



When do drug holidays delay resistance? A calibrated model of BRAF V600E melanoma

Wang E.Y

Submitted: March 20, 2026, Revised: version 1, June 14, 2026, version 2, July 20, 2026, version 3, August 30, 2026, version 4, September 14, 2026, version 5, September 18, 2026, version 6, September 24, 2026

Accepted: September 24, 2026

Abstract

Treatment scheduling seeks to kill sensitive tumor cells while limiting resistance, but drug holidays can suppress drug-dependent resistant populations while allowing sensitive cells to regrow and subsequently replenish resistance through the persister state. We developed a calibrated mechanistic model of BRAF V600E melanoma to identify the biological conditions under which this trade-off favors intermittent treatment. The model links MAPK signaling and protein degradation to sensitive, persister, and resistant cell populations. Across 1,200 virtual tumors, scheduling benefit depended on resistant-cell growth relative to turnover and on loss of resistant-cell fitness after treatment withdrawal. Comparisons at equal cumulative exposure showed that timing itself altered resistance outcomes, although shorter cycles were not consistently better. Intermittent treatment delayed resistant disease more often than it improved overall progression, demonstrating that resistance delay alone is insufficient evidence of better tumor control. The study therefore provides a conditional map of scheduling benefit and a reusable framework for investigating the temporal responses of sensitive and resistant populations. Its population parameters remain experimentally underconstrained, and its representation of average cell behavior does not reproduce rare-cell persistence. With tumor-specific calibration and independent validation, this framework could guide the development of standardized models that balance sensitive-cell killing against resistance control.

Keywords

BRAF V600E melanoma, Treatment scheduling, Intermittent therapy, Drug holidays, Acquired drug resistance, Drug-tolerant persister cells, MAPK signaling, Targeted protein degradation, Mathematical oncology, Mechanistic modeling

Ethan Y Wang, William A. Shine Great Neck South High School, 341 Lakeville Road, Great Neck, NY 11020, USA. ethanwang202511@gmail.com

1. Introduction

1.1 Targeted therapy and acquired resistance in BRAF V600E melanoma

Melanoma arises from melanocytes, the pigment-producing cells of the skin. Although it accounts for a small proportion of skin cancer diagnoses, it causes a disproportionately large share of skin cancer deaths because of its capacity for rapid growth and spread (1). Roughly half of cutaneous melanomas carry an activating mutation in the BRAF gene, most often the V600E substitution, which maintains activation of the mitogen-activated protein kinase (MAPK) pathway and drives tumor-cell proliferation (1).

Drugs targeting this specific dependency initially produce substantial responses. Flaherty et al. reported response rates near 80% in the phase I extension of vemurafenib and found that responding tumors showed the deepest suppression of ERK phosphorylation (2). Adding a MEK inhibitor improved outcomes further: Larkin et al. showed that combining cobimetinib with vemurafenib extended median progression-free survival from 6.2 to 9.9 months in a randomized trial of 495 patients (3).

These responses, however, are often not sustained. Within months, many tumors regain MAPK signaling. Villanueva et al. showed that resistant melanoma samples restore signaling through several routes, including increased receptor tyrosine kinase expression, increased RAS activity, or BRAF amplification (4). Inhibiting kinase activity does not remove the

protein, leaving cells with alternative routes through which signaling can persist or recover.

1.2 Drug-tolerant persistence and the emergence of resistance

Before stable resistance emerges, a small group of cells can survive treatment without acquiring genetic changes. Sharma et al. described these cells in lung cancer, finding that approximately 0.3% of a cell population survived a hundredfold excess of drug in a reversible, slowly dividing state termed drug-tolerant persistence. Approximately one-fifth of these cells eventually resumed proliferation while the drug remained present (5). In melanoma, Shaffer et al. used single-cell sequencing and live-cell barcoding to show that approximately one cell in 2,000 continued to divide under vemurafenib and eventually formed a resistant colony. These cells could be identified by their gene expression before treatment began (6).

Yang et al. observed the process directly. Using a biosensor for CDK2 activity, they tracked thousands of individual melanoma cells and found that 43% of A375 cells entered a quiescent state under dabrafenib, subsequently restored CDK2 activity, reentered the cell cycle, and divided. Increasing the dose tenfold reduced this fraction to 10% without eliminating it (7). Gerosa et al. characterized the associated signaling, reporting that vemurafenib reduced average pathway activity to approximately 5% of baseline within minutes. Nevertheless, 1–5% of cells retained ERK activity comparable to untreated cells at any given moment, with brief signaling pulses appearing in small patches (8).

Rambow et al. showed that residual disease during BRAF and MEK inhibitor treatment comprises transcriptionally distinct cell states rather than simply fewer of the original tumor cells (9). Marine et al. have argued more broadly that non-genetic resistance requires countermeasures distinct from those targeting mutational resistance (10). Together, these studies describe drug-tolerant populations that survive on residual signaling and can provide a source from which stable resistance emerges. Hata et al. demonstrated two routes to resistance in lung cancer by tracking resistant clones: some arose from rare cells already resistant before treatment, whereas others acquired resistance mutations within drug-tolerant cells during prolonged exposure (11).

1.3 Conflicting evidence on intermittent treatment

An alternative to improving the drugs themselves is to change how they are administered. The rationale comes from Das Thakur et al., who observed that melanoma cells resistant to vemurafenib grew poorly after drug withdrawal. Their resistance depended on increased production of BRAF V600E, which became a liability when the drug was absent (12). They tested this strategy in mice carrying patient-derived tumors, administering vemurafenib continuously to one group and the same dose on a four-week-on, two-week-off cycle to another. Every continuously treated mouse developed lethal drug-resistant disease within 100 days, whereas none receiving intermittent treatment developed resistant disease during 200 days of observation (12).

These findings motivated a randomized clinical trial. In SWOG S1320, Algazi et al. randomized 206 patients with BRAF V600 melanoma who were already benefiting from dabrafenib and trametinib to either continuous treatment or a five-week-on, three-week-off cycle (13). Contrary to the hypothesis, intermittent dosing produced shorter median progression-free survival than continuous dosing: 5.5 versus 9.0 months.

Intermittent treatment therefore produced contrasting outcomes in these experimental and clinical settings, which differed in treatment regimen, schedule, and endpoint. Understanding this discrepancy matters beyond the particular schedules tested, because it highlights the need to identify the biological conditions under which a drug holiday may help or harm a given tumor.

1.4 Mathematical models of treatment scheduling

The idea that treatment timing can be chosen to manage resistance rather than simply maximize cell killing was developed by Gatenby et al., who argued that resistant cells are often rare before treatment because resistance carries a fitness cost. Intensive treatment removes the sensitive cells that would otherwise compete with them (14). Zhang et al. subsequently translated this approach into a clinical protocol for prostate cancer, using each patient's tumor marker to determine when to pause and restart abiraterone. They reported that ten of eleven patients maintained stable disease, with a median time to progression exceeding 27 months while receiving less than half the standard cumulative dose (15).

Research in this area has also examined the biological conditions under which these strategies succeed. Strobl et al. found that cell turnover is a key determinant of whether adaptive therapy can control a tumor and, specifically, how strongly the fitness cost of resistance contributes to control (16). Their approach—varying a small number of biologically meaningful parameters across plausible ranges to identify where the strategy works—is the approach we follow here.

These models, however, do not explicitly represent the signaling biology that determines the extent of drug-induced pathway suppression. They describe cell populations growing and competing without modeling intracellular drug action. Although appropriate for their purposes, this leaves treatment scheduling in BRAF-mutant melanoma partly unresolved, because the persister state from which resistance can emerge depends on signaling dynamics.

1.5 Biological requirements for modeling BRAF V600E melanoma

Two structural choices determine whether the model can represent the pathways through which treatment fails: the topology of RAF activation and the response of the signaling cascade to inhibition.

The first concerns RAF activation. In cells with normal BRAF, RAF activity depends on signaling through receptor tyrosine kinases and RAS. Applying this arrangement to BRAF V600E would reproduce wild-type rather than mutant biology. Yao et al. showed that V600 mutants signal as active monomers—individual

protein molecules acting independently—even when RAS activity is low, and that ERK feedback suppresses RAS activity in these tumors (17). Making all RAF activation dependent on RAS would therefore misrepresent the signal that treatment must suppress. A separate dimer route, in which two RAF protein molecules associate as a pair, is also needed to represent signaling that persists despite monomer inhibition. Poulikakos et al. showed that inhibitor binding to one member of a RAF pair can activate the other, helping explain why dimers respond poorly to compounds that effectively inhibit individual monomers (18). Without a separate dimer route, the model cannot represent this source of residual signaling.

The second concerns the response of the RAF–MEK–ERK cascade. Representing every step as a simple first-order transfer can fail to capture the strong suppression observed experimentally. Huang and Ferrell found that the cascade exhibits an ultrasensitive response, with an effective Hill coefficient of four to five, even though its individual components are not themselves cooperative (19). This behavior allows partial inhibition upstream to produce a much larger reduction in downstream activity. The depth of suppression matters because Bollag et al. reported greater than 80% inhibition of ERK phosphorylation in biopsies from patients whose tumors regressed (20). A model that cannot suppress ERK below ~ one-quarter of baseline does not reach this clinically observed response range.

Inadequate representation of these features can distort the balance between tumor-cell killing,

persistence, and resistance. In earlier model formulations, this produced near-complete eradication or outright failure, with few intermediate outcomes. Because persistence and subsequent resistance motivate the scheduling question, the model must represent these intermediate states before schedule comparisons can be interpreted meaningfully.

1.6 Study objectives

We built a model that separates the monomer and dimer routes of BRAF V600E signaling, distinguishes protein degradation from kinase inhibition, and includes an explicit resistant compartment that cells can enter from the persister state. Rather than searching for a single best schedule, we mapped the boundary between tumors that remain controlled and those that progress, identified the biological parameters determining their position relative to that boundary, and examined whether changing the treatment schedule shifts the boundary and which tumors benefit.

2. Materials and Methods

2.1 Model structure and implementation

The model is a system of ordinary differential equations that simultaneously tracks the state of the MAPK signaling network within a representative cell, the concentration of each drug at the cell, and the sizes of three cell populations (Figure 1). Signaling variables are expressed as the fraction of each protein pool in its active form, and drug concentrations are expressed in nanomolar. Cell populations are

counted as absolute numbers rather than fractions, for the reasons given in Section 2.6. Time is measured throughout in hours.

