol), to determine the optimal thresholds leading to the longest average TTP for the virtual population. We found that while no protocol can prevent progression across the entire virtual population, there do exist protocols (threshold values) that can significantly extend TTP, on average, for all protocols with feedback. The longest average TTP, with the most non-progressing VPs, was predicted for the dual-feedback protocol incorporating both adaptive and toxicity constraints, as demonstrated by the simulated Kaplan-Meier survival curves in Fig. 9.

Although the proposed model is conceptual, employing a simplified representation of tumor heterogeneity (two subpopulations), drug pharmacokinetics, resistance (pre-existing only) and treatment-induced toxicity, the results are consistent with key principles of AT. Specifically, treatment pauses based on tumor size (though not toxicity) can actually maintain or even improve efficacy<sup>33</sup>, potentially through mitigating competition

between sensitive and resistant cancer cells. By pausing or reducing treatment upon reaching specific toxicity thresholds and resuming once resolved, patients may experience fewer adverse effects without compromising treatment outcomes. The importance of such an approach is further supported by the FDA's Project Optimus<sup>34</sup>, where drug developers are encouraged to move away from a "no regrets" MTD-like strategy towards a more nuanced, albeit harder to develop, protocol that balances efficacy and toxicity.

It should be emphasized that AT is primarily considered in the context of incurable disease where resistant subclones are present. If no resistant cells are present in the initial tumor, then an adaptive approach could theoretically reduce the possibility of cure by maintaining a residual population of sensitive cells rather than pursuing complete eradication. However, even in a setting where control and not cure is the goal, implementation of response or toxicity-guided treatment breaks may present communication challenges for both patients and clinicians<sup>35</sup>. Treatment interruptions may be interpreted as suboptimal treatment, which can make it difficult to convey that a pause is an intentional part of the therapeutic strategy. Misperceptions may also stem from a belief that more intensive or continuous therapy always yields better outcomes. For this reason, it is important that discussions with patients emphasize that pauses triggered by decreases in tumor size or by toxicity thresholds are intentional, designed to balance tolerability with maximizing long-term disease control, and should not be interpreted as therapeutic abandonment. To our knowledge, no studies have yet specifically explored patient perceptions of AT; understanding these views will be critical for clinical translation.

While beyond the scope of our present investigation, our model can be applied to a broader range of AT protocols. Like the first AT trial on castration resistant metastatic prostate cancer<sup>21</sup>, here we simulated therapy cycling based on fixed upper and lower thresholds of tumor burden, respectively. Mouse model experiments<sup>36–39</sup> and a current clinical trial on ovarian cancer<sup>40</sup> have used dose de-escalation as tumor burden drops and dose escalation as burden rises. Range-bounded AT lets the upper and lower bounds of on/off cycles drift upwards, taking advantage of stronger inter-cell type competition with increasing tumor burden<sup>41</sup>.

Future work could extend this model in several directions. First, threshold optimization is critical: the interplay between tumor size thresholds and toxicity thresholds is complex, and our results suggest the existence of "sweet spots" that maximize tumor control. Second, our analysis has focused on dose skipping but dose modulation is another strategy to manage toxicity. Extending our framework to incorporate both dose skipping and dose-modulation strategies will be an important next step. Third, the frequency of treatment decisions is likely to be critical: while we assumed daily

monitoring, less frequent assessments (weekly or monthly) may amplify overshoot and alter outcomes, as highlighted by Gallagher et al.<sup>32</sup>. Finally, integrating more realistic toxicity measures (e.g., neutropenia-based PK/PD models or composite toxicity indices) and incorporating patient-reported outcomes will be essential for translational relevance.

In conclusion, our modeling study demonstrates how incorporating toxicity constraints into AT can reduce treatment-associated toxicity without sacrificing efficacy. By balancing patient tolerability with treatment efficacy, such toxicity-informed adaptive protocols hold the promise of turning at least some types of incurable cancers into manageable chronic conditions. This approach may extend survival rates, improve patient quality of life, and potentially bring us closer to truly personalized cancer treatment.

## Methods

### Mathematical Model

The model analyzed in this study integrates drug pharmacokinetics and toxicity dynamics into an existing framework<sup>16</sup> of competitive interactions between treatment-sensitive ( $S(t)$ ) and treatment-resistant ( $R(t)$ ) cancer cells. The conceptual model assumes that the two subpopulations compete for, and grow logistically up to, a shared carrying capacity  $K$ . Sensitive cells can be killed by the drug, while resistant cells cannot. Drug concentration  $C(t)$  is described by a simple one compartment pharmacokinetic model with linear clearance. Toxicity  $T(t)$  increases with drug concentration  $C(t)$ , and resolves at a fixed linear rate.

