

## **Adaptive therapy for heterogeneous cancer: exploiting space and trade-offs in drug scheduling**

Jill Gallaher<sup>1</sup>, Pedro M. Enriquez-Navas<sup>2</sup>, Kimberly A Luddy<sup>2</sup>, Robert A. Gatenby<sup>1</sup>, and Alexander R. A. Anderson<sup>1</sup>

1H. Lee Moffitt Cancer Center, Integrated Mathematical Oncology (Tampa, FL, United States)

2H. Lee Moffitt Cancer Center, Cancer Imaging and Metabolism (Tampa, FL, United States)

### **Abstract**

Cancer treatment has greatly benefited from the introduction of both new agents (i.e. targeted therapy and check point inhibitors) and new strategies for conventional therapies such as chemotherapy and hormonal therapy. Most cancer types now have at least one effective treatment agent, and many tumors, such as breast and prostate cancer, have multiple available treatment options. However, even the most effective therapy is usually defeated as cancer cells deploy a wide range of molecular strategies to overcome therapy leading to disease progression. Here we propose that, while molecular dynamics govern response and resistance to therapy, evolutionary dynamics determine survival and proliferation of treatment-resistant cells. We hypothesize that understanding these evolutionary interactions may identify strategies to delay or prevent proliferation of resistant population thus prolong time to recurrence.

To simulate these interactions, we use an off-lattice, agent-based, framework to model competition between sensitive and resistant populations during therapy in a spatially competitive environment. Our model applies a classic evolutionary trade-off between fecundity (cellular proliferation) and survivorship (drug sensitivity) to the tumor populations. Thus, in the resource-limited tumor microenvironment, the cost of increased resource investment in a resistant

phenotype necessarily results in slower proliferation while a phenotype invested in fast proliferation is less likely to survive in drug.

Model simulations demonstrate that, in the absence of therapy, cells with slower growth but higher survivorship may become confined in the interior of the tumor during growth phases, which provides a spatial sanctuary during initial drug administration. Over time, therapy eliminates the treatment-sensitive population allowing the resistant cells to proliferate unopposed. We simulate the application of an anti-proliferative drug on varying ratios of mixed sensitive and resistant cells using two general treatment strategies: a continuous schedule of maximum tolerated dose or an evolution-informed schedule that incorporates dose-modulation and treatment vacations to sustain control of the tumor through competition between sensitive and resistant cell populations. We find tumors consisting only of sensitive cells can be cured with continuous treatment, but the presence of any significant population of resistant cells will lead to eventual recurrence. We identify two treatment strategies that control heterogeneous tumors: one emphasizes continuous dose modulation, and the other relies on treatment vacations. Both strategies control tumors over a wide range of resistant/sensitive population ratios but the average dose given is significantly lower with dose modulation. Cell migration and phenotypic drift limit the time to recurrence advantage of these strategies but some efficacy can be retained through a more vacation-oriented schedule.

## **1. Introduction**

Despite major advances in cancer therapies, most metastatic cancers remain fatal because tumor cells have remarkable capacity to evolve drug resistance leading to treatment failure and tumor progression. Most investigations of cancer treatment resistance have focused on identifying and targeting the molecular mechanisms that confer resistance. However, this has typically not led to substantial improvements in patient outcomes in large part because the cancer cells have access

to the vast information content of the human genome so that defeat of one resistance strategy simply results in deployment of another (1).

An alternative approach focuses on the population-level dynamics that govern proliferation of resistant tumor cells. While molecular investigations can define mechanisms of resistance, it is proliferation of resistant tumor cells that is ultimately responsible for tumor progression. And, this is governed by Darwinian evolutionary principles that define the fitness of each cell within the local environmental context. For example, cancer cells may become resistant to treatment through increased expression of efflux pumps. However, the synthesis, maintenance, and operation of these pumps require considerable investment of resources (up to 50% of the cell's total energy budget) (2). In the harsh microenvironment commonly found in cancers, this investment in survival will likely require diversion of resources that would ordinarily be devoted to invasion or proliferation. Thus, while a tumor cells may possess the molecular mechanisms necessary for therapy resistance, it is by no means certain that they will be able to proliferate at a rate sufficient to repopulate the tumor after treatment. In fact, proliferation of resistant cells is governed by complex interactions that include the cost/benefit ratio of the resistance mechanism(s) and competition with other tumor subpopulations.

Despite the growing recognition of the importance of heterogeneity and evolution in driving treatment failure, explicit inclusion of Darwinian principles in clinical trial design is rare (3-6). Here we examine intratumoral heterogeneity and the evolutionary dynamics that result from differences in the relative fitness of cells that are resistant or sensitive to therapy in the presence and absence of that therapy.

A common maxim in treatment of cancer is to "hit hard and fast" through maximum dose density strategies that administer the highest possible drug dose in the shortest possible time period. The rationale is that the greatest patient benefit is achieved through killing the maximum

possible number of cancer cells and preventing resistance before it occurs. As a result, initial clinical evaluation of a cancer drug (Phase 1 trial) always seeks to identify the maximum tolerated dose (MTD) – that is, the maximum administered drug dose that does not result in unacceptable toxicity. An evolutionary rationale for MTD is that rapid tumor cell killing reduces time for transition to a resistant phenotype.

The MTD principle has been standard of care for cancer treatment for several decades. It has not, however, resulted in consistent cures in patients with most disseminated cancers despite a number of efforts (7), such as an increase in the doses or number of drugs administered. An evolutionary flaw in this strategy is the assumption that resistant populations are not present prior to therapy. It is now clear that cancer cells can be insensitive even to treatments that they have never “seen” before through a variety of direct and indirect mechanisms. Because of this, MTD therapy designed to kill as many cancer cells as possible, although intuitively appealing, may be evolutionarily unwise. This is because of a well-recognized Darwinian dynamic from ecology termed “competitive release” observed, for example, when high doses of pesticide are applied for pest eradication (8). Because maximum application of pesticides applies strong selection pressure for resistance it eliminates all competing populations. Therefore, this strategy accelerates proliferation of resistant phenotypes and treatment failure.

Here we hypothesize that, while *selection* for resistance through application of cytotoxic therapy is inevitable, *proliferation* of those cells may be controlled using evolutionary principles. Importantly, multiple experimental models have shown that drug-resistant cancer cells proliferate slower than sensitive cells in the absence of drug (9-13). This is because resistant cancers, like most drug-resistant bacteria, incur a fitness cost that expend energy on the mechanisms involved (14). Here, we computationally investigate the hypothesis that these costs can be exploited to delay or prevent proliferation of resistance in cancer.

## 2. Methods

To investigate intratumoral evolutionary dynamics during treatment, we modify an off-lattice agent-based model built to investigate trait selection via spatial competition. We focus on the role of initial distributions of proliferation rates on both overall tumor growth and spatial distribution of subpopulations. We then examine outcomes from varying treatment strategies with an anti-proliferative drug focusing on progression of resistant populations and treatment failure.

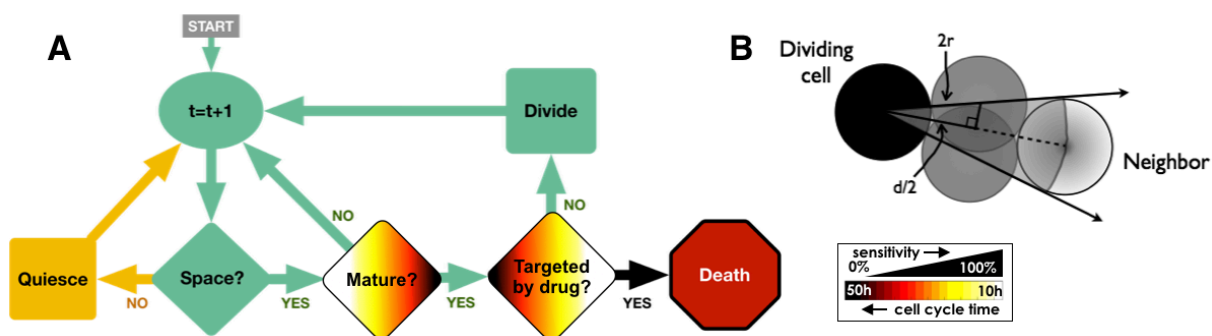


Figure 1. *Model rules. A) The flow chart details decision-making of a cell over each time step. Proliferation steps are shown in green, quiescence in yellow, and drug-induced death in red. The decisions that are multicolored reflect the phenotypic heterogeneity in cell cycle times and sensitivity to drug. Key shows the trade-off between cell cycle time and drug sensitivity. B) The off-lattice model assumes space limited proliferation, such that a dividing cell with insufficient space for non-overlapping offspring enters a quiescent state.*

The flowchart in **Figure 1A** shows the decisions made for each cell at each time point starting at the gray arrow. Green arrows and nodes represent proliferation decisions; yellow, quiescence; multicolored, decisions affected by heterogeneity; and red, exit of a cell due to drug-induced death. In this off-lattice model, we assume that cells take up space so that at carrying capacity, they enter quiescence due to contact inhibition. Space is deemed available if any contiguous set of integer angles are empty and sufficient to allow a cell to fit without overlap (described previously in (15)). When a cell is quiescent, we assume that the cell discontinues decision-making until space is available, i.e. it cannot be affected by drug and it cannot proliferate. We also assume that there is no cell death due to regular turnover and that resources are abundant when cells are not suppressed by lack of space.