The complete equations, initial conditions, equilibration procedure, integration and event rules, schedule functions, population transition rules, and endpoint definitions are provided in the Supplementary Methods. Supplementary Table 1 lists every parameter, its units, its role, and a one-line definition. The code that produces every number and figure in this paper is available in the project repository. Each run records the version of every library used, along with a hash of the source files, to support exact reproduction of the results.

2.2 MAPK signaling and feedback

The central structural decision concerns how the mutant protein drives the pathway. We follow the arrangement described by Yao et al. and Lito et al. rather than the conventional wild-type arrangement (17, 21).

BRAF V600E activates MEK through two routes that differ in their sensitivity to the drug. The first involves the mutant protein acting as a monomer, which does not require active RAS and is effectively blocked by a RAF inhibitor. The second involves dimers, whose formation requires GTP-loaded RAS and which the same inhibitor blocks much less effectively. Yao et al. found that inhibiting dimer-driven ERK activity required more than 30 times the concentration needed to inhibit monomer-driven activity (17).




Figure 1. Diagram of the model. The left panel shows the three cell populations, which are a bulk population, a drug tolerant persister population that cells enter under treatment and can leave again when signalling recovers, and a resistant population that cells enter from the persister state and cannot leave. There is no arrow into the resistant population from anywhere else: the equations carry no preexisting influx term and every simulation reported here starts with no resistant cells. Conversion into the persister state speeds up as the pathway is suppressed, while the return depends on pathway activity, so cells can only come back once signalling recovers. The right panel shows the signalling network. Ligand acts on the receptor and the receptor loads RAS with GTP, while BRAF V600E protein reaches MEK by two routes that differ in their sensitivity to drug, of which the monomer route does not require RAS and is blocked effectively by a RAF inhibitor, while the dimer route requires RAS to be loaded with GTP and resists the same inhibitor. ERK induces three things together, which are the receptor itself, a SPRY and MIG6 like brake on the step from receptor to RAS, and DUSP, so that suppressing ERK lowers all three at once. Because the brake falls further than the receptor does, the net effect is that signalling through the receptor arm becomes easier rather than harder, which is what raises RAS-GTP and lets the dimer route engage. The degrader removes BRAF protein rather than blocking its activity, while the receptor tyrosine kinase inhibitor blocks the receptor arm. The model carries three drugs in all, these two and a RAF inhibitor, which acts on the monomer and dimer routes and is the drug used for the signalling calibration in Figure 3 but is not drawn as a separate box. Every scheduling result reported here uses the degrader alone.

The relationship between these routes produces monomers, causing ERK activity to fall and the characteristic rebound in signaling after relieving the negative feedback that ERK treatment begins. Lito et al. described a normally exerts on upstream signaling. sequence in which the inhibitor blocks RAF Receptor-driven signaling then reactivates,

RAS-GTP rises, and RAF dimers form. Because these dimers resist the inhibitor, pathway activity settles at a new level above its initial minimum (21). Our model reproduces this sequence by making RAS activation responsive to feedback relief while retaining RAS-independent signaling through mutant RAF monomers.

Two further features follow from the data rather than modeling convenience. Gerosa et al. found that ERK drives expression of both the receptor and its negative regulators. Vemurafenib exposure reduced the abundance of six proteins—the EGF receptor, SPRY2, SPRY4, MIG6, DUSP4, and DUSP6—by 2.5- to 35-fold (8). Receptor abundance therefore falls during treatment, but transmission through the receptor arm nevertheless increases because the negative regulators decline more strongly. We also assign each transfer step in the cascade a cooperative form, following the measurements of Huang and Ferrell, because without this feature the model cannot reproduce the experimentally observed depth of pathway suppression (19).

2.3 Representation of inhibitors and protein degraders

The signaling measurements used for calibration originate from experiments with RAF inhibitors, either vemurafenib or dabrafenib, while the therapy whose schedule we study is a degrader. The model also includes a receptor tyrosine kinase inhibitor to represent blockade of the receptor-mediated signaling route that remains active when the monomer route is inhibited. However, no scheduling result reported here uses that

inhibitor. Representing kinase inhibition and protein degradation as a single generic process would simplify calibration but obscure the distinction central to scheduling: a degrader's effect persists until the protein is resynthesized, rather than ending when the drug is cleared.

A kinase inhibitor acts while it occupies its target, whereas a degrader removes the protein and its effect persists until the cell produces more. Alabi et al. measured this recovery directly, reporting 30% recovery of BRAF levels and MAPK phosphorylation during the 24 hours after the compound was washed out following a 24-hour treatment (22). Protein resynthesis therefore introduces a timescale that must be considered when scheduling degrader doses.

The model includes an explicit pool of BRAF protein, with the degrader acting on the protein rather than directly on the signal it produces. Degradation potency is drawn from the same study, which reports half-maximal degradation at 6.8 nM and more than 95% maximal degradation in SK-MEL-28 cells. The compound also degrades the p61 splice variant of BRAF V600E, a dimeric form resistant to vemurafenib, although approximately tenfold less potently (22).

Because degraders act by bringing two proteins together, their dose–response relationships can be nonmonotonic. At high concentrations, the compound can saturate each partner separately rather than bridge them. This loss of activity, known as the hook effect, produces the bell-shaped curves described by Gadd et al. (23). We represent this effect explicitly (Figure 2),

although the particular functional form is a modeling assumption rather than a fitted result. In the model, suppression is consequently deepest at an intermediate concentration; increasing the dose beyond that point reduces activity.




Figure 2. Degradar dose response, showing the hook effect. (A) The fraction of BRAF held in the three way complex with the degrader and the ligase, which is the form that leads to degradation, plotted against degrader concentration. (B) The deepest suppression of ERK activity reached at each concentration, relative to untreated. Suppression of ERK activity is deepest at an intermediate concentration and becomes weaker at high concentration, because a bifunctional molecule that saturates each of its two targets separately forms fewer of the three way complexes it needs in order to work. The consequence for dosing is that increasing the concentration beyond the optimum is counterproductive, which is the opposite of what a model treating the degrader as an ordinary inhibitor would predict.

Because the calibration experiments, simulated therapy, and published scheduling studies involve different treatments, Table 1 identifies the treatment represented in each part of the study. This distinction makes clear which comparisons concern drug action and which concern schedule structure.

Three distinctions follow from this table. The signaling parameters are estimated from RAF inhibitor experiments, while the therapy whose schedule we study is a degrader. This transfer rests on the drugs acting on the same pathway through different mechanisms represented

separately in the model, rather than on their being interchangeable. The receptor tyrosine kinase inhibitor is included in the model and its effects are reported in the pathway-suppression results, but no population result uses it because it contributes little suppression on its own. Finally, the comparisons with Das Thakur et al. and SWOG S1320 concern schedule structure—four weeks on and two weeks off, or five weeks on and three weeks off—rather than the drugs administered. These results should therefore not be interpreted as direct simulations of either published experiment.

Table 1. Treatments used in model calibration, scheduling simulations, and experimental comparisons

Part of the study	Treatment simulated	Measurement anchored to
Signalling calibration: ERK suppression, rebound and feedback at 24 hours	RAF inhibitor, vemurafenib, 1 micromolar	Gerosa et al. (8)
Signalling calibration: single cell escape at 1 and 10 micromolar	RAF inhibitor, dabrafenib	Yang et al. (7)
Degrader pharmacology: potency, maximal degradation, recovery after drug removal	BRAF degrader at constant concentration	Alabi et al. (22)
Outcome distribution, schedule comparison, response map, ablation, prior sensitivity, exposure matching	BRAF degrader alone at 1,000 nanomolar, gated by the duty cycle	not fitted to any single experiment
Held out comparison by time to resistant disease	as above	Das Thakur et al. (12), vemurafenib in mice
Held out comparison by progression free survival	as above	Algazi et al. (13), dabrafenib with trametinib in patients

2.4 Tumor populations and state transitions

Cells occupy one of three states. The bulk population proliferates at a rate dependent on ERK activity and is susceptible to treatment-induced killing. The persister population tolerates the drug and divides slowly while retaining the ability to revert to the bulk state. The resistant population has restored its signaling and is no longer suppressed by the drug. Oren et al. showed that persisters are not uniformly quiescent: individual cells ranged from never dividing to dividing seven times over 14 days (24). We therefore assign the persister state a small positive growth rate rather than treating it as fully arrested.

Two features of this arrangement warrant explicit explanation because, in each case, the simpler alternative produces behavior is unsupported by the measurements. First, conversion into the persister state depends on treatment rather than occurring at a constant

rate. A constant rate drives even an untreated population toward a large equilibrium persister fraction, contrary to observations in untreated cells. Making conversion depend on the degree of pathway suppression removes this problem and is consistent with the observations of Sharma et al. and Yang et al. that drug-tolerant states arise during treatment and can reverse after withdrawal (5, 7).

Second, reversion from the persister state depends on pathway activity. Yang et al. describe escapees as cells that restore CDK2 activity before reentering the cell cycle (7). Because cell-cycle reentry requires mitogenic signaling, the model makes reversion dependent on ERK activity. This reproduces the observed direction of dose dependence: the fraction of escaping cells decreases as dose increases in Yang et al.'s data, whereas a constant reversion rate produces the opposite pattern.

Movement from the persister state into the resistant state follows the structure observed by Sharma et al., in which $\sim 20\%$ of persisters eventually resume proliferation in the presence of drug and can subsequently be maintained under continued treatment (5). The model structure admits a preexisting resistant population, because Hata et al. showed that resistance can arise both from preexisting resistant cells and from drug-tolerant cells during prolonged treatment (11), but every simulation reported here starts with no resistant cells, so all resistance in these results arises through the persister state.

2.5 Drug dependence of resistant cells

The drug-dependence parameter scales how fast resistant cells grow when the drug is absent relative to when it is present. We include it because Das Thakur et al. found that withdrawing vemurafenib from mice carrying resistant tumors produced clear signs of regression within ten days while simultaneously increasing MEK and ERK phosphorylation (12). Their interpretation is that resistant cells overproduce BRAF V600E and that the resulting excess signaling becomes harmful once the drug no longer restrains it.