The resulting system of equations is as follows:

$$S'(t) = \alpha_S S \left(1 - \frac{S+R}{K}\right) - \delta CS, \quad (1)$$

$$R'(t) = \alpha_R R \left(1 - \frac{\beta S+R}{K}\right), \quad (2)$$

$$C'(t) = \text{source} - \lambda C, \quad (3)$$

$$T'(t) = \mu C - \gamma T. \quad (4)$$

Here  $\alpha_S$  and  $\alpha_R$  are the growth rates of the sensitive and resistant cells, respectively, and  $\beta$  is the competition suppression factor that sensitive cells impose on resistant cells. The drug is assumed to be instantaneously injected into the model through the source term, has a clearance rate of  $\lambda$ , and a killing efficacy of  $\delta$  on the sensitive cells. Toxicity accumulates proportionally to the drug concentration with parameter  $\mu$ , and decays exponentially with parameter  $\gamma$ . Baseline values of parameters, together with their meanings, are summarized in Table 1. We remark that in our formulation, resistance to the drug incurs a cost so that resistant cells cannot proliferate faster than sensitive cells. Mathematically, we express the growth rate of resistant cells as:  $\alpha_R = \epsilon \alpha_S$ , where  $0 \leq \epsilon \leq 1$ .

In the interest of generality and wide applicability, and because here we are not modeling any specific cancer type, drug, or experimental dataset, we rescale our model variables as follows. The tumor cell subpopulations  $S(t)$  and  $R(t)$  are expressed as fractions of an assumed carrying capacity of 100, making them interpretable as percentages of sensitive and resistant tumor cells. The drug concentration  $C(t)$  is dimensionless, reflecting its role as a generic chemotherapeutic agent. Likewise, toxicity  $T(t)$  is measured in arbitrary units to encompass a broad range of potential side effects. These rescalings provide a flexible framework within which we can explore treatment dynamics across various scenarios. The baseline parameters were selected to demonstrate the dynamical differences between the four protocols detailed in 'Defining Treatment Protocols' subsection.

Equations (1)–(4) are solved in MATLAB® using ode23s, a stiff solver that implements a modified Rosenbrock formula of order two. Unless otherwise indicated, the scripts used to solve and analyze the toxicity model are available at <https://github.com/jgevertz/toxicity>.

### Defining treatment failure

Following the RECIST (Response Evaluation Criteria in Solid Tumors) criteria used in clinical practice, which typically defines disease progression as at least a 20% increase in tumor size from some baseline<sup>42</sup>, we define therapy failure as the total tumor burden rising 50% above its value at treatment initiation. In our model, the tumor burden starts at 50% of the carrying capacity, and treatment failure is thus defined as a tumor burden of 75% of the carrying capacity (that is, the initial tumor burden is 50, and treatment failure occurs when  $S+R \geq 75$ ). TTP is defined as the time from treatment initiation until the tumor burden reaches this treatment failure size. We emphasize that in a clinical setting, an oncologist would not know the tumor's carrying capacity and, instead, would make decisions based on measurable tumor burden obtained via imaging or other methodologies. To reflect this, efficacy-based decisions in our model are related to the tumor size at therapy initiation rather than to its (unknown) carrying capacity. In our numerical simulations, the tumor's initial composition is set at 99% sensitive cells and 1% resistant cells at 50% carrying capacity, or  $S(0) = 49.5$  and  $R(0) = 0.5$ . Treatment starts on day 1, so both the drug concentration and toxicity are assumed to be zero initially, that is,  $C(0) = T(0) = 0$ .

### Defining treatment protocols

Next, we present our approach to integrating both efficacy-based (tumor size) and toxicity-based (drug-induced toxicity level) feedback into the decision rules for starting, pausing, and resuming treatment. The resulting four protocols are summarized in Fig. 1.

**Daily protocol.** A fixed dose of 1 unit is administered at daily intervals (time  $t = 1$  day) until either treatment failure or the total elapsed time hits 100 days, whichever occurs first. Tumor burden is monitored daily, just before treatment administration, to check for disease progression. This protocol is summarized in Fig. 1A.