Based on a number of experimental observations, we assume that the cost of resistance results in a decrease in the cell's proliferation rate (9-13). We assume that the fastest proliferation rate has 100% sensitivity, and in a linear fashion, the slowest proliferation rate has 0% sensitivity. This provides a simple trade-off between cell fitness without drug and its fitness when the drug is present.

Despite the treatment strategy, a cell's response to drug exposure depends on its sensitivity and the dose applied. We define the drug sensitivity, due to the imposed trade-off with the ability to reproduce, as a decreasing linear function of cell cycle time:

$$s(T) = \frac{T_{max} - T}{T_{max} - T_{min}}, \quad (1)$$

where  $T$  is the time it takes to go through a cell cycle, and  $T_{min}$  and  $T_{max}$  are the minimum and maximum allowed cell cycle times, which are 10 and 50 hours, respectively. We represent drug efficacy as a typical Hill function for simplicity and generality (16,17):

$$E(D) = \frac{D^n}{D^n + K_{0.5}^n}, \quad (2)$$

where  $n=1.5$  and  $K_{0.5}=0.125$ . These values are not calibrated to any specific half maximal activity, but generically assume a response function that gives nearly 100% probability of death at the highest drug concentrations, dropping off slowly at mid ranges and then quicker as it gets to lower concentrations. The probability of death,  $P_{death}$ , combines the sensitivity and dose effects:

$$P_{death}(T, D) = s(T)E(D) = \frac{s(T)D^n}{D^n + K_{0.5}^n}. \quad (3)$$

We explore treatment strategies other than conventional MTD, which we term “adaptive therapy.” In general, we vary the dose  $D$  from its previous value  $D_0$  based on the relationship of

the current number of cells  $N$  to the previous number of cells  $N_0$  according to the following equation:

$$D(\alpha, \beta) = \begin{cases} 0 & \text{if } N < 0.5N_0 \\ (1 + \alpha)D_0 & \text{if } N > (1 + \beta)N_0 \\ (1 - \alpha)D_0 & \text{if } N \leq (1 - \beta)N_0 \\ D_0 & \text{otherwise} \end{cases}, \quad (4)$$

where  $\alpha$  and  $\beta$  are fixed values that can affect responsiveness of the change in dose to the change in population size. In **Eqn. 4**, if the number of cells is below half of the original, a treatment vacation is enforced. If the change in number of cells is within a specified window the dose remains the same as before. Otherwise, there will be a fractional increase in the dose by  $\alpha$  if the fractional increase in population size is greater than  $\beta$  and a fractional decrease in dose by  $\alpha$  if the fractional decrease in the population size is greater than  $\beta$ . The dose begins at the MTD, and never exceeds the MTD and each new dose is determined every 3 days in accordance with the experiments reported by Enriquez-Navas et al (18). If a cell is targeted by the drug, there is a 15-30h delay, randomly chosen, before it is removed from the domain, to account for the time it takes for apoptosis to occur (19,20).

### 3. Results

#### **Fitness differences and space limitation affect competition and selection**

The cost of resistance and the selection force of space limitation can be demonstrated through a combination of experimental and mathematical models. The MCF7Dox cell line was selected for overexpressing MDR1 and its resistance to Doxorubicin. These cells were labeled with red fluorescence protein (RFP). The MCF7Dox cells are plated with the MCF7 sensitive parental line (labeled with green fluorescent protein, GFP) at a 1:1 ratio to observe the relative dynamics of growth. The cells are passaged and measured by fluorescence intensity every 3 days for 46 days (**Fig. 2A**). The sensitive cells are seen to dominate the population over time, which demonstrates

that the resistant cells pay a proliferative cost at expense of developing the resistance mechanisms. Critically, once the sensitive cells take over a significant portion of the population, they suppress further outgrowth of resistant cells.




Figure 2. *Fitness differences and space limitation affect growth and selection. A) Sensitive (MCF7) and resistant (MCF7Dox) cells to doxorubicin are grown together in vitro and demonstrate the fitness cost of resistance. The plot shows over time the percent of drug-sensitive and drug-resistant cells. B) Model population dynamics showing fraction of sensitive phenotype (40h cell cycle) and resistant (60h cell cycle) cells at a ratio of 1:1 with 100 cells and grown up to 20,000 cells. In one example, the cells were initially spread sparsely (dish), and in the other, cells are initially seeded in a tight cluster configuration. C) The spatial distribution for each example is shown at two time points,  $t_1=7$  days and  $t_2=30$  days, and for the cluster also at  $t_3=90$  days. A zoomed-in view shows the distribution of proliferating (yellow) and quiescent (green if sensitive, red if resistant) cell types.*

We can recreate a similar scenario in our simple model system (**Fig. 2B**) by just assuming that the population mix consists of a faster growing sensitive phenotype and a slower growing resistant phenotype. We use the average growth rates found at the beginning of the experiment, which was a 40 hour cycle time for MCF7 and a 60 hour cycle time for MCF7Dox. We initially seeded 100 cells scattered throughout a circular domain with a 1.5 mm radius to grow distinct colonies, reflecting the distribution in the *in vitro* model. We grow the population to 20,000 cells and find that, over time, the sensitive cells quickly take over the space and prevent further proliferation of the resistant colonies. However, the *in vitro* environment with ample space and nutrients is distinct from the densely populated and resource depleted tumor environment. In solid tumors, the cells are less free to move and proliferate, which adds additional selection pressure on the cell populations. If we initialize the simulation with cells arranged in a more confined solid mass configuration but with the same conditions, we find that there is a significant



change in the population dynamics; the limited space causes much more competition between different populations and takes a much greater amount of time to grow to the same population size. The spatial layout (Fig. 2C) shows that when the cells can grow as separate colonies, the “tumor edge” is large and convoluted. With the cells arranged in a dense cluster, there is greater selection at the boundary, and the sensitive cells quickly take over the expanding front. Because proliferation is limited to the edge of the mass, the population grows much slower.

### **Phenotypic spread reduces coexistence during growth**

Sensitivity to drug is often made simpler than the situation in reality, being thought of in terms of binary states where a cell is simply sensitive or resistant to a drug, and it would respond with either death or survival. However, real tumors have a more nuanced mix of phenotypes that differ in their sensitivity to a drug. Here we investigate how various mixtures of sensitive and resistant cells compete during growth.

We begin with 100 cells with phenotypes drawn from a normal distribution of drug sensitivities (each tumor has a unique mean cell cycle time  $\tau$  and a standard deviation  $\sigma_\tau$ ) and stop the simulation when the number of cells reach 15,000, which is around 1.5 mm in diameter, representing a micrometastasis. The initial and final population distributions are shown in Figure 3 for each case, with the percent makeup of the dominant phenotype labeled. We consider an array of initial distributions with increasing mean and standard deviation in the cell cycle times (corresponding to decreasing sensitivities). The difference between the initial and final compositions can tell us something about the competition between the phenotypes. Toward the top of the array, where there is little heterogeneity, the compositions remain rather homogeneous, i.e. the sensitive tumors remain mostly sensitive and the resistant tumors retain resistance. However, with more heterogeneity in the initial compositions, the more sensitive cells are seen to outcompete the more resistant cells. In many cases, the competition is so great that

the final proportion of resistant cells is imperceptibly small. **Figure 3A-D** shows the final spatial configuration of some of the more extreme cases. They all show that the more resistant cells get trapped in the interior of the tumor by the more sensitive cells, which take over the invading edge. More coexistence between phenotypes is seen in B and D, but given enough time the more proliferative phenotypes that are bulging forward at the edges should eventually take over. Overall, the trend shows that with a given distribution of phenotypes, the faster proliferators will take over the population, but the timing will depend on the relative fitness difference between the competitors.