The parameter ranges from one, indicating that resistant cells grow equally well with or without the drug, through zero, indicating that growth stops after withdrawal, to negative values, indicating regression. Values below zero represent the drug addiction observed by Das Thakur et al. The sign of this parameter therefore distinguishes resistant-cell regression after withdrawal from settings in which

withdrawal arrests growth or allows it to continue.

2.6 Population scaling and extinction rules

Expressing populations as fractions of a carrying capacity allows a model to report surviving fractions on the order of 10^{-32} . Such fractions have no biological meaning at realistic cell counts, yet a continuous model that permits them allows a fraction of a cell to regrow into a tumor. This is a recognized difficulty in continuous population models and is usually addressed by imposing a floor.

We therefore work in absolute cell numbers, anchored to the approximately 18,000 cells that Shaffer et al. plated before applying the drug (6), and set any population that falls below one cell to zero. Eradication in this model consequently means that no cells remain, rather than that the population has reached an arbitrarily small positive value.

Because the schedules being compared differ in how often integration is interrupted, the timing of floor application matters. The floor is applied in two places. First, it is applied to the recorded trajectory at every point on the output grid. This grid comprises the same 1,600 points spanning the observation horizon for every schedule, so all endpoints, including eradication, are evaluated on a common grid. Second, the floor is applied to the state carried forward from one integration segment to the next. Segments end at every duty-cycle switch, so the frequency of this application depends on the schedule: none under continuous dosing and several hundred under the shortest cycles. The classification procedure is therefore

schedule-independent by construction, whereas flooring of the propagated state is not.

Supplementary Table 6 reports the effect of this second application by repeating the schedule comparison with flooring removed from the propagated state, applied instead at fixed schedule-independent intervals, and applied at thresholds of 0.1 and ten cells. Supplementary Table 7 repeats the analysis for the exposure-matched blocks in Section 3.9, where two intermittent arms are compared with one another rather than with continuous dosing, and the segmentation asymmetry within a contrast is therefore greatest.

2.7 Calibration data and parameter estimation

We fit the model simultaneously to measurements from three studies, so that the signaling, drug-response, and cell-state components have to agree with the data together rather than being fitted separately.

The first study supplies the signaling measurements. Gerosa et al. treated A375 melanoma cells with 1 μ M vemurafenib and followed pathway activity over the subsequent 24 hours. They also published simulation output from their rule-based model of the same system (8). From this material, we obtain eight quantities: the fractions of phosphorylated ERK and MEK in untreated cells, the depth of the initial decline in ERK activity, its level after 24 hours, the corresponding level of phosphorylated MEK, and the fold changes in the receptor, SPRY negative regulators, and DUSP over the same period. We additionally use their growth-factor challenge experiment in cells already adapted to vemurafenib, taking

the height and timing of the resulting spike in ERK activity. These two measurements are the only available targets that constrain the dimer route; without them, the RAS-driven arm of the model would lack calibration targets.

The second study supplies the drug-response measurements. Fallahi-Sichani et al. measured dose-response relationships across a panel of cancer drugs and showed that potency alone is an incomplete description of cellular response, because curve steepness and maximum effect vary independently of it (25). We use this work to set the concentration producing half-maximal effect and the steepness of the response.

The third study supplies the degrader pharmacology. Alabi et al. characterized a vemurafenib-based BRAF degrader and reported half-maximal degradation at 6.8 nM, more than 95% maximal degradation, maximal degradation within four hours at 100 nM, and 30% recovery of BRAF protein during the 24 hours after compound washout (22). We treat the reported degradation potency as an emergent property of the model rather than a binding parameter. Potency measured at a fixed time reflects both how readily the three-way complex forms and how rapidly the protein is degraded; setting the binding constant equal to the reported potency would conflate these distinct quantities.

We also report two single-cell escape measurements from Yang et al., obtained at 1 μ M and 10 μ M dabrafenib (7), without fitting them, for the reasons given in Section 4.4. The target set therefore contains 18 quantities, of

which 16 are fitted and two are reported for comparison. Supplementary Table 2 identifies each category; references to 16 targets elsewhere in this paper refer to the fitted targets.

Supplementary Table 2 provides all targets and their sources, with each labeled according to how it was obtained. Some are values reported in the literature, some were calculated from the simulation output published by Gerosa et al., and one is an imposed requirement rather than a measurement: that the adapted state persists for several days rather than decaying, a condition not constrained by the available data. These distinctions affect the interpretation of fit quality. Agreement with an experimental measurement, published simulation output, or an imposed assumption provides different forms of evidence, and labeling each target makes the basis of calibration explicit.

Fitting used sequential approximate Bayesian computation following Toni et al. (26). Sequential populations make the approach feasible with broad priors, because drawing candidates directly from the prior and rejecting those outside the final tolerance would yield very few accepted sets. We ran 12 populations of 300 accepted parameter sets, evaluating 42,404 candidate sets in total.

2.8 Parameter identifiability

We report identifiability using conditional loss profiles, in which each parameter is swept across its full prior range while the others are held at their best-fitting values. We record how much of that range still produces an acceptable fit.

We use this measure rather than posterior spread because the latter can be misleading for this sampler. Sequential methods narrow their sampling distribution at each generation, concentrating particles around the region to which the sampler has converged. In a flat region of parameter space, this can produce a tight cluster even for parameters that the data do not constrain. In our analysis, posterior spread suggested that all 33 parameters were tightly constrained, whereas nine retained an acceptable fit across their entire prior range.

2.9 Virtual tumor generation and parameter sampling

Each simulated tumor, which we call a virtual subject, is defined by a calibrated parameter set drawn from the posterior, a phenotype selected from one of four hypothetical tumor types, and sampled values for parameters that the calibration data do not constrain. Every subject has an identifier, and the same subjects are used for every schedule and model variant, allowing outcomes to be compared within a tumor rather than only between groups.

The four tumor types are hypothetical phenotypes rather than empirically derived patient groups. Tsoi et al. and Verfaillie et al. describe transcriptional programs separating proliferative from invasive melanoma states and report higher receptor tyrosine kinase expression and greater therapy tolerance in less differentiated states. Neither study, however, reports rate constants for a model of this kind (27, 28). The parameter ranges therefore encode the reported direction of each difference rather than measured magnitudes.

The parameters unconstrained by the calibration data form a coherent group: all concern cell-population dynamics rather than signaling. Changing any of them alters every fitted observable by less than two parts in ten million, consistent with numerical integration noise. Rather than fitting them and reporting a posterior that merely reproduces the prior, we sweep them across their ranges and treat them as the axes of the map, because population dynamics is central to the scheduling question.

2.10 Treatment schedules and exposure matching

We compare a continuous reference arm with five intermittent schedules: four weeks on and two weeks off, as used by Das Thakur et al.; five weeks on and three weeks off, as used in SWOG S1320; two weeks on and one week off; one week on and one week off; and three days on and three days off. Drug concentration is identical in every arm, following the approach used by Das Thakur et al., so arms containing holidays receive less cumulative drug exposure. This confounds treatment timing with total exposure. We address this confound through a separate comparison in which cumulative exposure is held fixed and only the cycle period varies.

We compare these six schedules rather than sweeping a larger grid of doses and intervals because the mechanism under study concerns responses to treatment and withdrawal. Different analyses use different observation horizons: 180 days for the outcome distribution under continuous therapy, 400 days for the schedule comparison and ablation to allow the longest cycles to complete, and 378 days for

the exposure-matched comparison. The latter is an exact multiple of every cycle period included in that comparison. Each horizon is specified alongside its analysis in the Supplementary Methods.

2.11 Outcome classification and statistical analysis

We classify each simulated tumor by outcome and then examine which parameters differ between the classes, using three methods with different assumptions.

The first is regionalized sensitivity analysis, developed by Spear and Hornberger, which compares the distribution of each parameter between controlled and progressing tumors using a two-sample Kolmogorov–Smirnov statistic (29). The second is the PAWN index of Pianosi and Wagener, which measures sensitivity by comparing cumulative distribution functions (30). We chose this method because the outcome distribution is strongly skewed and would not be adequately summarized by a variance-based index. The third is partial rank correlation, as described by Marino et al. for systems biology models (31).

For threshold rules, we use shallow decision trees and check the number of outcome events per predictor before fitting. This check is informed by Peduzzi et al.'s finding that regression coefficients can become unreliable below approximately ten events per variable and may acquire significance in the wrong direction (32).

Every predictor entering a threshold rule is a baseline biological property fixed before

treatment begins. We avoid using pathway suppression measured under the schedule being evaluated, because doing so would partly restate the treatment response rather than identify biological conditions that predict it.

Where a hazard ratio is reported, we also assess the proportional-hazards assumption using scaled Schoenfeld residuals in the form described by Grambsch and Therneau. Where the assumption fails, we additionally estimate the association within separate follow-up windows, using episode splitting with left truncation so that each window includes only tumors still at risk when it begins. Inference is based on a bootstrap interval for the correlation between residuals and time, clustered by the calibrated parameter set from which each tumor was generated. This preserves consistency with the other analyses, because tumors sharing a parameter set are related simulations rather than independent draws.

Two conventions govern reporting. First, the 1,200 tumors derive from 300 calibrated parameter sets, with four tumors per set. Reported intervals are therefore obtained by resampling whole parameter sets rather than individual tumors, and cross-validation folds are grouped so that no parameter set appears in both the training and test sets. Second, asserted effects are reported as effect sizes with intervals rather than solely as p values, because analytic significance can depend strongly on the number of tumors chosen for simulation. Where p values are reported, they are derived either from permutation tests that exchange whole parameter sets or from tests assessing evidence of an association. A large p value

indicates insufficient evidence against the null hypothesis and is not, by itself, evidence that an effect is absent.

2.12 Comparisons with independent experimental datasets

Three datasets were excluded entirely from calibration and used only for testing: the murine scheduling experiment of Das Thakur et al., the persister measurements of Sharma et al., and the resistant-colony counts of Shaffer et al.

In particular, we did not tune the resistant compartment to reproduce Shaffer's measurement. Although such adjustment could improve agreement, it would make the measurement part of calibration rather than an independent test. Excluding it from fitting preserves its value for evaluating the model.