**Adaptive protocol.** Treatment is administered in daily 1-unit doses until the total tumor size,  $S(t) + R(t)$ , falls below the  $\text{Rx}_{\text{off}}$  threshold. At this point, therapy is paused until the tumor size regrows and again exceeds the upper  $\text{Rx}_{\text{on}}$  threshold, prompting treatment to resume. As with the daily protocol, tumor burden is monitored daily, just before potential treatment administration, to determine if disease progression has occurred. This protocol is summarized in Fig. 1B.

**Toxicity feedback.** This protocol can be incorporated into either of the previous two protocols described above, thus acting as an override switch for pausing therapy. Specifically, if toxicity levels are above a designated threshold  $\text{Tox}_{\text{off}}$  when the next treatment decision is to be made, then treatment is halted regardless of the tumor size. Therapy can only resume when the toxicity level at the time of the next treatment decision drops below the  $\text{Tox}_{\text{on}}$  threshold. In the daily protocol with toxicity feedback, dosing (re)starts when toxicity is below  $\text{Tox}_{\text{on}}$ , and will continue until toxicity rises above  $\text{Tox}_{\text{off}}$ . In the adaptive protocol with toxicity feedback, however, treatment (re)starts when both toxicity is below  $\text{Tox}_{\text{on}}$  and the tumor size is above  $\text{Rx}_{\text{on}}$ . Treatment will pause when either toxicity is above  $\text{Tox}_{\text{off}}$  or the total tumor size is below  $\text{Rx}_{\text{off}}$ . These conditions are tested daily, prior to dose administration. The cycles will continue until treatment failure or total simulation time hits 100 days. Figure 1C and D summarize adding toxicity feedback to the daily and adaptive protocols, respectively.

Each administered dose is assumed to be 1 unit of drug, and the potential treatment period is simulated over 100 days. During this time, the simulation may or may not predict disease progression by day 100. Time to progression (TTP) and tumor composition at TTP are recorded at the end of each simulation. If the tumor did not progress in the 100-day simulation, then TTP is recorded to be 150 (a randomly selected value greater than 100 to indicate treatment did not fail during the simulation time). Simulations will follow one of the four treatment protocols described above.

## Global sensitivity analysis framework

To assess the robustness of model predictions to variability in parameter values, we conduct a global sensitivity analysis using the extended Fourier Amplitude Sensitivity Test (eFAST). eFAST is a variance-based decomposition technique capable of efficiently handling nonlinear models<sup>43,44</sup>. The method quantifies the sensitivity of a model's output to variations in input parameters, by computing both a first-order and a total-order sensitivity index. The sensitivity of a particular input, the ratio of the total variability in the output attributed to changes in that input, is found by averaging over all other inputs. Parameters with high sensitivity indices are identified as influential to the model output, while those with low sensitivity indices may be regarded as negligible. The template used for eFAST implementation can be found here: <http://malthus.micro.med.umich.edu/lab/usadata/>

Herein, we assess the global sensitivity of TTP across the four treatment protocols for the following model parameters over the following ranges: the sensitive cell growth rate ( $0.5 \leq \alpha_S \leq 1.5$ ), the resistant cell growth rate relative to the sensitive cell growth rate ( $0.2 \leq \epsilon \leq 1$ ,  $\alpha_R = \epsilon \alpha_S$ ), the rate of drug-induced cancer cell death ( $0.5 \leq \delta \leq 1.5$ ), the competition parameter ( $0.5 \leq \beta \leq 3.5$ ), and the toxicity recovery rate ( $0.2 \leq \gamma \leq 0.8$ ). The output of this sensitivity analysis then informs a set of parameter sweeps, wherein we quantify the impact that simultaneously varying two model parameters has on TTP in each of the four treatment protocols.

We remark that the rate of toxicity accumulation  $\mu$  is omitted from the global sensitivity analysis in this proof-of-concept study. This is because we are not mapping toxicity to a specific biomarker or drug, so the scale of  $T(t)$  is arbitrary, and decisions to initiate or suspend therapy depend only on  $T(t)$  crossing the thresholds  $\text{Tox}_{\text{off}}$  and  $\text{Tox}_{\text{on}}$  (see Table 1). Moreover, the differential equation governing  $T(t)$  in Eqs. (1)–(4) can be integrated to yield the following analytical solution:

$$T(t) = \mu \int_0^t e^{-\gamma(t-s)} C(s) ds.$$

From this, it follows that  $T(t)$  and the toxicity thresholds  $\text{Tox}_{\text{off}}$  and  $\text{Tox}_{\text{on}}$  can be all rescaled by  $\mu$  without affecting model dynamics or decision-making criteria. Instead, we systematically vary the toxicity thresholds to elucidate their effect on treatment outcomes (see 'Assessing Therapy Outcome Robustness to Protocol Thresholds' subsection).