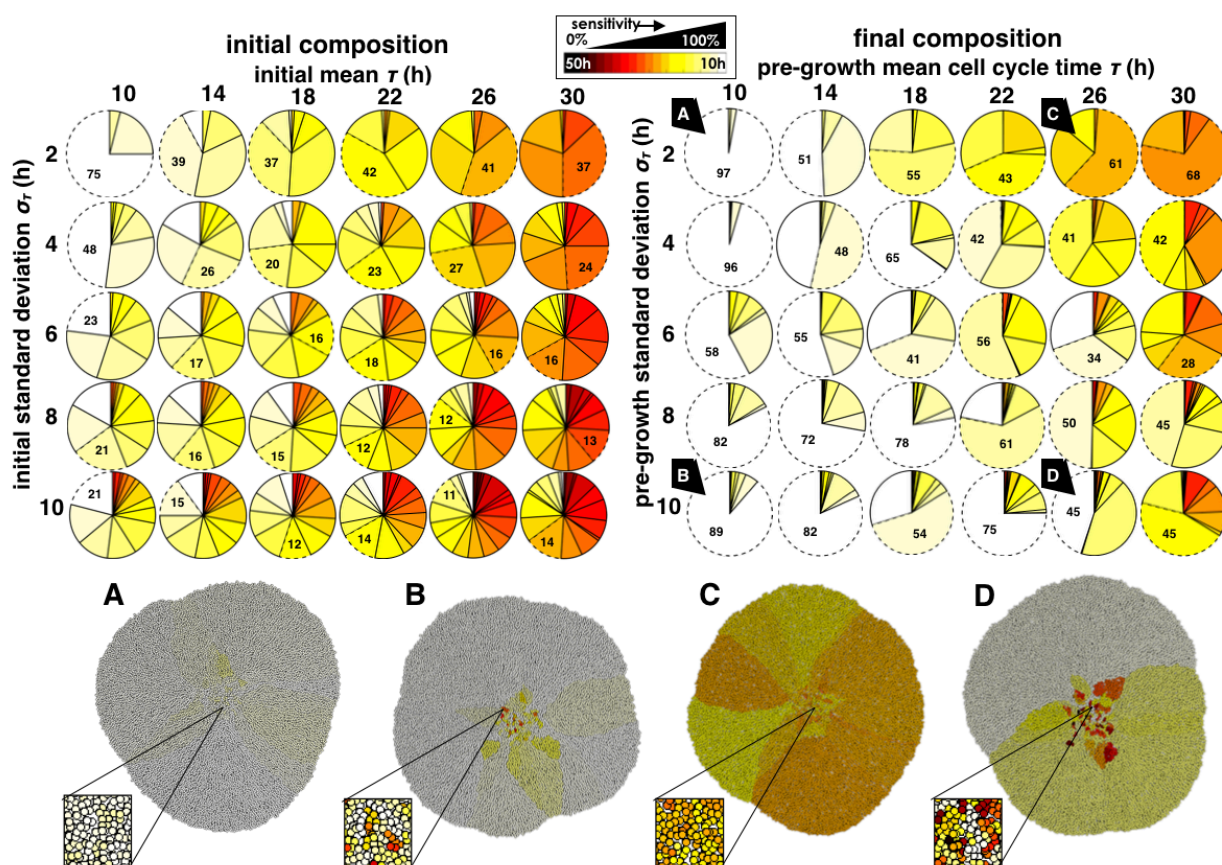


Figure 3. Growing tumors from many different initial compositions of cell phenotypes. We display in pie charts the fraction of the population occupied by cells with different proliferation rates at the beginning (left) and end (right) of the simulation. In each, the initial distribution is defined by a mean cycle time  $\tau$  and a standard deviation  $\sigma_r$ . A pie piece with a dashed border and a number shows the dominant clone and the percent occupied at that time point. The color key gives the cycle time and sensitivity over the range of colors. Final spatial distributions for selected cases A-D are displayed below.

## Continuous treatments can cure some tumors; adaptive treatments can control most tumors

Treatment is applied to the previously grown tumors (number original of cells  $N_0=15,000$ ) and stops when either the tumor is cured (i.e. the number of current cells  $N=0$ ), recurs (i.e.  $N=4/3N_0=20,000$  cells), or the tumor reaches an age of 2 years post-treatment. We assume tumor control if the final number of cells is below 10,000, which is half the value that determines recurrence and accounts for fluctuations during adaptive treatments. We do not simulate any short-term pharmacokinetics, so for the continuous therapy, we simply apply the maximum tolerated dose (MTD) the entire duration and each dose change in the adaptive strategy is instantaneously realized. Only when a cell is capable of dividing does it become sensitive to drug toxicity and die with the given probability in **Eqn. 3**. Further, we define several metrics to describe adaptive treatments: the average dose,  $D_{ave}$ , gives the average dose over the total time period, and the equivalent dose,  $D_{eq}$ , gives the number of equivalent number of MTD doses that equal the sum of the actual doses given over the total time period.

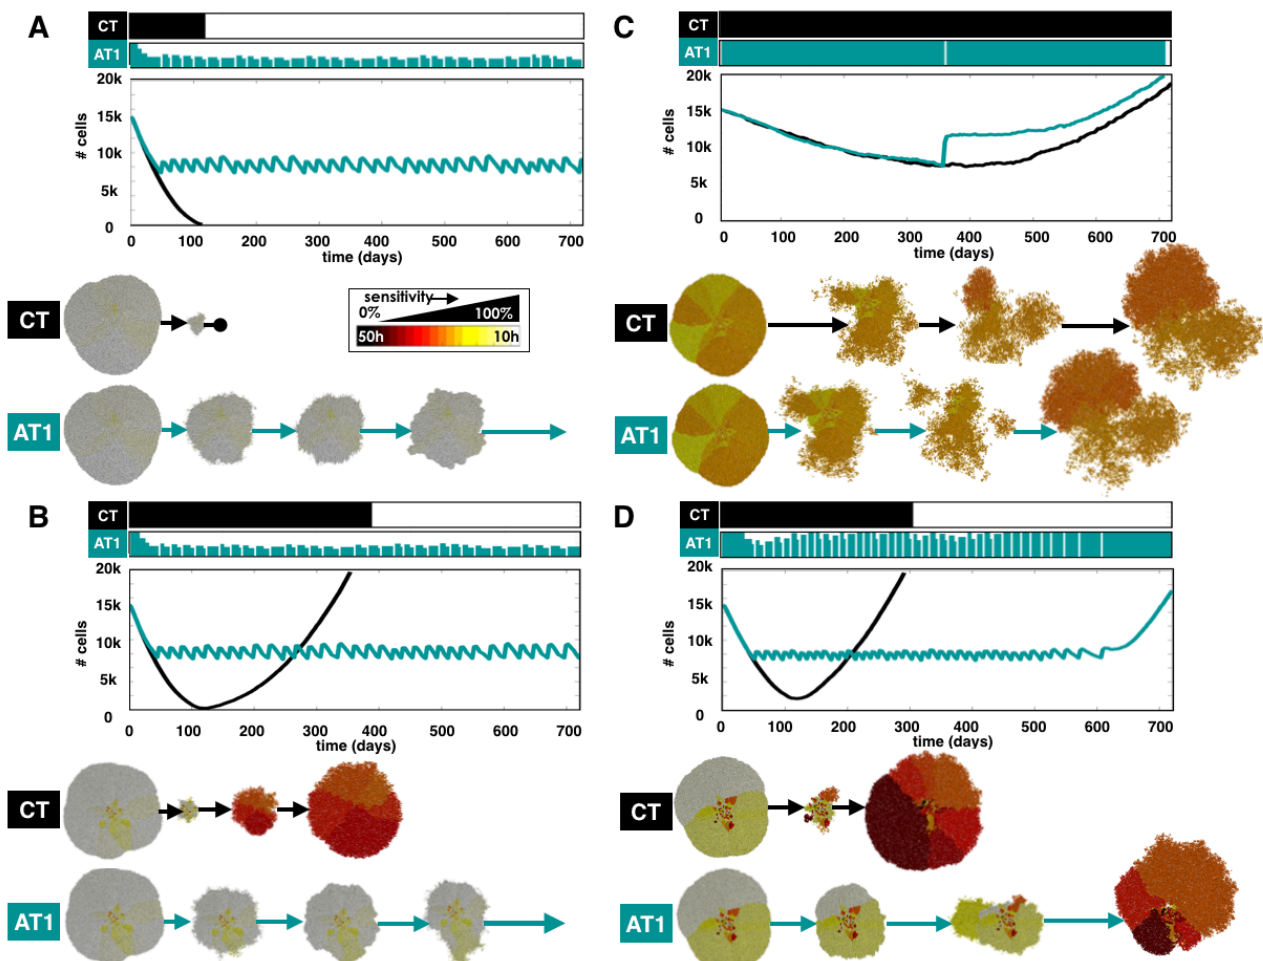


Figure 4. *CT and AT schedules applied to different compositions of tumors: initial distributions before growth with mean sensitivity  $s$  and standard deviation  $\sigma_s$  are A)  $s=100\%$ ,  $\sigma_s=5\%$ ; B)  $s=100\%$ ,  $\sigma_s=25\%$ ; C)  $s=60\%$ ,  $\sigma_s=5\%$ ; and D)  $s=60\%$ ,  $\sigma_s=25\%$ . The top panels show the dose schedules (in % of MTD over time) for each treatment strategy. Below those, the population dynamics for each strategy is shown. The spatial configurations for the two strategies at several different time points are shown below the plots, starting from the same tumor configurations.*

We start by comparing CT with an adaptive scheme that will change the dose by 25% ( $\alpha=0.25$ ) if the population size changes by 5% ( $\beta=0.05$ ) and otherwise follows Eqn 4. In general, there is not one therapy strategy that works best for all tumors, but the response to the treatment depends on the tumor composition. In **Fig. 4** we show the range of response with both treatment strategies: continuous therapy (CT) and adaptive therapy (AT), using the tumors grown in **Fig. 3** (labeled A-D). We see that homogeneous tumors consisting of only sensitive cells (**Fig. 3A**) are curable with CT and can be controlled with AT. While cure by CT takes less time (114 days), AT gives a lower average dose ( $D_{ave}=34\%$ ), but a larger total dose (81 doses vs. 39 doses) over a longer time period (2 years). Tumor B consists primarily of sensitive cells but is initiated with a broader

range of phenotypes, including some treatment-resistant phenotypes that have ended up in the tumor interior surrounded and suppressed by sensitive cells. When treated continuously, the tumor initially responds very well, but this is followed by recurrence of a completely resistant population after around 1 year. However, the same tumor under the adaptive strategy can be controlled with a lower average dose ( $D_{ave}=35\%$ ) and lower total dose (85 doses vs. 122 doses) over the full 2 years. Tumor C is mostly a homogeneous tumor with an average around 60% sensitivity to the drug, which should on average yield a net negative growth rate under drug. There is only a slight response to the treatment at the maximum dose as the cells mean survival rate is around 40%, which also causes the tumor to become diffuse due to the variability amongst neighbors. There is enough phenotypic variation in the population that the most sensitive cells die off and the average sensitivity shifts to allow a net positive growth rate for CT, which leads to tumor progression. Using the AT strategy, the change in the population size during treatment is so slow that the dose never modulates. When the population halves ( $N=0.5N_0$ ) around 1 year, most of the sensitive cells, which were coexisting with the less sensitive cells, have already been eliminated. At this point, the treatment vacation begins, and because the tumor is diffuse, there is a large portion of non-quiescent cells, so we see a swift increase in growth as resistant cells (only slightly <50% sensitivity) take over. Overall, this causes a net loss of 1 month's time over the CT strategy. There is no real advantage of one treatment strategy over the other as both essentially give the same total dose (230 doses) over the full 2-year time period. In tumor D, when there is more phenotypic spread, the sensitive cells again outcompete the more resistant cells during the growth phase. We see here that both treatments lead to eventual recurrence, but AT keeps the total number of cells controlled longer than CT. Notably, AT gives a higher total dose (192 doses vs. 101 doses), but it recurs later (~2 years vs 300 days) and at average of 79% of the MTD.