3. Results

3.1 Pathway suppression and signaling rebound

From an untreated steady state, a saturating concentration of the degrader alone reduces ERK activity to 4.7% of baseline. Adding a receptor tyrosine kinase inhibitor reduces it to 4.1%, whereas the inhibitor alone produces little suppression, leaving ERK activity at 98.5% of baseline. Every population result reported below uses the degrader alone. Across the 1,200 calibrated tumors, median residual ERK activity is 0.19% of baseline. Both degrader-alone estimates exceed the 80% inhibition of ERK phosphorylation reported by Bollag et al. in biopsies from patients whose tumors regressed (20), placing the modeled

suppression within the range associated with clinical responses.

The simulated signaling response follows the sequence described by Lito et al. (Figure 3). Applying a RAF inhibitor at 1 μ M rapidly reduces ERK activity, after which RAS-GTP

risks and ERK activity partially rebounds. The dimer contribution to RAF activity increases from approximately 13% at the start of the plotted trajectory to a peak of approximately 62%, before declining to approximately 39% at 72 hours.




Figure 3. Signaling response to a RAF inhibitor at 1 μ M. (A) ERK activity and RAS-GTP relative to their untreated levels. ERK activity falls rapidly after inhibitor application and then partially rebounds as RAS-GTP increases. (B) The dimer contribution to total RAF activity increases from approximately 13% at the start of the plotted trajectory to approximately 62% near 15 hours, then declines to approximately 39% at 72 hours.

3.2 Agreement with calibration targets

Twelve of the 16 fitted targets are matched within one assumed standard deviation, as recorded in Supplementary Table 2 (Figure 4). Except where a published interval was available, these standard deviations represent chosen tolerances rather than measurement uncertainties. The agreement should therefore be interpreted as falling within stated tolerances rather than experimental error. The

four targets outside these tolerances are the depth of the initial ERK minimum, the level of phosphorylated MEK at 24 hours, and two features describing the timing and magnitude of the ligand-challenge response.

The fit yields a median normalized χ^2 statistic of 0.62 across the fitted targets. This value requires careful interpretation: with 16 targets and 33 fitted parameters, the conventional

degrees of freedom are negative, so the statistic cannot be interpreted as a conventional reduced χ^2 . A value below one indicates residuals smaller than the assumed tolerances, rather than an unusually good fit or evidence that the parameters are well constrained.




Figure 4. Calibration. For each target the blue marker and line show the median and 95 percent range produced by the accepted parameter sets, in grey where the target was not fitted, and the orange marker and line show the measured value with its assumed uncertainty. Twelve of the sixteen fitted targets are matched within one of the assumed standard deviations listed in Supplementary Table 2, which are tolerances we chose rather than measurement uncertainties, except where a published interval was available. Two further quantities, the fraction of cells escaping treatment at one and at ten micromolar, are reported for comparison but were not fitted, for the reason given in Section 4.4. The model misses both by about four orders of magnitude, at ratios of 1.1×10^{-4} and 7.6×10^{-5} , which falls outside the axis; those two rows carry a left pointing marker at the edge with the ratio beside it rather than a model interval.

3.3 Parameter identifiability

When each parameter was swept across its prior range while the others were held fixed, 11 of the 33 parameters retained an acceptable fit over less than a quarter of that range and were therefore meaningfully constrained by the data

(Figure 5). Nine parameters retained an acceptable fit across their entire prior range and were unconstrained by this criterion. An analysis of correlations between the fitted observables yielded 5.78 effectively independent constraints from 16 targets, giving

a ratio of effective constraints to parameters of $\sim$ one to six.

The constrained parameters include the rates governing RAS activation and deactivation, the strength of the autocrine ligand signal, several cascade parameters, and the affinity that

determines degrader potency. Adding the ligand-challenge targets provides constraints on the dimer route that are absent from the other calibration data. These targets therefore constrain the signaling route that motivated the model's restructuring.




Figure 5. How well the data constrain each parameter, measured by sweeping each parameter across its full prior range while holding the others at their fitted values and recording the fraction of that range which still produces an acceptable fit. Lower bars indicate better constrained parameters. Eleven of thirty three parameters accept less than a quarter of their range. We report this measure rather than the spread of the fitted posterior, because for a sequential sampler the posterior spread reflects how far the sampler travelled rather than what the data require.

3.4 Outcomes under continuous treatment

Across 1,200 simulated tumors receiving continuous therapy for 180 days, 75.2%

progressed, 22.8% remained controlled, and 1.9% were eradicated. Time to progression

ranged from 2.4 days to the observation horizon, with a median of 36.2 days (Figure 6).

The outcomes are therefore graded rather than divided into two widely separated groups, allowing the boundary between control and

progression to be examined. The smaller outcome class contains 297 tumors, sufficient to support a boundary analysis with several predictors at the events-per-predictor ratio recommended by Peduzzi et al. (32).




Figure 6. Outcome distribution under continuous therapy, across 1,200 simulated tumors followed for 180 days. (A) The proportion of tumors falling into each outcome class within each of the four hypothetical tumor types, where a tumor is counted as eradicated if its population falls below one cell, as controlled if it neither reaches one cell nor returns to its starting size, and as progressed if it returns to its starting size after reaching a minimum. (B) The distribution of time to progression, with the median marked. The 903 tumors that progressed are the grey histogram. The 297 that never progressed cannot be placed on this axis and are drawn separately in green at the horizon, where they are held for the purposes of the median; they are not a tall final bin of the same distribution. Among the tumors that progressed the distribution is graded rather than divided into two separated groups, which is what allows a boundary between the two outcome classes to be located at all. The marked median is a Kaplan Meier median, which is well defined here because all censoring occurs at the horizon.

3.5 Net resistant growth and the tumor-control boundary

Regionalized sensitivity analysis, the PAWN index, and decision-tree classification identified the same parameter as separating tumors that remained controlled from those that progressed (Figure 7). The difference between the resistant-cell growth rate and the baseline rate of cell turnover produced a Kolmogorov–Smirnov statistic of 0.931 in the regionalized sensitivity analysis, compared with 0.782 for resistant-cell growth alone and 0.320 for turnover alone; every other parameter had a

statistic below 0.14 (Figure 8). Resampling whole calibrated parameter sets resulted in corresponding 95% intervals of 0.917–0.951, 0.753–0.823, and 0.266–0.392. All three quantities separated the outcome classes under a permutation null obtained by exchanging whole parameter sets, with $p < 0.001$ based on 2,000 permutations. We report intervals and permutation-based inference rather than analytic p values, which fall below 10^{-200} for separations this large and primarily reflect the number of tumors simulated rather than the biological importance of the association. The

PAWN index, which is computed against time to progression rather than against the outcome classes, resulted in the same first six.

A decision tree fitted to net resistant-cell growth classified outcomes with 96% accuracy under cross-validation, using a single threshold

of $\sim 0.0008 \text{ h}^{-1}$. This corresponds to a net doubling time of ~ 36 days after accounting for baseline cell turnover. Adding further parameters did not improve classification, indicating that the boundary is $\sim$ one-dimensional.




Figure 7. The boundary between control and progression, in the parameter space that determines it. Each point is one simulated tumor, positioned by the net growth rate of its resistant population, meaning the resistant growth rate minus the baseline rate of cell turnover, on the horizontal axis, and by the degree to which its resistant cells depend on the drug for their own fitness on the vertical axis. Blue points did not progress and orange points did. Panel (A) shows continuous therapy and panel (B) shows the same tumors under a four weeks on and two weeks off schedule. Both panels are drawn from the schedule comparison at a four hundred day horizon, which is longer than the one hundred and eighty day horizon used for the outcome distribution in Section 3.4, so the percentages here are not comparable with the ones reported there. The dashed line in each panel is the threshold found by a decision tree fitted to that panel alone, and the accuracy printed beside it is cross validated over five folds and given against the majority class base rate. Two features are visible. The first is that the boundary is essentially vertical, so outcome is decided by net resistant growth and not by drug dependence. The second is that the intermittent schedule does not move the boundary but blurs it: the fraction of tumors not progressing is 19.2 percent under continuous dosing against 18.8 percent under the intermittent schedule, the fitted threshold moves from 0.30 to 1.36 in units of ten to the minus three per hour, and the cross validated accuracy falls from 0.99 to 0.85 against a base rate of 0.81. Counted individually, 71 tumors are rescued from progression and 76 are pushed into it, which is the net negative effect reported in Section 3.11.

Two qualifications are important when interpreting this accuracy. First, progression is defined as the total population returning to its starting size, and the resistant compartment drives that return. A threshold based on its net

growth rate therefore follows largely from the model structure rather than constituting an unexpected biological finding. The result establishes that the model behaves as its equations require and locates the boundary

between outcomes. Second, the 1,200 tumors derive from 300 calibrated parameter sets, with four tumors per set, so a classifier could, in principle, perform well by recognizing a parameter set already represented in training. Grouping cross-validation folds by calibrated parameter set prevents related tumors from appearing in both training and test sets. This changes accuracy by only 0.0003, less than the variation between fold assignments (Table 2).




Figure 8. Agreement between three sensitivity methods on which parameters govern the outcome. Bars show the Kolmogorov Smirnov statistic from regionalized sensitivity analysis, which is computed against the two outcome classes, together with the PAWN index from cumulative distribution functions and the absolute partial rank correlation coefficient, both of which are computed against time to progression rather than against the classes. Bars are shown for each swept parameter and for the two composites, net resistant growth and the conversion to reversion ratio, which are formed from swept parameters after sampling. The first two agree that net resistant growth dominates, while partial rank correlation ranks the rate of transition into resistance first, which reflects that it measures association with how quickly a tumor progresses rather than with whether it does.

All four configurations used the same 1,200 assignments. Under the original split, every test tumor, with a baseline classification accuracy of 0.7525; the standard deviations are tumors in the training set. Removing this calculated across five random fold leakage changed accuracy by less than the

variation between fold assignments. The last row in Table 2 shows the configuration originally reported: an unstratified grouped split produced folds whose class proportions ranged from 0.196 to 0.288, compared with 0.246–0.250 under stratification. Its apparent accuracy penalty of 0.007 therefore reflects fold composition rather than leakage.

Table 2. Outcome-classification accuracy under alternative cross-validation configurations

Cross validation	Accuracy	Test tumors sharing a parameter set with the training half
five fold, as first used	0.9583	100.0 percent
five fold, folds reshuffled	0.9562, sd 0.0032	98.9 percent
grouped by parameter set, stratified	0.9580, sd 0.0012	0.0 percent
grouped by parameter set, not stratified	0.9517	0.0 percent

This finding agrees with Strobl et al.'s conclusion that turnover is a key determinant of whether adaptive therapy can control a tumor (16). The agreement is notable because their model represents competition between two cell populations without explicit signaling, whereas ours incorporates signaling calibrated to melanoma phosphoprotein data.