## Defining virtual patient framework

To extend our numerical simulations and two-dimensional parameter sweeps to better understand variability across parameter space, we utilize a virtual population approach<sup>45</sup>. In our computational framework, each virtual patient (VP) in the virtual population is represented by a set of parameter values related to characteristics of tumor growth and response to treatment. That is:

$$\text{Virtual Patient} = (\alpha_S, \alpha_R, \delta, \beta, \gamma).$$

The baseline parameter values listed in Table 1 can be thought of as corresponding to an 'average' patient. To construct a VP, a value for each parameter in the set  $(\alpha_S, \alpha_R, \delta, \beta, \gamma)$  is drawn as a simple random sample from a corresponding lognormal distribution. The distribution's shape parameters are chosen such that the peak approximately occurs at the baseline VP parameter value reported in Table 1, and the width approximately matches the range used in the global sensitivity analysis. This process is repeated 100 times to form a heterogeneous virtual population, or 100 sets of the five VP parameters.

The carrying capacity  $K$ , rate of drug clearance  $\lambda$ , and rate of toxicity build up  $\mu$  are fixed across all VPs. This decision was made to preserve the consistency of key modeling assumptions and to prevent confounding of sensitivity results due to parameter interdependencies. Since tumor size at treatment initiation was set at  $0.5K$  and treatment failure at  $0.75K$ , varying  $K$  would not affect time to progression. We fixed  $\mu$  because of its relation

with  $\gamma$  – higher toxicity build-up rates  $\mu$  can be offset by correspondingly higher toxicity clearance rates  $\gamma$ . Likewise, the drug's cell-killing effect can be preserved even with a high clearance rate  $\lambda$  if the rate of treatment-induced cell death  $\delta$  is also high.

## Data availability

Data sharing is not applicable to this article as no datasets were generated directly during the current study.

## Code availability

Programming scripts in MATLAB are freely available at <https://github.com/jgevertz/toxicity>.

Received: 1 August 2025; Accepted: 27 November 2025;

Published online: 08 January 2026

## References

1. Bright, C. J. et al. Data resource profile: the systemic anti-cancer therapy (SACT) dataset. *Int. J. Epidemiol.* **49**, 354 (2020).
2. Mahumud, R. A. et al. Systemic therapy and radiotherapy related complications and subsequent hospitalisation rates: a systematic review. *BMC Cancer* **24**, 826 (2024).
3. Goldstein, D. et al. Cancer-related fatigue in women with breast cancer: outcomes of a 5-year prospective cohort study. *J. Clin Oncol.* **30**, 1805–1812 (2012).
4. Jairam, V. et al. Treatment-related complications of systemic therapy and radiotherapy. *JAMA Oncol. Am. Med. Assoc.* **5**, 1028–1035 (2019).
5. Yazbeck, V. et al. An overview of chemotoxicity and radiation toxicity in cancer therapy. *Adv. Cancer Res.* **155**, 1–27 (2022).
6. Carlotto, A., Hogsett, V. L., Maiorini, E. M., Razulis, J. G. & Sonis, S. T. The economic burden of toxicities associated with cancer treatment: review of the literature and analysis of nausea and vomiting, diarrhoea, oral mucositis and fatigue. *Pharmacoeconomics. Springer* **31**, 753–766 (2013).
7. Wang, X. S. & Woodruff, J. F. Cancer-related and treatment-related fatigue. *Gynecol. Oncol.* **136**, 446–452 (2015).
8. Kuter, D. J. Treatment of chemotherapy-induced thrombocytopenia in patients with non-hematologic malignancies. *Haematologica* **107**, 1243 (2022).
9. Akbarali, H. I., Muchhal, K. H., Jessup, D. K. & Cheatham, S. Chemotherapy induced gastrointestinal toxicities. *Adv. Cancer Res.* **155**, 131–166 (2022).
10. Weinstein, Z. B. et al. Modeling the impact of drug interactions on therapeutic selectivity. *Nat. Commun.* **9**, 3452 (2018).
11. Tracey, A. Oncologists, advocates, FDA call for an end to MTD and the "more is better" era in cancer drug dosing. *Cancer Letter [Internet]* **48**, 1 [https://cancerletter.com/clinical/20220429\\_1/](https://cancerletter.com/clinical/20220429_1/) (2022).
12. West, J., Ma Y. & Newton, P. K. Capitalizing on competition: an evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment. *J. Theor. Biol.* **455**, 249–260 (2018).
13. West, J. et al. Towards multidrug adaptive therapy. *Cancer Res.* **80**, 1578–1589 (2020).
14. Bacevic, K. et al. Spatial competition constrains resistance to targeted cancer therapy. *Nat. Commun.* **8**, 1995 (2017).
15. Thomas, D. S., Cisneros, L. H., Anderson, A. R. & Maley, C. C. In silico investigations of multi-drug adaptive therapy protocols. *Cancers* **14**, 2699 (2022).
16. Strobl, M. A. others. Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Res.* **81**, 1135–1147 (2021).
17. Lenormand, T., Harmand, N. & Gallet, R. Cost of resistance: an unreasonably expensive concept. *Rethinking Ecol.* **3**, 51–70 (2018).
18. Kam, Y. et al. Sweat but no gain: inhibiting proliferation of multidrug resistant cancer cells with "ersatzdroges. *Int. J. Cancer* **136**, E188–E196 (2015).