Analyzing the population dynamics of the whole array of tumor compositions, we can discern why different treatments fail or succeed. **Figure 5** shows how the sensitive and resistant

populations change with CT (**Fig. 5A**) and AT (**Fig. 5B**) for the same array of tumors from **Fig. 3**. The arrays show, for a single run, how the makeup of each tumor changes over the course of treatment. The asterisk marks, for CT, which tumors are cured and, for AT, which tumors are controlled. We see that for both cure and control the sensitive population must be maintained over the treatment period. In some cases, the resistant population starts to appear at the end, but the phenotypes representing most of the population are sensitive for the entire simulation.

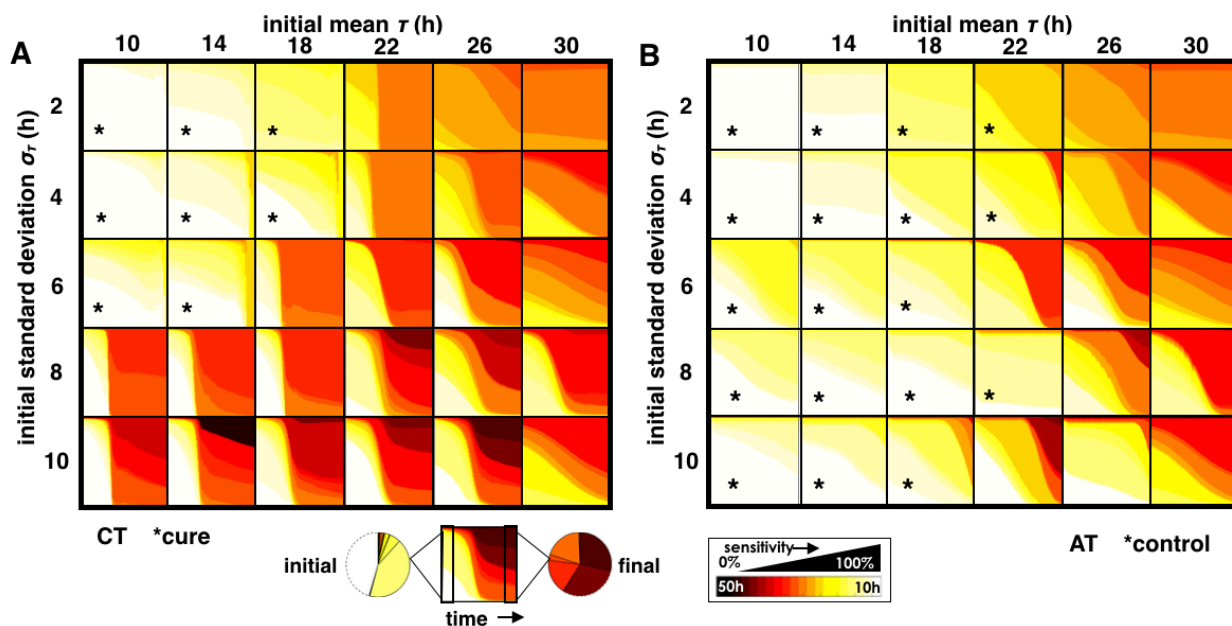


Figure 5. Changes of phenotypic compositions over time for each pre-growth normal distribution with mean cell cycle times  $\tau$  from 10-30h (from left to right), and heterogeneity (standard deviation in cell cycle time  $\sigma_\tau$  of 2-10h from top to bottom). Each square shows the distribution of phenotypes over time as depicted by the key shown below the figure. A) For CT, cure is labeled with an asterisk. B) For AT ( $\alpha=0.25$ ,  $\beta=0.05$ ), control is labeled with an asterisk.

We can check the full range of space to compare the strategies. In Fig. 6A, we compare the time to recurrence (TTR) for each strategy averaging over 3 runs. However, this only works when the tumor does recur. Therefore, we need to set the values for cure and for control because they don't recur over the duration of the simulation. We set the cure TTR to 1000 days and the control TTR to 730 days (the simulation stop time). That means that any TTR for CT at 1000 days was cured 3 out of 3 times, and any TTR for AT at 730 days was controlled 3 out of 3 times. We see that in the more homogeneous and sensitive region, the tumor is cured, but everywhere else, the tumor

recurs. For AT, in the most sensitive regions, regardless of the amount of heterogeneity, the tumor is controlled. For the rest of the array, AT appears to give a larger TTR over CT.

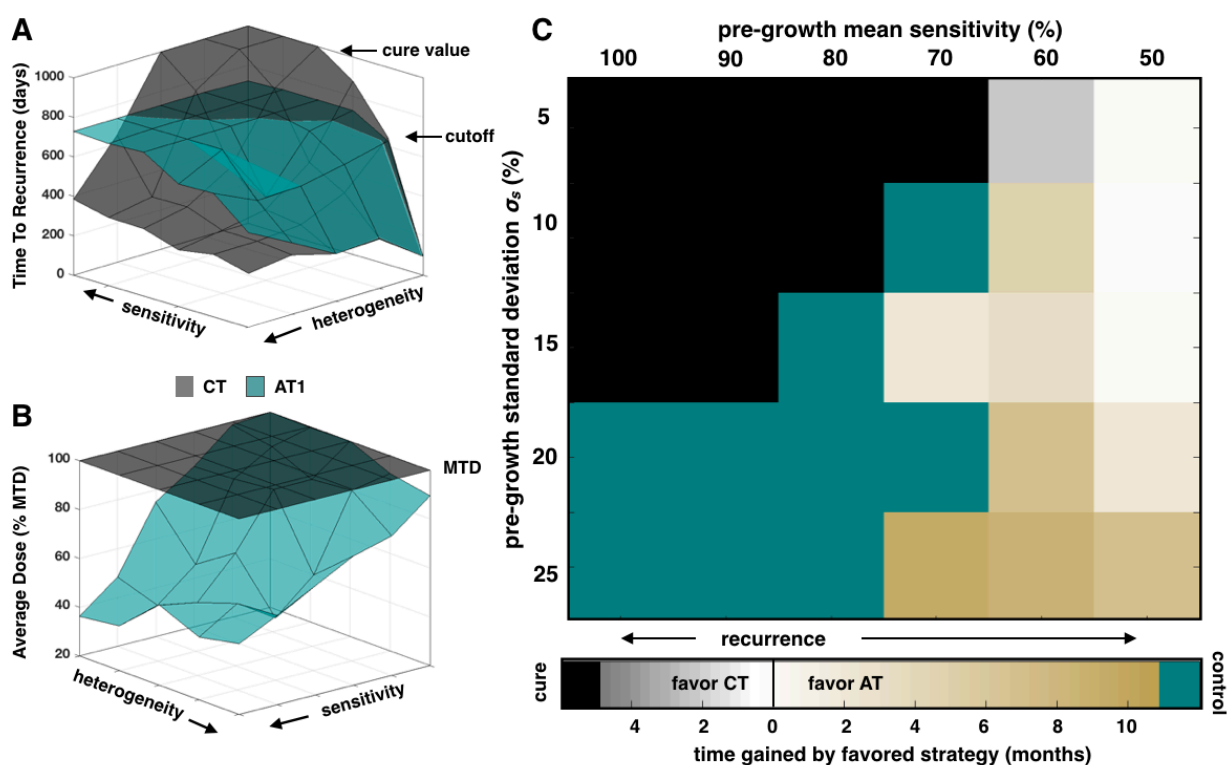


Figure 6. Comparing CT and AT ( $\alpha=0.25$ ,  $\beta=0.05$ ) treatments. A) For the array of tumor compositions, the time to recurrence (TTR) is shown as the average of 3 runs. Since the cured tumors have no recurrence time, we designate the TTR as 1000 days. The AT controlled tumors also do not recur, so their value is set at the simulation cutoff (730 days). B) The average dose over the full simulation, averaged over 3 runs. C) Best 2 out of 3 winning strategies for different tumor compositions. Each square is colored according to the best-case scenario based on 3 simulations according to the key at the bottom. Cure by CT is favored first, control is favored next. If recurrence is most likely, the treatment with the most time gained before recurrence is favored, and the color indicates the average time gained by the winning strategy over the losing strategy (time gained =  $TTR_{favored} - TTR_{unfavored}$ ).

There is also a large difference in the average dose given for each treatment strategy. A comparison of the average dose (average over 3 simulations) for the array of tumor compositions is given in Fig. 6B. Of course, CT gives the MTD over the full duration. For AT, we find that for the most sensitive tumors the average dose can be reduced to 34% of the MTD. This value increases for more resistant tumors. For the most sensitive and homogeneous tumors, the average dose approaches the MTD.