Partial rank correlation resulted in a different ordering. The correlation ranks the rate at which persisters become resistant first and net resistant growth second. This difference is consistent with the methods addressing different outcomes: partial rank correlation measures association with time to progression as a continuous quantity, whereas regionalized sensitivity analysis and the PAWN index assess the classified outcome. The former therefore addresses how quickly a tumor progresses, while the latter address whether it progresses within the observation period.

3.6 Drug dependence under continuous treatment

The parameter describing whether resistant cells lose fitness after drug withdrawal shows no association with outcome under continuous therapy, with a Kolmogorov–Smirnov statistic of 0.060 and a p value of 0.37. This behavior is expected because the parameter scales growth in the absence of drug, whereas continuous therapy never withdraws treatment. The parameter therefore has no opportunity to act, which is precisely why the scheduling question cannot be answered from the reference arm alone.

3.7 Biological conditions governing scheduling benefit

We predicted in advance that drug dependence would dominate the benefit a tumor receives from an intermittent schedule. Across all 1,200 tumors, however, its rank correlation with benefit was only -0.113 , compared with -0.307 for the bulk growth rate. The prediction, as originally stated, was therefore not supported.

The effect is conditional. In the four-week-on schedule, the first drug holiday begins on day 28. Resistance already established before that point cannot be delayed, while a delay cannot be observed if resistance never develops within the observation period. Grouping tumors by when resistant disease appears under continuous therapy identified 512 that reached it before day 28, 381 that reached it on day 100 or later or remained censored, and 307 in the intermediate window. Two conditions must be distinguished. The first is temporal eligibility, which follows from the schedule rather than constituting a biological finding: intermittent and continuous dosing were identical until day 28. Accordingly, 687 tumors have identical times to resistant disease under both schedules, comprising all 512 early cases, 170 that never

reach resistant disease under either schedule within follow-up, and five within the intermediate window. The second condition is biological: among tumors still at risk when treatment is first withdrawn, benefit depends on how much fitness resistant cells lose without the drug.

Within the intermediate group, this association was strong, with a rank correlation of -0.666 , and benefit varied systematically with drug dependence (Figure 9). Tumors whose resistant cells were indifferent to withdrawal gained a median of 3.3 days before resistant disease appeared, with 82% showing some benefit. The two intermediate drug-dependence categories gained 13.0 and 25.5 days.




Figure 9. Benefit of intermittent dosing, restricted to the tumors where a drug holiday is able to act. (A) The additional time before resistant disease appears under a four weeks on and two weeks off schedule, relative to continuous dosing, against the drug dependence of the resistant population, with colour indicating net resistant growth. (B) The median benefit of each schedule relative to continuous dosing. Both panels are restricted to the 307 tumors whose time to resistant disease under continuous treatment lies between day 28 and day 100. The lower edge is the first drug holiday of the four weeks on and two weeks off schedule; the upper edge is a choice, not a horizon, and the observation period runs to 400 days. Of the 381 tumors excluded above it, 211 have a non-zero benefit under the four week schedule. In panel B the same window is applied to all five schedules so that every bar covers the same tumors. Three of the five take their first holiday earlier than day 28, so for those the window discards the tumors in which they act soonest and flatters them: the grey bars give each schedule's median over its own first holiday window instead, and on that basis the ordering is reversed. The ranking of record is the restricted mean over the whole cohort in Table 5.

The most drug-dependent category requires particular care because 21 of its 84 tumors did not reach resistant disease within the observation period. Assigning censored times to the observation horizon resulted in a median gain of 53.5 days, whereas restricting the comparison to pairs in which both arms reached resistant disease resulted in 29.0 days. Neither calculation fully characterizes the benefit: the first substitutes the horizon for an unknown event time, while the second excludes tumors whose resistance may have been most effectively delayed. The two calculations are not symmetric. A tumor held at the horizon enters at the time remaining rather than at the time it would have taken, which is below its true value, so 53.5 days is a lower bound on the median gain rather than an inflation of it, up to the quarter day spacing of the output grid. Restricting to pairs in which both arms reached resistant disease removes the longest delays

and understates it further. The gain in this category is therefore at least 53.5 days. The remaining categories were less affected. Across the intermediate group, 8.1% of tumors possessed at least one censored arm; estimates for the other three categories changed by no more than 1.1 days, and the rank correlation changed from -0.666 to -0.620 .

Restricting analysis to uncensored pairs did not resolve censoring, because it selects tumors that reach resistant disease within follow-up and alters the distribution of drug dependence. We therefore also fitted a Cox proportional-hazards model for time to resistant disease, retaining censored observations, with drug dependence as the covariate and standard errors clustered by calibrated parameter set (Table 3). The proportional-hazards assumption and its implications are assessed below.

Table 3. Hazard ratio for drug dependence under each schedule, with standard errors clustered on the calibrated parameter set

Schedule	Hazard ratio for drug dependence	95 percent interval
continuous	0.92	0.83 to 1.02
five weeks on, three off	1.19	1.07 to 1.32
four weeks on, two off	1.24	1.12 to 1.38
two weeks on, one off	1.49	1.33 to 1.67
three days on, three off	2.99	2.65 to 3.37
one week on, one off	3.21	2.82 to 3.66

A hazard ratio above one indicates that resistant cells less affected by treatment withdrawal reach resistant disease sooner. The continuous arm provides a check on the method: the drug is never withdrawn, so the parameter describing fitness after withdrawal

has no opportunity to act, and its confidence interval includes one, as expected.

A single hazard ratio assumes that the association remains constant throughout follow-up. There is reason to question that assumption here, because the effect of drug

dependence cannot emerge until treatment is withdrawn. A scaled Schoenfeld residual test supports this concern. Because the 1,200 tumors comprise four related simulations from each of 300 calibrated parameter sets, inference uses a bootstrap interval clustered by parameter set for the correlation between scaled residuals and ranked event time. This accounts for dependence between simulations rather than treating events as independent.

Under continuous dosing, there was no evidence against proportional hazards: the correlation was -0.038 , with an interval of -0.100 to 0.022 , which included zero. Under all five intermittent schedules, the intervals excluded zero. The correlations were 0.196 and 0.179 for the four- and five-week cycles, 0.102 for the two-week cycle, and -0.142 and -0.272 for the one-week and three-day cycles, respectively (Supplementary Table 4 and Supplementary Figure 1). The differing signs are consistent with the later onset of holidays in longer cycles and their earlier onset in shorter

cycles; both patterns indicated that the association changed over time.

Estimating the association separately within follow-up windows clarified this variation. Under continuous dosing, the hazard ratios remained near one across the four windows, at 0.94 , 0.88 , 0.81 , and 1.11 . Under the four-week cycle, the ratio increased from 1.12 in the first 50 days to 3.56 after 200 days; under the five-week cycle, it increased from 0.98 to 3.57 . Under the three-day and one-week cycles, whose first holidays occur earlier, the ratios already exceeded four in the first window and subsequently tended to decline. These patterns are inconsistent with a constant hazard ratio and suggest that the association reflects the timing and accumulation of treatment breaks (Supplementary Table 5). We therefore report the single ratios as summaries over follow-up rather than constant effects. The schedule comparisons in the next section use differences in restricted mean survival time and do not require proportional hazards.

Table 4. Sensitivity of the scheduling-benefit region to parameter-sampling assumptions

Sampling configuration	Window size (%)	Gain over 7days all tumors (%)	Gain over 7 days, inside window (%)	Median gain inside the window (Days)	Axis governing benefit	Correlation with drug dependence
as published	22.7	25.3	73.5	18.9	drug dependence	-0.645
every log axis halved in width	56.0	62.0	91.1	20.0	net resistant growth	-0.569
every log axis widened threefold	12.7	13.0	68.4	12.9	drug dependence	-0.511
uniformly rather than log uniformly	9.3	7.3	71.4	16.1	not determined	-0.360

Scheduling benefit is thus concentrated within a conditional region of parameter space, identifying where the intervention may help rather than establishing a treatment

recommendation. Across four sampling configurations, this region remained present and the ordering of schedules was stable, but its size varied from 9.3% to 56.0% of simulated

tumors (Table 4). The association with drug dependence persisted, although the screening analysis did not identify the same governing parameter in every configuration. Because the population parameters defining this region were varied rather than measured, its size reflected sampling assumptions and cannot be interpreted as prevalence in melanoma.

Each configuration was run on 300 tumors, one per calibrated parameter set. The mechanism remained consistent across configurations: a conditional region existed in all four, 68–91% of tumors within it benefited, the median gain remained between 13 and 20 days, and the correlation with drug dependence retained its direction and $\sim$ magnitude. The size of the region, however, was not stable and is therefore not reported as a prevalence estimate. The governing parameter was not determined for the final configuration because its window contained too few tumors [28] for the screening procedure. Narrowing the priors enlarges rather than shrinks the window by concentrating resistant growth rates near the threshold at which a holiday can act. The four configurations therefore do not represent a progression from optimistic to conservative assumptions.

3.8 Effects of scheduling on resistance and progression

An unanticipated result concerns the difference between two endpoints often treated as interchangeable: time to resistant disease and overall time to progression.

Four of the five intermittent schedules increased the estimated time to resistant disease (Figure 10). Using restricted mean survival time over the observation horizon and paired comparisons, because every schedule was run on the same tumors, the three-day cycle gained 33.6 days, the one-week cycle 29.7 days, and the two-week cycle 13.6 days. Paired confidence intervals excluded zero in each case (Table 5). The four-week cycle gained 5.5 days, with an interval of 0.0–10.6 days, indicating a marginal benefit, whereas the five-week cycle used in SWOG S1320 showed no benefit, with an estimated difference of -0.1 days. For overall time to progression, however, only the three-day cycle showed an improvement with an interval excluding zero. The two schedules used in the published experiments therefore produced the smallest estimated delays in resistant disease among the intermittent schedules tested.