19. Enriquez-Navas, P. M., Wojtkowiak, J. W. & Gatenby, R. A. Application of evolutionary principles to cancer therapy. *Cancer Res.* **75**, 4675–4680 (2015).

20. Kam, Y., Das T., Minton, S. & Gatenby, R. A. Evolutionary strategy for systemic therapy of metastatic breast cancer: balancing response with suppression of resistance. *Women's Health* **10**, 423–430 (2014).

21. Zhang, J., Cunningham, J. J., Brown, J. S. & Gatenby, R. A. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nat. Commun.* **8**, 1816 (2017).

22. Kozlowska, E., Haltia, U.-M., Puszynski, K. & Färkkilä, A. Mathematical modeling framework enhances clinical trial design for maintenance treatment in oncology. *Sci. Rep.* **14**, 29721 (2024).

23. Shochat, E. & Rom-Kedar, V. Novel strategies for granulocyte colony-stimulating factor treatment of severe prolonged neutropenia suggested by mathematical modeling. *Clin. Cancer Res. Am. Assoc. Cancer Res.* **14**, 6354–6363 (2008).

24. Cassidy, T., Humphries, A. R., Craig, M. & Mackey, M. C. Characterizing chemotherapy-induced neutropenia and monocytopenia through mathematical modelling. *Bull. Math. Biol.* **82**, 1–26 (2020).

25. Fathoni, M. I. A., Adi-Kusumo, F., Gunardi, G. & Hutajulu, S. H. Dynamics of a breast cancer model for neutropenia case due to chemotherapy effects. *Int. J. Differ. Eqns.* **2021**, 3401639 (2021).

26. Craig, M. Towards quantitative systems pharmacology models of chemotherapy-induced neutropenia. *CPT Pharmacometr. Syst. Pharmacol.* **6**, 293–304, (2017).

27. Farhang-Sardroodi, S., La Croix, M. A. & Wilkie, K. P. Chemotherapy-induced cachexia and model-informed dosing to preserve lean mass in cancer treatment. *PLoS Comput. Biol.* **18**, e1009505 (2022).

28. Hadjandreou, M. M. & Mitsis, G. D. Mathematical modeling of tumor growth, drug-resistance, toxicity, and optimal therapy design. *IEEE Trans. Biomed. Eng.* **61**, 415–425 (2013).

29. Dua, P., Dua, V. & Pistikopoulos, E. N. Optimal delivery of chemotherapeutic agents in cancer. *Comp. Chem. Eng.* **32**, 99–107 (2008).

30. Barbolosi, D. & Iliadis, A. Optimizing drug regimens in cancer chemotherapy: a simulation study using a PK–PD model. *Comp. Biol. Med.* **31**, 157–172 (2001).

31. McGehee, C. & Mori, Y. A mathematical framework for comparison of intermittent versus continuous adaptive chemotherapy dosing in cancer. *npj Syst. Biol. Appl.* **10**, 140 (2024).