Ignoring the dose, the descending order of desired outcomes is to: cure, maintain, then gain the most time before recurrence. **Fig. 6C** delineates these winning strategy choices for the

best 2 out of 3 trials over the array. We find that the top left side of the plot gives cure by CT as the best and most likely outcome, where the initial distribution is mostly sensitive and homogeneous. On the lower left and some along the edge of the cure region, control by AT is most likely over cure and recurrence, which is a region where the initial distribution is still mostly sensitive but has some heterogeneity. In any other case we have recurrence, and the color bar on the bottom indicates the time gained by the winning strategy. We see that for the top left region, treatments produce nearly the same time to recurrence. That is, the response with such low sensitivities is so poor that dose modulations and treatment vacations are rare, so that the AT treatment actually is nearly the same as the CT treatment (constant MTD). The top right edge of the cure region favors CT by a few months, but the middle and lower right regions of the array favor the AT strategy, because enough sensitive cells in the initial distribution will have trapped the resistant cells during growth and are maintained during treatment.

### **A more vacation-oriented strategy can control the tumor at the expense of higher doses**

The AT schedule used so far, which will now be referred to as AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), was chosen somewhat arbitrarily. Changing  $\alpha$  and  $\beta$  changes the amount the dose is modulated if the population changes by a threshold amount, respectively, and these values can have an impact on AT control. We test another schedule that boosts both values, requiring a larger change in population size (10%) to produce a larger change in the dose modulation (50%), which we will call AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ).



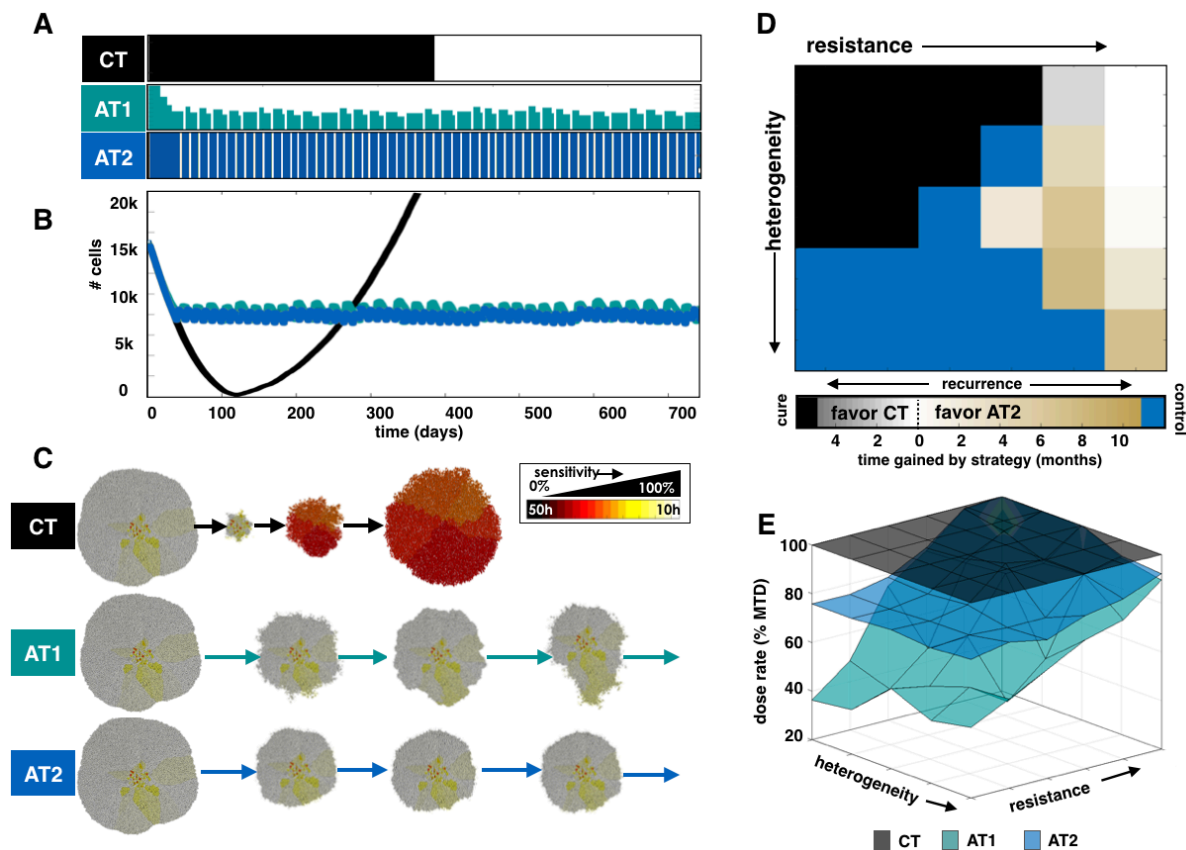


Figure 7. Comparing CT, AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), and AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ) treatments, where adaptive dose schedules are given by Equation 4. A-C) Using a tumor from a pre-growth normal distribution of  $s=100\%$  and  $\sigma_s=25\%$ , we compare these treatment strategies. A) The dose schedules (in % of MTD) for each treatment. B) The population dynamics during each treatment. C) The spatial layout at various time points. D) Best 2 out of 3 winning strategies for different tumor compositions using AT2. Each square is colored according to the best-case scenario based on 3 simulations according to the key at the bottom. Cure by CT is favored first, control is favored next. If recurrence is most likely, the treatment with the most time gained before recurrence is favored, and the square is colored according to the average time gained by the winning strategy over the losing strategy. E) The average dose  $D_{ave}$  over the treatment period for the array of initial distributions for CT, AT1, and AT2.

We compare all treatment strategies (CT, AT1, and AT2) in Fig. 7 using a tumor with a pre-growth distribution that centers around a sensitive phenotype ( $s=100\%$ ) but with a large degree of heterogeneity ( $\sigma_s=25\%$ ). At the start of treatment the resistant cells are trapped in the interior by the sensitive cells. **Figure 7A** gives the dose schedules, **Fig. 7B** shows the population dynamics, and **Fig. 7C** reveals the spatial layouts at various time points. The tumor recurs with the CT strategy while both AT strategies control the tumor for the full 2 years. We can see in the AT1 dose schedule that the dose lowers at the start then incorporates both vacations and dose adjustments to control the tumor. In the AT2 schedule, however, we find that the population does not change fast enough to invoke a dose change, so control is achieved solely by having treatment vacations.

Yet even with practically the same outcome, the previous schedule (AT1) gives a much lower effective dose to the patient. Whilst the CT treatment recurs at 366 days after 122 doses of MTD, the adaptive therapies allow for tumor control over 2 years with an average of  $D_{ave}=35\%(76\%)$  of the MTD resulting in  $D_{eq}=85(185)$  doses for AT1(AT2). With this new schedule, we find similar winning outcomes when compared to CT as with the previous schedule (**Fig. 6C**), but with extended control for tumor compositions that are less sensitive and heterogeneous. The average doses for the whole array of tumor compositions are given for CT, AT1, and AT2 in **Fig. 7E**. We find that while the treatment vacations do reduce the dose, modulating the dose substantially lowers the overall average dose in the regions that are controlled by adaptive strategies. When the proliferating cells that are exposed to the drug are very sensitive, the full dose is not needed and cell kill is still achieved with a smaller dose. However, when the less sensitive phenotypes are proliferating, the increased dose in AT2 grants better control.

### **More treatment vacations are needed to control heterogeneous invasive tumors**

For the adaptive strategy to work, the sensitive cells must impede the proliferation of the resistant cells by competing for space and trapping the resistant cells inside the tumor. However, if the cells can move, the spatial structure that keeps the cells quiescent and hidden from the drug is disturbed. We examine the effect of cell migration by allowing cells to move in persistent random walk as follows. Each cell is given a persistence time (it will move for a length of time randomly chosen from a normal distribution with  $80\pm 40$  minutes), after which it turns at a random angle and starts again. All cells move at a modest speed of  $5\ \mu\text{m}/\text{h}$  for as long as the cell is not in the quiescent state and doesn't contact another cell, in which case the cells simply turn by a random angle and start again with a new persistence time.