Table 5. Effects of intermittent treatment on restricted mean time to resistant disease and overall progression

Schedule	Gain in time to resistant disease	95 percent interval	Gain in time to progression	95 percent interval
three days on, three off	+33.6 days	28.3 to 39.0	+18.3 days	12.1 to 24.4
one week on, one off	+29.7 days	23.1 to 36.7	+3.5 days	-5.0 to 11.9
two weeks on, one off	+13.6 days	8.8 to 18.7	-1.8 days	-7.5 to 4.5
four weeks on, two off	+5.5 days	0.0 to 10.6	-3.5 days	-11.4 to 3.8
five weeks on, three off	-0.1 days	-5.5 to 5.7	+2.1 days	-6.4 to 10.9

Gains are differences in restricted mean survival time over the observation horizon relative to the continuous arm. Because every schedule was run on the same 1,200 tumors, each interval is obtained by resampling calibrated parameter sets once per replicate and recomputing both arms using the same draw, preserving the pairing. Table 5 highlights the distinction between the endpoints: four of the five schedules increased the estimated time to resistant disease, whereas only one improved overall time to progression with an interval excluding zero.

The model also distinguishes the direction of individual responses from the average effect. Under the five-week cycle, 308 tumors reached resistant disease later than under continuous dosing and 155 reached it earlier, a ratio of $\sim$ two to one, while the average effect across all tumors was $\sim$ zero. A schedule can therefore delay resistance in more tumors than it accelerates resistance without producing an average delay. We report the direction and magnitude of the effect separately, because the number of tumors benefiting does not, by itself, establish an overall benefit.




Figure 10. Median additional time before resistant disease appears, in days, for each schedule and each category of resistant cell drug dependence, within the tumors where scheduling can act. Categories run from resistant cells that regress strongly when the drug is withdrawn, at the top, to resistant cells whose growth is unaffected by withdrawal, at the bottom. The gradient down each column is the central result of the study, since it shows that the same schedule produces a delay of many weeks in one biological setting and a delay of a few days in another. The cells are medians of the per tumor difference, in which a tumor that never reaches resistant disease under the schedule enters at the four hundred day horizon, and each cell carries the number of its tumors in that state. The ten cells of the neutral and independent rows have no censored tumors in either arm and are therefore exact. In the ten that do, the entered value is below the true difference, so those cells are lower bounds rather than observed delays; under the one week cycle 44 of its 84 tumors do not reach resistant disease within the horizon, so the Kaplan Meier median is not reached at all. Summarising the same differences as a restricted mean to four hundred days rather than as a median, that row reads 187.8 days for the three day cycle, 118.6 for the four week cycle, 106.7 for the five week cycle, 223.4 for the one week cycle and 119.2 for the two week cycle.

3.9 Schedule comparisons at matched exposure

The preceding comparisons involved arms receiving different amounts of drug, because concentration was held constant while arms containing holidays spent less time on treatment. An improved outcome could therefore reflect holiday timing, lower cumulative exposure, or both. Matching cumulative exposure by increasing concentration in the intermittent arm did not resolve this distinction in the model. Drug effects on cell populations are mediated through pathway suppression, which is nearly saturated across the relevant concentration range: mean suppression was 0.9896 at 667 nM and 0.9875 at 1,500 nM. Thus, a 2.25-fold increase in concentration changed the driver of the population equations by $\sim$ two parts per thousand and left the endpoints of almost every tumor unchanged.

We therefore matched exposure by holding the fraction of time on treatment fixed and varying only the cycle period. Two blocks were compared, with treatment administered for either two-thirds or one-half of each cycle. Both were simulated over 378 days, an exact multiple of every cycle period, ensuring identical cumulative exposure within each block. The reference arms were the longest cycles at each duty fraction: four weeks on and two weeks off in the two-thirds block, and three weeks on and three weeks off in the one-half block.

Within both blocks, shorter cycles delayed resistant disease relative to these references. Among tumors whose outcomes differed between arms, resistance occurred later in

75.2%, 78.6%, and 74.0% of comparisons in the two-thirds block, and in 83.2%, 79.1%, and 68.8% in the one-half block. These percentages excluded tied outcomes, which accounted for 176–546 of the 1,200 tumors, depending on the comparison. We therefore also report differences in restricted mean survival time, retaining every tumor and accounting for censoring. Against the same references, these differences were 7.8, 14.8, and 14.3 days in the two-thirds block, and 20.8, 24.4, and 12.3 days in the one-half block. Intervals clustered by calibrated parameter set excluded zero in every case. Supplementary Table 3 provides every pairwise contrast within each block, identifies the reference arm, and reports the numbers of later, earlier, tied, and total outcomes.

The effect was not monotonic with cycle period. In the one-half block, the one-week and three-day cycles outperformed the one-day cycle by 8.5 days (interval: 2.7–14.3) and 12.1 days (interval: 8.2–16.1), respectively. In the two-thirds block, the two shortest cycles were not distinguishable, with a difference of 0.5 days and an interval of -1.4 to 2.6 days. These results indicate that cycle timing must suit the modeled biology rather than follow a general rule that shorter cycles are better.

3.10 Sensitivity to extinction handling

The arms within each block differed in how finely integration was segmented. The largest asymmetry occurred between the one-day cycle, with 378 duty-cycle switches, and the three-week reference, with 18. We therefore repeated every within-block contrast under the seven extinction-floor configurations described in Section 2.6, using 300 tumors, one per

calibrated parameter set, with the four phenotypes balanced. Under six configurations—including the published implementation, removal of flooring from the propagated state, continuous detection by root finding, a fixed schedule-independent interval and its matched control, and a 0.1-cell threshold—no contrast changed sign. The largest change among the 18 contrasts was 2.63 days.

This change occurred in the one-day cycle, whose advantage over the three-week reference decreased from 12.7 to 10.1 days when flooring was removed from the propagated state. The floor eradicated three of the 300 tumors that otherwise reached resistant disease; only ten of the 2,700 tumor–arm pairs changed. Thus, flooring modestly overstated the benefit of the most frequently segmented schedule.

In the 300-tumor subsample, five of the six contrasts against a block reference had intervals excluding zero. The exception was the one-day contrast, whose interval under the published implementation was -0.7 to 24.9 days. To determine whether its full-sample benefit persisted independently of flooring, we repeated the comparison across all 1,200 tumors with the floor were removed from the propagated state. The advantage was 10.8 days, with an interval of 4.5 – 17.2 days, compared with 12.3 days and an interval of 5.9 – 18.5 days under the published implementation. Both intervals excluded zero. The full-sample overstatement was therefore 1.5 days, smaller than the 2.6 days observed in the subsample. The one-week and three-day contrasts remained unchanged to the reported precision,

at 20.8 and 24.4 days under both implementations.

The seventh configuration, a ten-cell rather than one-cell floor, behaved differently. In the one-day cycle, floor activations increased from 53 to 84 per tumor, and the number classified as eradicated increased from 31 to 89 of 300 tumors, an increase of 58, compared with increases of 3–17 in the other arms. Its advantage over the three-week reference increased from 12.7 to 52.1 days, making it the best rather than the third-best arm in its block. The ordering within the one-half-duty block consequently reversed: the three-day cycle's 13.8-day advantage over the one-day cycle turned into a 15.0-day deficit, with an interval of -28.5 to -1.2 days.

The central finding that timing matters at fixed exposure persisted: all six contrasts against a block reference remained positive under every configuration. However, the nonmonotonic relationship between benefit and cycle period was sensitive to the ten-cell threshold. We interpreted that threshold as a change in the extinction assumption, because it removes populations containing up to ten cells rather than fewer than one. Its effect on schedule ranking was nevertheless important and is reported explicitly. Supplementary Table 7 provides every contrast under every configuration, together with eradication rates and floor activations for each arm.

3.11 Tumor burden and treatment-failure rankings

A natural way to choose between schedules is to score them by the tumor burden they leave

behind, a quantity clinicians measure directly. Comparing this ranking with the ranking by how often each schedule avoids the modeled failure state gives a rank correlation of 0.486, with a p value of 0.33. The rankings therefore differ, and their association is not statistically established.

The clearest example is the four-week cycle used by Das Thakur et al., which ranks fourth of six by tumor burden and last of six by failure probability. Its net effect is unfavorable when tumors are counted individually: it prevents failure in 71 tumors while causing failure in 76 that would otherwise avoid it. Its position differs between the two rankings, reinforcing the need to report these endpoints separately rather than treating either as a substitute for the other.

There was therefore no generally best schedule in these results; whenever we describe a schedule as best, we specify the endpoint. The three-day cycle performed best by delay to resistant disease and probability of failure. After accounting for censoring, it was also the only schedule whose time to progression was distinguishable from continuous dosing. The ranking by tumor burden, however, differed from both resistance-delay and failure-probability rankings.

3.12 Resistant-compartment ablation

To test the proposed mechanism, we repeated the comparison with the transition from the persister state to the resistant state switched off, removing durable resistance from the model. We specified the expected result before running the comparison.




Figure 11. Effect of removing the resistant compartment. The same 600 tumors were simulated with the transition from the persister state into the resistant state intact and then with that transition switched off, under both continuous and intermittent dosing. With the transition intact, 26.5 percent of tumors show any delay at all under intermittent dosing and the benefit is associated with drug dependence. With it switched off, no tumor benefits and the association disappears entirely, which establishes that the effect operates through the resistant population rather than through any other part of the model.

Of 600 tumors, 26.5% showed some delay from intermittent dosing in the full model, whereas none did after resistance was removed (Figure 11). This fraction counted any delay, rather than a gain exceeding seven days, whose prevalence interpretation we withdrew in Section 3.7. We report it as the baseline for comparison with the zero-benefit ablation result, rather than as an estimate of prevalence in melanoma.

The association between drug dependence and benefit also disappeared. The rank correlation changed from -0.166 , with a 95% interval of -0.247 to -0.086 obtained by resampling calibrated parameter sets, to 0.016 , with a p value of 0.70 . Without a resistant compartment, intermittent dosing shortened time to progression by an average of 136 days, because the holiday allowed sensitive-cell regrowth without an offsetting benefit from suppressing resistant cells.

3.13 Agreement with independent experimental datasets

The murine experiment was reproduced in direction but not in magnitude. Among tumors that progressed under continuous therapy within 100 days and whose resistant cells were drug-dependent—the situation described by Das Thakur et al.—the model reproduced the continuous-treatment outcome but not the intermittent-treatment outcome. At most, 11% of these tumors remained free of resistant disease beyond 200 days, compared with 100% in the murine experiment.