32. Gallagher, K., Strobl, M. A., Anderson, A. R. & Maini, P. K. Deriving optimal treatment timing for adaptive therapy: matching the model to the tumor dynamics. *medRxiv* **87**, 146 (2025).

33. Zhang, J., Cunningham, J., Brown, J. & Gatenby, R. Evolution-based mathematical models significantly prolong response to abiraterone in metastatic castrate-resistant prostate cancer and identify strategies to further improve outcomes. *Elife* **11**, e76284 (2022).

34. Gao, W. et al. Realizing the promise of project optimus: challenges and emerging opportunities for dose optimization in oncology drug development. *CPT Pharmacometr. Syst. Pharmacol.* **13**, 691–709 (2023).

35. Patell, R. et al. Patient perceptions of treatment benefit and toxicity in advanced cancer: a prospective cross-sectional study. *JCO Oncol. Pract.* **17**, e119–e129 (2021).

36. Enriquez-Navas, P. M. et al. Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Am. Assoc. Adv. Sci.* **8**, 327 (2016).

37. Wang, J., Zhang, Y., Liu, X. & Liu, H. Optimizing adaptive therapy based on the reachability to tumor resistant subpopulation. *Cancers* **13**, 5262 (2021).

38. Strobl, M. A. et al. To modulate or to skip: De-escalating PARP inhibitor maintenance therapy in ovarian cancer using adaptive therapy. *Cell Syst.* **15**, 510–525 (2024).

39. Hockings, H. et al. Adaptive therapy exploits fitness deficits in chemotherapy-resistant ovarian cancer to achieve long-term tumor control. *Cancer Res.* **85**, 3503–3517 (2025).

40. Mukherjee, U. A. et al. Study protocol for Adaptive ChemoTherapy for Ovarian cancer (ACTOv): a multicentre phase II randomised controlled trial to evaluate the efficacy of adaptive therapy (AT) with carboplatin, based on changes in CA125, in patients with relapsed platinum-sensitive high-grade serous or high-grade endometrioid ovarian cancer. *BMJ Open* **14**, e091262 (2024).

41. Brady-Nicholls, R. & Enderling, H. Range-bounded adaptive therapy in metastatic prostate cancer. *Cancers MDPI* **14**, 5319 (2022).

42. Villaruz, L. C. & Socinski, M. A. The clinical viewpoint: definitions, limitations of RECIST, practical considerations of measurement. *Clinical cancer research. Am. Assoc. Cancer Res.* **19**, 2629–2636 (2013).

43. Saltelli, A., Tarantola, S. & Chan, K.-S. A quantitative model-independent method for global sensitivity analysis of model output. *Technometrics. Taylor Francis* **41**, 39–56 (1999).

44. Marino, S., Hogue, I. B., Ray, C. J. & Kirschner, D. E. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *J. Theoret. Biol.* **254**, 178–196 (2008).

45. Craig, M., Gevertz, J. L., Kareva, I. & Wilkie, K. P. A practical guide for the generation of model-based virtual clinical trials. *Front. Syst. Biol.* **3**, 1174647 (2023).

## Acknowledgements

We gratefully acknowledge that this research was supported by the Fields Institute for Research in Mathematical Sciences (Toronto, Canada) through a workshop entitled: Mathematical Modeling of Cancer Treatments, Resistance, Optimization. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the Institute. The authors thank Dr. Helen Byrne for insightful discussions during the workshop. This research was supported in part by the Natural Sciences and Engineering Research Council (NSERC) Discovery Grant program RGPIN-2018-04205 (KWP).

## Author contributions

M.D. proposed the research question. J.L.G., H.V.J., I.K., K.P.W., J.B., Y.P.H., E.S., V.V., and M.D. were involved with the initial project formulation. J.L.G., H.V.J., I.K., and K.P.W. developed the algorithms, wrote the code, ran numerical simulations, designed figures, and prepared and edited all manuscript drafts. J.L.G., H.V.J., I.K., K.P.W., J.B., Y.P.H., E.S., V.V., and M.D. have read and approved the manuscript.

## Competing interests

Author HVJ serves as Associate Editor of this journal and had no role in the peer-review process or decision to publish this manuscript. JLG, IK, KPW, JB, YPH, ES, VV and MD declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41540-025-00635-6>.

**Correspondence** and requests for materials should be addressed to Jana L. Gevertz or Mark Davies.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026