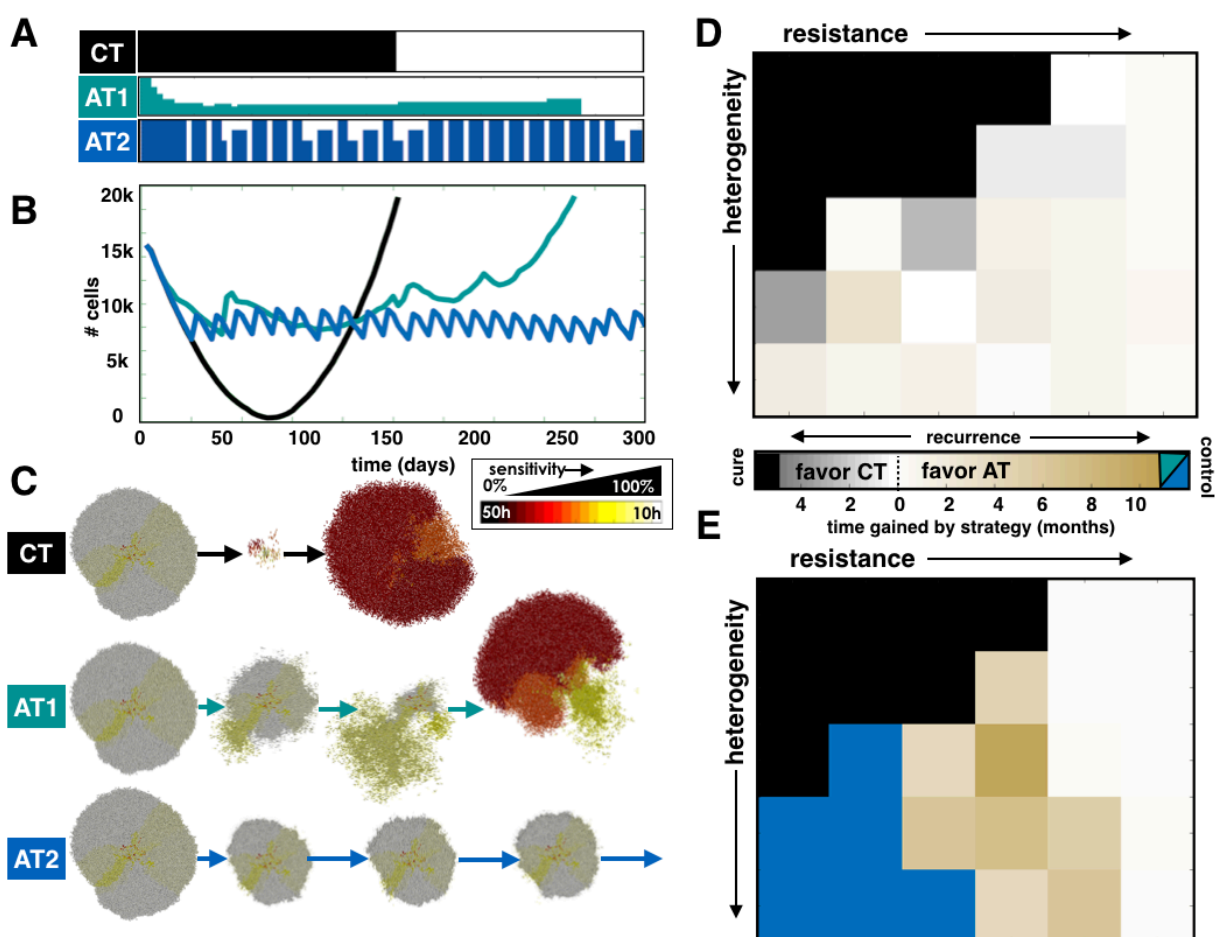


Figure 8. Allowing migration ( $5 \mu\text{m}/\text{h}$ ) reduces the efficacy of control by AT. A-C) An example of a tumor (with a pre-growth distribution of  $s=100\%$  and  $\sigma_s=25\%$ ) subject to strategies CT, AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), and AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ). We show the dose schedules (A), the population over time (B), and the spatial layout (C) for each strategy. D-E) The best 2 out of 3 winning strategies for different tumor compositions using AT1 (D) and AT2 (E). Each square is colored according to the best-case scenario based on 3 simulations according to the key. Cure by CT is favored first, control is favored next, and if recurrence is most likely, the treatment with the most time gained before recurrence is favored, and the square is colored according to the average time gained by the winning strategy over the losing strategy.

**Figure 8** shows an example with the same initial conditions as the previous example ( $s=100\%$  and  $\sigma_s=25\%$ ), but with migrating cells. Again, the dose schedules are given (**Fig. 8A**), the population over time (**Fig. 8B**), and the spatial layout (**Fig. 8C**). We see that for the CT case, compared to the case without migration (**Fig. 7**), we have a much quicker decline in the population followed by a much faster relapse at 153 days. For this strategy, the patient gets 51 doses at MTD. Due to the spatial spreading of the population from migration, the drug is able to get to the cells quicker. This also means that the resistant cells will escape sooner so there is a faster return of growth. For AT1, the tumor population doesn't get small enough to trigger a vacation, but dose is

observed to modulate. The dose reduces then rises back up at the end when the resistant, unresponsive cells outgrow. This schedule ends up with a recurrent tumor at 258 days after  $D_{eq}=101$  doses at MTD delivered at  $D_{ave}=30\%$  MTD. Using AT2, the tumor is controlled for the 2 year time period (not fully shown). The dose is observed to modulate because the larger amount of proliferating cells (and therefore susceptible to the drug) causes larger fluctuations in the population size triggering a dose change. But, by employing the quicker switching between fast growth during the vacations and fast death with higher doses, the tumor is able to keep the sensitive population at the invasive edge. This schedule gives  $D_{ave}=67\%$  MTD and  $D_{eq}=163$  doses.

The full sweep of the array space is shown in **Fig. 8D** comparing CT vs. AT1 and in **Fig. 8E** comparing CT vs. AT2. We find that while most of the original region cured by CT remains, allowing migration has destroyed the favorability of AT1 over CT. However, in comparing CT vs. AT2, we find that the regions previously controlled with AT2 are somewhat preserved. The main point is that with a small increase in migration rate, AT1 loses the ability to control the tumor, whereas control is partly preserved with AT2. As we increase the migration rate further to 10  $\mu\text{m}/\text{h}$ , neither adaptive treatment is favored to control the tumor (**Fig. S1A-B**).

### **Treatment vacations help delay recurrence of heterogeneous tumors with phenotypic drift**

We have considered treatment responses to tumors that have pre-existing heterogeneity but assume that progeny inherit the same phenotypes from their initial parental lineage. However, this ignores that a cell's phenotype can vary over generations. To test how phenotypic drift affects the response to treatment strategy, we allow a cell the opportunity to change its proliferation rate slightly from its mother cell upon division. In particular, there is a 10% probability at each division that a cell will either increase its cycle time by 1 h, decrease its cycle time by 1 h or stay the same (ensuring that the IMT stays within the allowed range of 10-50h).

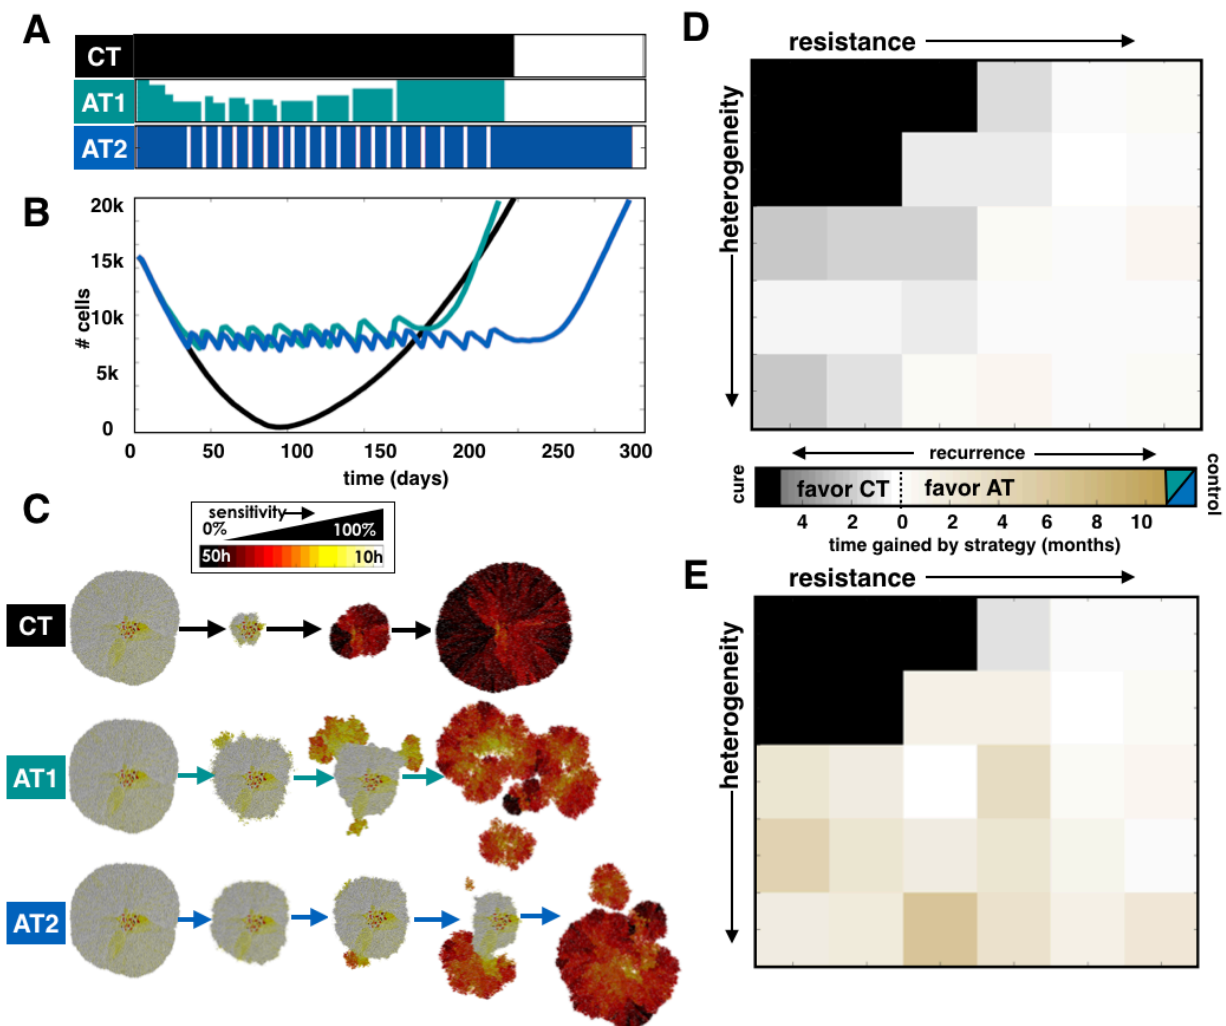


Figure 9. Allowing phenotypic drift reduces the efficacy of AT. A-C) An example of a tumor (with a pre-growth distribution of  $s=80\%$  and  $\sigma_s=25\%$ ) subject to strategies CT, AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), and AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ). A) Shows the schedules, B) the population dynamics, and C) the spatial layout for each strategy. D-E) The best 2 of 3 winning strategies for CT vs. AT1 (D) and for CT vs. AT2 (E). Each square is colored according to the best-case scenario based on 3 simulations according to the key. Cure by CT is favored first, control is favored next. If recurrence is most likely, the treatment with the most time gained before recurrence is favored, and the square is colored according to the average time gained by the winning strategy over the losing strategy.