A mechanism contributing to this discrepancy is identifiable within the model: a drug holiday

also allows the sensitive population to regrow, and this population – in turn – seeds new resistance when dosing resumes. We cannot determine whether the mice avoided this outcome through immune clearance during the holiday, spatial constraints absent from the model, or slower tumor regrowth *in vivo* than represented in the swept range.

The model underestimates the persister fraction reported by Sharma et al. by ~ 44 -fold, producing a fraction of 6.7×10^{-5} compared with their 3×10^{-3} . Their measurement concerns lung cancer cells treated with an EGFR inhibitor rather than melanoma, so we treated it as an order-of-magnitude check rather than a calibration target. Nevertheless, the discrepancy remains substantial.

Comparison with the resistant-colony frequency reported by Shaffer et al. requires an important qualification. Their value of 5×10^{-4} counts colonies was obtained under a dose chosen to kill only a fraction of the cells. Most cells that did not form colonies were arrested rather than dead, so colony frequency is not equivalent to a surviving fraction. Our model reports the latter, and these quantities are therefore not directly comparable. With this qualification, the model's median surviving fraction at 28 days is 0.59 , and its lowest value across all tumors is 0.04 ; ~ 80 times the reported colony frequency.

4. Discussion

4.1 Biological interpretation of scheduling benefit

Treatment scheduling is beneficial only under

specific biological conditions. Two requirements must hold together. First, tumor growth must be sufficiently slow that sensitive cells do not rebound substantially during a drug holiday and resistance is not already established before the first break. Second, resistant cells must lose fitness when the drug is withdrawn. Neither requirement alone is sufficient. We do not estimate the proportion of tumors that meet these conditions because the size of the schedule-sensitive region depends on assumptions about unmeasured population parameters: across the four sampling configurations, it varied from 9.3% to 56.0%.

Identifying these conditions provides more useful guidance than ranking schedules alone. It indicates when a drug holiday is unlikely to help and offers a testable basis for evaluating its potential benefit. Measuring whether resistant cells regress during treatment withdrawal, together with how rapidly sensitive cells regrow, could guide the selection of schedules for experimental evaluation.

4.2 Resistance delay versus overall progression

The distinction between delaying resistance and delaying overall progression offers a possible explanation for the contrasting findings of Das Thakur et al. and Algazi et al. Intermittent dosing prevented lethal drug-resistant disease in the murine study (12), whereas it worsened progression-free survival in the human trial (13). The model illustrates how these outcomes can diverge: treatment withdrawal can suppress drug-dependent resistant cells while permitting sensitive cells to regrow, so a delay in resistance need not translate into better tumor control.

This interpretation remains a hypothesis rather than a reproduction of either experiment. The model does not reproduce the clinical disadvantage under the specific S1320 schedule, and its divergence between resistance and progression outcomes occurs under shorter cycles. Examining whether resistance emerged later in the intermittent arm of S1320 despite earlier overall progression would test the proposed explanation. However, the model omits immune responses, spatial constraints, and species-specific pharmacokinetics; it also simulates a single degrader rather than the trial's two-drug regimen following an eight-week lead-in. These differences limit direct comparison with the experimental and clinical outcomes.

4.3 Interpretation of calibration and parameter uncertainty

The available data meaningfully constrain only 11 of the 33 fitted parameters, indicating that much of the model remains underdetermined. We therefore assess parameter identifiability directly and report the ratio of effective constraints to parameters, rather than interpreting narrow posterior intervals as evidence of precision. Such underdetermination is common in systems biology: Gutenkunst et al. found that parameter sensitivities span many orders of magnitude across 17 models, with only a few combinations of parameters well determined (33).

The measurements primarily constrain signaling and drug action, while the cell-population dynamics governing treatment outcomes remain unmeasured. Population parameters are therefore varied across plausible

ranges rather than estimated from the calibration data. Consequently, the study identifies the biological conditions under which scheduling may help; it does not establish how frequently those conditions occur in melanoma or predict outcomes for individual patients.

4.4 Population heterogeneity and model discrepancies

The unresolved discrepancies reflect a structural limitation: the model represents average population behavior and cannot capture subsets of cells with distinct signaling activity or treatment responses. Shaffer et al. observed nine resistant colonies among 18,000 cells (6), while Gerosa et al. found near-normal ERK activity in 1–5% of cells at any moment and signaling pulses in approximately 16% over 16 hours (8). Yang et al. likewise found that a tenfold dose increase reduced cellular escape from 43% to 10% without eliminating it (7). These observations depend on variation between cells that a single ERK activity value for each population cannot represent.

This limitation also creates a quantitative conflict. Reproducing Yang et al.'s 43% escape rate at 1 μM while average pathway activity is suppressed by more than 99% requires a proliferation threshold so low that growth becomes effectively independent of MAPK activity. Such an adjustment would undermine the mechanism through which pathway inhibition arrests proliferation. Both escape measurements are therefore reported but excluded from fitting; reproducing them would require a model that explicitly represents cellular heterogeneity.

4.5 Limitations

The model represents cell populations by their average behavior, limiting its ability to reproduce the rare-cell responses discussed in Section 4.4. These structural discrepancies cannot be resolved through parameter adjustment alone; incorporating cell-to-cell variability and stochastic transitions into resistance may help address them. The preexisting route to resistance that Hata et al. describe is not exercised here, since every simulation starts with no resistant cells, so these results speak only to resistance that emerges through the persister state. The resistant population is also represented by a single state and growth rate, although distinct resistance mechanisms may respond differently to treatment withdrawal (4).

Calibration relies on cultured A375 cells, whereas the scheduling experiments involve animals and patients with different growth dynamics. Sweeping growth rates avoids imposing an unsupported conversion factor, but absolute time estimates remain conditional on the selected growth regime, and paired schedule comparisons do not establish clinical transferability. The four hypothetical tumor phenotypes reflect the direction of reported biological differences rather than measured parameter magnitudes, so phenotype-specific conclusions remain illustrative (27, 28). Finally, persistence of the adapted state is imposed as a calibration assumption rather than established experimentally, because the measurements of Gerosa et al. (8) end at 24 hours.

4.6 Experimental validation and model development

The most informative next experimental step would be to measure the growth or regression of resistant cells after treatment withdrawal, because their loss of fitness is a key determinant of whether a drug holiday benefits a given tumor. Das Thakur et al. demonstrated the feasibility of this measurement in mice (12).

Computationally, introducing stochastic transitions into resistance could help address the limitations discussed in Section 4.4 and enable a more direct comparison with Shaffer's resistant-colony counts. The framework could also be extended to other drug combinations by modifying their pharmacodynamic parameters and molecular targets, with changes to the model structure where the treatments engage additional pathways or resistance mechanisms.

4.7 Perspectives: a template for tumor-specific scheduling models

The broader contribution of this study is a framework for asking when treatment timing matters. Its goal is to identify schedules that kill sensitive tumor cells while limiting the emergence and expansion of resistant cells. These aims can compete: a treatment break may suppress drug dependent resistant cells while allowing sensitive cells to regrow. By comparing schedules at equal cumulative exposure and examining resistance and overall progression separately, the model shows why delaying resistance alone does not establish that a schedule controls the tumor better.

The framework could serve as a model template for other tumor types and treatments. Initially, applying it would require experiments that measure how sensitive and resistant populations grow during treatment, respond to withdrawal, and behave when treatment resumes. As evidence accumulates, standard parameter ranges may emerge for defined tumor types and treatments. Researchers could then use those ranges for preliminary modelling without measuring every parameter anew, test whether conclusions hold across plausible ranges, and validate important predictions with independent data.

A longer term goal would be to build models standardized by tumor type that identify schedules balancing sensitive cell killing against resistance control. Common definitions of cell populations, treatment exposure, and outcomes would make findings easier to compare across studies. The present study offers a starting framework for that work; its schedule rankings cannot yet be assumed to transfer to other tumors or patients.

5. Conclusion

This study provides a framework for determining when treatment scheduling can preserve sensitive-cell killing while limiting the emergence and expansion of resistant populations. The calibrated BRAF V600E melanoma model connects intracellular signaling and drug action to the dynamics of sensitive, temporarily drug-tolerant, and resistant cells. Comparisons at equal cumulative drug exposure show that timing itself influences resistance outcomes, while removing the resistant compartment eliminates

the scheduling benefit. Together, these analyses identify resistant-cell growth and loss of fitness after treatment withdrawal as key conditions governing the value of drug holidays.

Treatment scheduling must balance suppression of drug-dependent resistant cells against the consequences of treatment breaks. Not only can sensitive cells regrow during a drug holiday, but they can also enter the persister state when treatment resumes and subsequently replenish the resistant population, potentially offsetting the benefit of withdrawal. This balance determines whether delaying resistance improves overall tumor control and explains why shorter cycles need not be better. It also provides a framework for interpreting conflicting experimental and clinical findings. Schedules should therefore be evaluated according to the responses of both cell populations, with resistance delay and overall progression considered together.

The model offers a methodological foundation for studies of other tumor types and treatments, provided that their signaling pathways, drug responses, and population dynamics are appropriately represented and independently validated. Experiments measuring how sensitive and resistant populations respond over time to treatment, withdrawal, and treatment resumption could establish standardized

parameter ranges for defined tumor types and treatments, progressively reducing the need to measure every parameter anew. Although uncertain population parameters and the absence of cell-to-cell variability limit current predictions, the framework establishes a route toward tumor-specific scheduling models that balance sensitive-cell killing against resistance control.

Acknowledgments

The author thanks Nicole Spinelli and James Truglio of the William A. Shine Great Neck South High School science research program, who supervised this project and provided sustained guidance, mentorship, and feedback on the research design and on successive drafts of this manuscript. Author contributions: the author conceived the study, designed and implemented the computational model, carried out the parameter calibration, uncertainty quantification, and all analyses, produced all figures, and wrote the manuscript. Nicole Spinelli and James Truglio provided project supervision and editorial feedback and did not carry out the computational work. The author used AI based tools (Anthropic's Claude) for two purposes, which were retrieving and locating published literature and editing some of the writing in this manuscript. The author reviewed and verified all analyses and all text presented here.

Data and code availability

The complete model source, every analysis script, and every result table underlying the figures and tables in this manuscript are provided as a supplementary archive accompanying this submission. The model specification is given in full in the Supplementary Methods, and the calibration targets with their provenance are in Supplementary Table 2. Each analysis reported here can be reproduced by running the corresponding script against the shipped tables. One

intermediate file, the full record of accepted and rejected candidates from the sequential calibration, is omitted from the archive for size and is available from the author on request; no result reported here depends on it.