In **Fig. 9**, we show an example with initial conditions of  $s=80\%$  and  $\sigma_s=25\%$ . The treatment schedules for CT, AT1, and AT2 are shown in **Fig. 9A** as a percent of the MTD. The population dynamics are shown in **Fig. 9B**. Finally, the spatial configurations for each treatment at various time points are displayed in **Fig. 9C**. In all cases, treatments eventually fail. CT, as in the previous case with migration, automatically selects for the resistant cells and delivers the MTD for 297 days, which corresponds to 99 doses of MTD. AT1 lasts nearly the same amount of time as the CT strategy, but by using a smaller dose ( $D_{ave}=66\%$  over 285 days with  $D_{eq}=63$  doses). Because of the

drifting phenotype, the sensitive cells can now become more resistant over generations. It can be seen that for AT1 there are individual hotspots on the tumor's edge that drift to more resistant phenotypes and outgrow while the more sensitive cells are killed off. This results in separation of the mass into several individual tumor clumps around these regions. AT2 is able to control the tumor a little longer than AT1 (387 days) by having treatment vacations and no dose modulation. During this time,  $D_{ave}=85\%$  and  $D_{eq}=109$  doses. Because of the sharp contrast between no drug and full MTD, even an initial hotspot with increased resistance gets suppressed by the sensitive cells during the vacation, at least for a while.

The array of different tumor compositions is evaluated to compare CT vs. AT1 (**Fig. 9D**) and CT vs. AT2 (**Fig. 9E**). The region originally cured by CT is significantly reduced, and in both cases the ability to control the tumor through adaptive means is completely eliminated. CT actually does a better job with just a few months gain over AT1 in most regions. But again, the AT2 strategy delays recurrence for longer over the CT strategy for heterogeneous tumor compositions. Regardless, neither treatment sustains the control for long, as the net gain with AT2 is at the most 6 months. Increasing the drift rate further to occur at every division destroys any control whatsoever using CT, AT1, or AT2 (**Fig. S1C-D**).

### **Why do treatment vacations maintain tumor sensitivity for longer?**

When migration and phenotypic drift were not allowed, dose modulation granted a way to control a tumor while lowering the dose rate. However, an emphasis on treatment vacations can partly preserve this control. To understand why, it is helpful to introduce the effective net proliferation rate:  $\rho_{eff}=\rho(1-P_{death}/0.5)$ , where  $P_{death}$  is the probability of death, and  $\rho$  is the proliferation rate unperturbed by treatment. This value gives a better understanding of how dose and sensitivity combine to give a net proliferation rate across the range of phenotypes. It essentially modulates the proliferation rate by the probability of death due to a certain dose (% of MTD), such that if

$P_{death}=0$ ,  $\rho_{eff}=\rho$ , if  $P_{death}=0.5$ ,  $\rho_{eff}=0$  and at  $P_{death}=1.0$  the cells die off as fast as they divide. The effective net proliferation is shown in the background heat map of **Figure 10**. The highest sensitivities have both the highest growth rates (without drug) and the highest death rates (at MTD), which are the darkest colors at the top row of the heat map. Remaining in the red zone means maintaining net cell kill over proliferation.

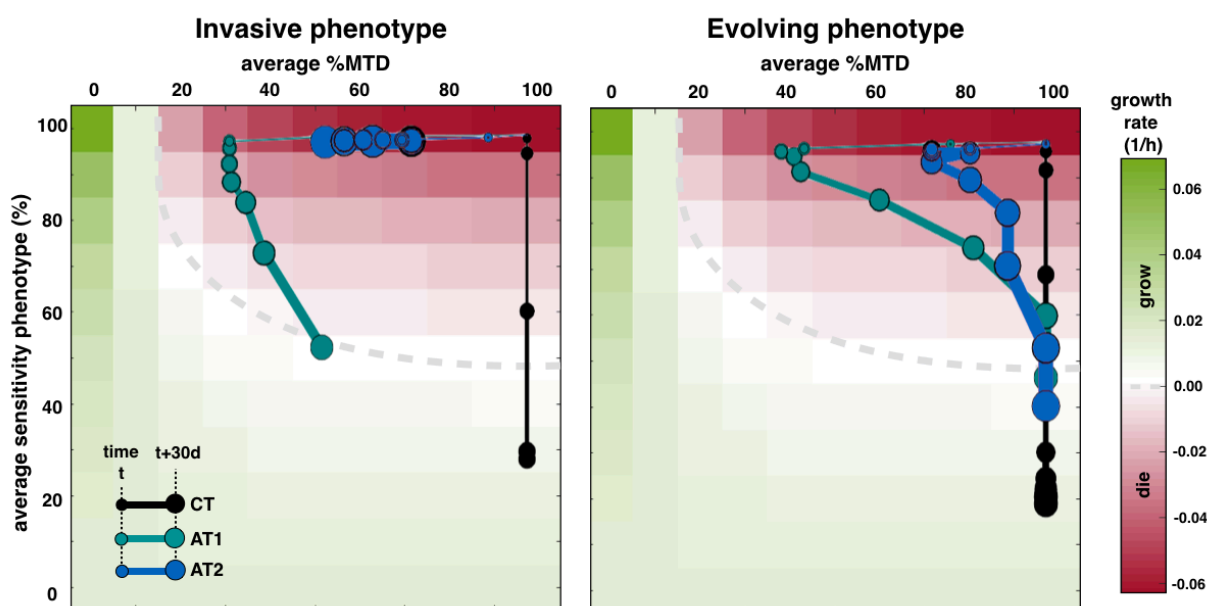


Figure 10. Effect of modulating dose on the growth rates of different phenotypes (sensitivities) for different treatment strategies: CT, AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), and AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ). We use the trajectories for the invasive (left) and drifting (right) phenotypes from Figure 7A and Figure 8A, respectively. The background heat map displays the effective growth rate of a phenotype, which modulates the given growth rate by the probability of death due to a specific dose. The gray dashed line indicates where net growth is zero. Each colored line shows how the sensitivity of the tumor population changes with the given dose for each strategy. A point is placed for each month in time at the average sensitivity and the average dose of each treatment over that month. Points get larger in size at each represented time point.

If we plot the previous examples as trajectories through the space of average sensitivity over average dose, we can see why failures occur for the invasive (left) and evolving (right) phenotypes (same examples as from **Figure 8A** and **Figure 9A**). Using the CT strategy means a constant dose. This is shown with the black line that does not vary in the x-direction. If all phenotypes were to lie above the dashed line (net zero proliferation), cure is possible. However, because there are a variety of existing phenotypes, the mean sensitivity decreases over time. As the most sensitive cells are eliminated and space becomes available the line is crossed where the

more resistant cells can have a net positive growth and the tumor recurs. With the invasive phenotype AT1 kills off the most sensitive clones sequentially as the dose is modulated to keep the tumor at a constant size. The dose modulation strategy will shift to doses lower than the MTD to find the lowest dose that still kills the most sensitive phenotypes. Eventually, the most sensitive cells will be killed, but because the dose is low the next most proliferative cells will have a net positive proliferation rate and outgrow. The increased growth then causes a dose increase to keep proliferation in check. This essentially "picks off" sequentially sensitive phenotypes with ever increasing doses until the more resistant cells take a hold of the invasive edge. At that point, the resistant cells have a small (albeit positive) growth rate and there is not much chance for control. With the AT2 strategy that relies more on treatment vacations, there is a sharp transition between the quick drug-free expansions and MTD-induced contractions. In this case, with the most sensitive cells at the proliferating edge, the net decline during treatment is approximately equal to the net increase when the dose is zero. As long as mostly sensitive cells continue to lie on the outer rim, the tumor can be controlled.

With a drifting phenotype a cell's sensitivity can change upon division to become more sensitive or more resistant. This allows movement along the y-axis. For all treatment strategies, the drug kills off the more sensitive cells, so the phenotypes eventually end up further into the resistant region where the net proliferation rate is always positive. While this does not change the outcome much for CT, the outgrowth of just a few resistant phenotypes causes a quicker dose increase for AT1, which causes a quicker decline of the most sensitive cells. This results in an eventual recurrence. By using the AT2 strategy that shifts between extreme doses, the fastest proliferators are maintained during vacations that are also very sensitive during drug application in order to suppress other cells. Unfortunately, the effect is not sustained for very long.