6. References

1. Davies, H., Bignell, G. R., Cox, C., Stephens, P., Edkins, S., Clegg, S., et al. (2002). Mutations of the BRAF gene in human cancer. *Nature*, 417(6892), 949-954. <https://doi.org/10.1038/nature00766>
2. Flaherty, K. T., Puzanov, I., Kim, K. B., Ribas, A., McArthur, G. A., Sosman, J. A., et al. (2010). Inhibition of mutated, activated BRAF in metastatic melanoma. *New England Journal of Medicine*, 363(9), 809-819. <https://doi.org/10.1056/NEJMoa1002011>
3. Larkin, J., Ascierto, P. A., Dréno, B., Atkinson, V., Liskay, G., Maio, M., et al. (2014). Combined vemurafenib and cobimetinib in BRAF-mutated melanoma. *New England Journal of Medicine*, 371(20), 1867-1876. <https://doi.org/10.1056/NEJMoa1408868>
4. Villanueva, J., Vultur, A., Lee, J. T., Somasundaram, R., Fukunaga-Kalabis, M., Cipolla, A. K., et al. (2010). Acquired resistance to BRAF inhibitors mediated by a RAF kinase switch in melanoma can be overcome by cotargeting MEK and IGF-1R/PI3K. *Cancer Cell*, 18(6), 683-695. <https://doi.org/10.1016/j.ccr.2010.11.023>
5. Sharma, S. V., Lee, D. Y., Li, B., Quinlan, M. P., Takahashi, F., Maheswaran, S., et al. (2010). A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1), 69-80. <https://doi.org/10.1016/j.cell.2010.02.027>
6. Shaffer, S. M., Dunagin, M. C., Torborg, S. R., Torre, E. A., Emert, B., Krepler, C., et al. (2017). Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658), 431-435. <https://doi.org/10.1038/nature22794>
7. Yang, C., Tian, C., Hoffman, T. E., Jacobsen, N. K., & Spencer, S. L. (2021). Melanoma subpopulations that rapidly escape MAPK pathway inhibition incur DNA damage and rely on stress signalling. *Nature Communications*, 12, 1747. <https://doi.org/10.1038/s41467-021-21549-x>
8. Gerosa, L., Chidley, C., Fröhlich, F., Sanchez, G., Lim, S. K., Muhlich, J., et al. (2020). Receptor-driven ERK pulses reconfigure MAPK signaling and enable persistence of drug-

adapted BRAF-mutant melanoma cells. *Cell Systems*, 11(5), 478-494.e9.

<https://doi.org/10.1016/j.cels.2020.10.002>

9. Rambow, F., Rogiers, A., Marin-Bejar, O., Aibar, S., Femel, J., Dewaele, M., et al. (2018). Toward minimal residual disease-directed therapy in melanoma. *Cell*, 174(4), 843-855.e19.

<https://doi.org/10.1016/j.cell.2018.06.025>

10. Marine, J. C., Dawson, S. J., & Dawson, M. A. (2020). Non-genetic mechanisms of therapeutic resistance in cancer. *Nature Reviews Cancer*, 20(12), 743-756.

<https://doi.org/10.1038/s41568-020-00302-4>

11. Hata, A. N., Niederst, M. J., Archibald, H. L., Gomez-Caraballo, M., Siddiqui, F. M., Mulvey, H. E., et al. (2016). Tumor cells can follow distinct evolutionary paths to become resistant to epidermal growth factor receptor inhibition. *Nature Medicine*, 22(3), 262-269.

<https://doi.org/10.1038/nm.4040>

12. Das Thakur, M., Salangsang, F., Landman, A. S., Sellers, W. R., Pryer, N. K., Levesque, M. P., et al. (2013). Modelling vemurafenib resistance in melanoma reveals a strategy to forestall drug resistance. *Nature*, 494(7436), 251-255. <https://doi.org/10.1038/nature11814>

13. Algazi, A. P., Othus, M., Daud, A. I., Lo, R. S., Mehnert, J. M., Truong, T. G., et al. (2020). Continuous versus intermittent BRAF and MEK inhibition in patients with BRAF-mutated melanoma: a randomized phase 2 trial. *Nature Medicine*, 26(10), 1564-1568.

<https://doi.org/10.1038/s41591-020-1060-8>

14. Gatenby, R. A., Silva, A. S., Gillies, R. J., & Frieden, B. R. (2009). Adaptive therapy. *Cancer Research*, 69(11), 4894-4903. <https://doi.org/10.1158/0008-5472.CAN-08-3658>

15. Zhang, J., Cunningham, J. J., Brown, J. S., & Gatenby, R. A. (2017). Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature Communications*, 8(1), 1816. <https://doi.org/10.1038/s41467-017-01968-5>

16. Strobl, M. A. R., West, J., Viossat, Y., Damaghi, M., Robertson-Tessi, M., Brown, J. S., et al. (2021). Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Research*, 81(4), 1135-1147. <https://doi.org/10.1158/0008-5472.CAN-20-0806>

17. Yao, Z., Torres, N. M., Tao, A., Gao, Y., Luo, L., Li, Q., et al. (2015). BRAF mutants evade ERK-dependent feedback by different mechanisms that determine their sensitivity to

pharmacologic inhibition. *Cancer Cell*, 28(3), 370-383.

<https://doi.org/10.1016/j.ccell.2015.08.001>

18. Poulikakos, P. I., Zhang, C., Bollag, G., Shokat, K. M., & Rosen, N. (2010). RAF inhibitors transactivate RAF dimers and ERK signalling in cells with wild-type BRAF. *Nature*, 464(7287), 427-430. <https://doi.org/10.1038/nature08902>

19. Huang, C. Y., & Ferrell, J. E. (1996). Ultrasensitivity in the mitogen-activated protein kinase cascade. *Proceedings of the National Academy of Sciences*, 93(19), 10078-10083. <https://doi.org/10.1073/pnas.93.19.10078>

20. Bollag, G., Hirth, P., Tsai, J., Zhang, J., Ibrahim, P. N., Cho, H., et al. (2010). Clinical efficacy of a RAF inhibitor needs broad target blockade in BRAF-mutant melanoma. *Nature*, 467(7315), 596-599. <https://doi.org/10.1038/nature09454>

21. Lito, P., Pratilas, C. A., Joseph, E. W., Tadi, M., Halilovic, E., Zubrowski, M., et al. (2012). Relief of profound feedback inhibition of mitogenic signaling by RAF inhibitors attenuates their activity in BRAFV600E melanomas. *Cancer Cell*, 22(5), 668-682. <https://doi.org/10.1016/j.ccr.2012.10.009>

22. Alabi, S., Jaime-Figueroa, S., Yao, Z., Gao, Y., Hines, J., Samarasinghe, K. T. G., et al. (2021). Mutant-selective degradation by BRAF-targeting PROTACs. *Nature Communications*, 12(1), 920. <https://doi.org/10.1038/s41467-021-21159-7>

23. Gadd, M. S., Testa, A., Lucas, X., Chan, K. H., Chen, W., Lamont, D. J., et al. (2017). Structural basis of PROTAC cooperative recognition for selective protein degradation. *Nature Chemical Biology*, 13(5), 514-521. <https://doi.org/10.1038/nchembio.2329>

24. Oren, Y., Tsabar, M., Cuoco, M. S., Amir-Zilberstein, L., Cabanos, H. F., Hütter, J. C., et al. (2021). Cycling cancer persister cells arise from lineages with distinct programs. *Nature*, 596(7873), 576-582. <https://doi.org/10.1038/s41586-021-03796-6>

25. Fallahi-Sichani, M., Honarnejad, S., Heiser, L. M., Gray, J. W., & Sorger, P. K. (2013). Metrics other than potency reveal systematic variation in responses to cancer drugs. *Nature Chemical Biology*, 9(11), 708-714. <https://doi.org/10.1038/nchembio.1337>

26. Toni, T., Welch, D., Strelkowa, N., Ipsen, A., & Stumpf, M. P. H. (2009). Approximate Bayesian computation scheme for parameter inference and model selection in dynamical

systems. *Journal of the Royal Society Interface*, 6(31), 187-202.

<https://doi.org/10.1098/rsif.2008.0172>

27. Tsoi, J., Robert, L., Paraiso, K., Galvan, C., Sheu, K. M., Lay, J., et al. (2018). Multi-stage differentiation defines melanoma subtypes with differential vulnerability to drug-induced iron-dependent oxidative stress. *Cancer Cell*, 33(5), 890-904.e5.

<https://doi.org/10.1016/j.ccell.2018.03.017>

28. Verfaillie, A., Imrichová, H., Atak, Z. K., Dewaele, M., Rambow, F., Hulselmans, G., et al. (2015). Decoding the regulatory landscape of melanoma reveals TEADS as regulators of the invasive cell state. *Nature Communications*, 6, 6683. <https://doi.org/10.1038/ncomms7683>

29. Spear, R. C., & Hornberger, G. M. (1980). Eutrophication in Peel Inlet, II. Identification of critical uncertainties via generalized sensitivity analysis. *Water Research*, 14(1), 43-49.

[https://doi.org/10.1016/0043-1354\(80\)90040-8](https://doi.org/10.1016/0043-1354(80)90040-8)

30. Pianosi, F., & Wagener, T. (2015). A simple and efficient method for global sensitivity analysis based on cumulative distribution functions. *Environmental Modelling & Software*, 67, 1-11.

<https://doi.org/10.1016/j.envsoft.2015.01.004>

31. Marino, S., Hogue, I. B., Ray, C. J., & Kirschner, D. E. (2008). A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254(1), 178-196. <https://doi.org/10.1016/j.jtbi.2008.04.011>

32. Peduzzi, P., Concato, J., Kemper, E., Holford, T. R., & Feinstein, A. R. (1996). A simulation study of the number of events per variable in logistic regression analysis. *Journal of Clinical Epidemiology*, 49(12), 1373-1379. [https://doi.org/10.1016/s0895-4356\(96\)00236-3](https://doi.org/10.1016/s0895-4356(96)00236-3)

33. Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R., & Sethna, J. P. (2007). Universally sloppy parameter sensitivities in systems biology models. *PLoS Computational Biology*, 3(10), e189. <https://doi.org/10.1371/journal.pcbi.0030189>