#### **4. Discussion**



Our simulations demonstrate that there is not a one-size-fits-all solution drug strategy for all tumors. We modify an off-lattice agent based model (15) to incorporate treatment response with the central goal to compare tumor kill that reduces the takeover of resistant phenotypes to standard therapy at the MTD. We find that conventional MTD application of cancer therapy can cure homogeneous tumors that consist entirely or almost entirely of cells that are sensitive to the applied treatment. While this is improbable at advanced stages, it is observed clinically as some tumors that histologically appear homogeneous, such as testicular cancer and some lymphomas, are routinely cured by conventional chemotherapy. However, cure is not typically achievable in cancers that contain cells that are resistant therapy due to their inherent phenotypic properties or environmentally-mediated mechanisms. This is consistent with clinical observations that most common epithelial tumors exhibit multiscalar heterogeneity (i.e. on both radiological and pathological images) and are not curable once they become metastatic.

In tumors that are not curable by available therapy, our models find that evolution-based strategies that exploit the fitness cost of resistance can delay treatment failure and tumor progression. Recent experimental studies have shown adaptive therapy can enforce prolonged tumor control (18), however, none of the prior theoretical models have included spatial dynamics and thus assume that the different tumor subpopulations are well mixed (11,12). This is in contrast to radiological and pathological images of tumor that show marked spatial heterogeneity (21-24). Spatial structure can affect the emergence of resistant phenotypes due to space limitation (15,25) and limited drug perfusion (26,27). We present some insight into how spatial context might influence disease control, which reflects similar ideas presented in a recent study with growing bacterial colonies of *E. Coli* (25). In this study, Fusco *et al.* showed how less fit resistant clones can be trapped due to spatial constraints, but high intensity drug exposure leads to the competitive release of dormant mutants. Our model shows the same nature of competition, and we are able to test multiple treatment strategies on different compositions of phenotypes.

We consider phenotypic heterogeneity over a wide spectrum of tumor compositions. A larger variation in cell phenotypes beyond just sensitive and resistant (with corresponding fitness differences) can drive changes in the dominating population at different doses. By systematically testing an array of different phenotypic distributions we can delineate different regions where some strategies work better than others. We also find that some AT strategies are better than others for different situations. For example, with cell migration and phenotypic evolution, we can apply a strategy with less dose modulation and more emphasis on treatment vacations to keep the population sensitive. The extreme dose changes keeps the tumor responding quickly during both the growth phase and the treatment phase to either maintain the spatial structure in the case of migration or to prevent the evolution of less sensitive cells in the case of phenotypic drift.

An ongoing clinical trial tests these evolutionary principles in patients with metastatic castration resistant prostate cancer to try to prevent the evolution of resistance to abiraterone therapy (28). In this trial, abiraterone is discontinued when the blood Prostate Specific Antigen (PSA) concentration falls below 50% of the initial value and does not resume until the PSA returns to the pre-treatment level. PSA changes during therapy are shown for 2 patients in **Fig 11A** and **Fig. 11B**. It is important to note that: 1) each patient serves as its own control to calibrate the PSA as a relative value, and 2) the adaptive schedules effectively personalize the treatment schedule to patient response so that while one patient has only 2 treatments in a year, another gets 3 in 10 months. Yet both are still under control.

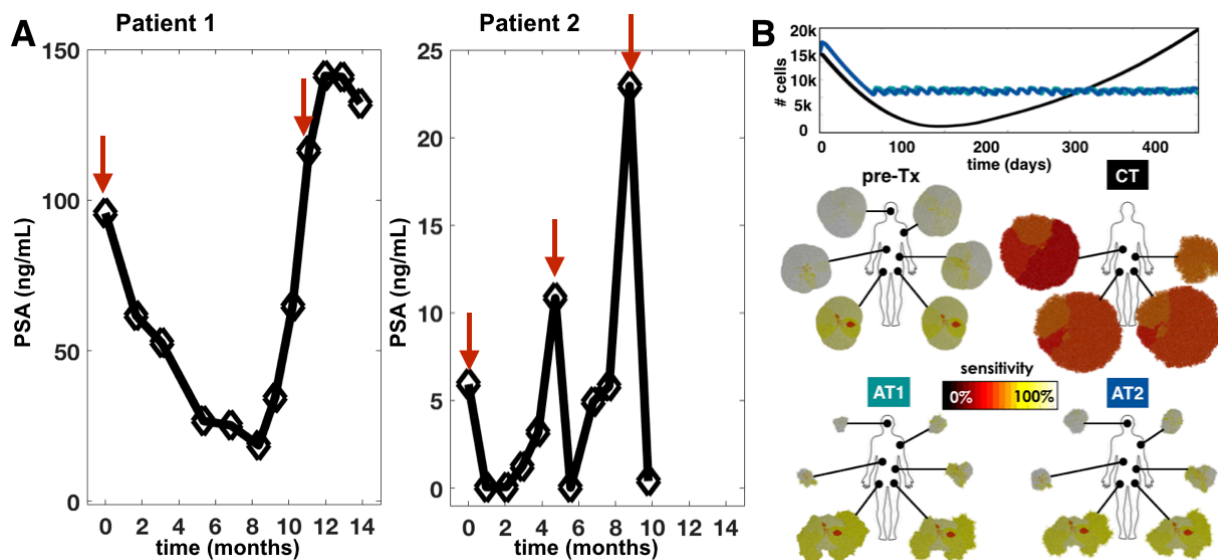


Figure 11. A) PSA over time for two patients enrolled in the adaptive therapy trial for metastatic castrate resistant prostate cancer (28). The red arrows show an applied dose of abiraterone given when the PSA returns to pre-treatment levels. B) The model checks that tumors with dissimilar compositions can be controlled by adaptive scheduling. We tested the response for CT, AT1 ( $\alpha=0.25$ ,  $\beta=0.05$ ), and AT2 ( $\alpha=0.50$ ,  $\beta=0.10$ ) using the sum of all tumor cells to determine the dose for the AT treatments. The total tumor burden dynamics from all metastatic sites are shown at the top, while the individual spatial compositions at the end of the simulation are given below for each treatment strategy.

Comparing widely disseminated cancer, where the distribution of tumor cell subpopulations will likely vary among the metastatic sites, to our model, which only considers the response of single tumors may seem disconnected. It is thus important to ask how much variation in clinical response to CT and AT should be expected. Testing the treatments on a set of tumor metastases with mixed compositions shows that control can be achieved (**Fig. 11C**). By using the total numbers of cells across all tumors to represent a systemic measure of burden, we find that with the CT strategy the more sensitive tumors respond with complete eradication, while the less sensitive tumors eventually recur. In contrast, both AT schedules can still control the disease. While the more sensitive tumors shrink, the less sensitive ones grow in size keeping the total population relatively constant. The full dose schedules are shown in **Fig. S2** along with a case with metastases that are heterogeneous but have all the same composition. In the latter situation, we get very similar results to what would happen with single independent tumors.

The methods for actually measuring tumor burden non-invasively and often enough to direct treatment decisions are not well developed for all cancers. In the ongoing adaptive therapy

trial on prostate cancer (28), PSA is used as a marker for disease burden. PSA may not be a perfect indicator, but it is readily available, standardized, and utilized (29). Other biomarkers are used for surveillance of various cancers with degrees of specificity and sensitivity (30). As technologies for serum biomarkers, circulating tumor cells, and cell-free DNA continue to be developed (31), challenges can be addressed to better manage advanced and disseminated cancers through periodic blood draws.

The model presented here is greatly simplified and abstracted to study the impact of spatial intratumoral heterogeneity on controlling tumor progression using adaptive therapy. We have ignored pharmacokinetics, spatially weighted dose dependence, microenvironmental influence, and dose-dependent acquired drug resistance. All of these are possible extensions and modulations of this work, but we have chosen a simple starting point to understand how a spectrum of cell phenotypes compete for space under different drug strategies.

Metastatic and drug-resistant cancer has remained tricky to navigate. Evolutionary-based treatments take into account tumor heterogeneity to avoid drug resistance takeover while also sparing the patient's immune and vascular system from harsh and damaging insults. However, we find that some tumors are simply too resistant at the start of treatment that the drug is not effective. If we were to consider several targetable traits, a more complex adaptive strategy could be used with multiple drugs. With multiple targets, phenotypes can be steered by one therapy to render the population more vulnerable to second therapy (32-34). In addition, given that sensitivity of resistant populations may be restored after sufficiently long treatment vacations (10), the drifting phenotype dynamic can be exploited. Further integration of clinical trials, *in vitro* and *in vivo* models, and mathematical models will help improve adaptive strategies that promote tumor control aimed at keeping the patient alive longer and with the best quality of life.

## Acknowledgements

## Bibliography

1. Saunders NA, Simpson F, Thompson EW, Hill MM, Endo-Munoz L, Leggatt G, et al. Role of intratumoural heterogeneity in cancer drug resistance: molecular and clinical perspectives. *EMBO Mol Med*. EMBO Press; 2012 Jun 25;4(8):675–84.
2. Kam Y, Das T, Tian H, Foroutan P, Ruiz E, Martinez G, et al. Sweat but no gain: Inhibiting proliferation of multidrug resistant cancer cells with “ersatzdroges.” *Int J Cancer*. 2014 Sep 2;136(4):E188–96.
3. Hiley CT, Swanton C. Pruning Cancer's Evolutionary Tree with Lesion-Directed Therapy. *Cancer Discovery*. American Association for Cancer Research; 2016 Feb 4;6(2):122–4.
4. Gillies RJ, Verduzco D, Gatenby RA. Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. *Nat Rev Cancer*. 2012.
5. Korolev KS, Xavier JB, Gore J. Turning ecology and evolution against cancer. *Nat Rev Cancer* [Internet]. Nature Publishing Group; 2014 Apr 17;14(5):371–80. Available from: <http://dx.doi.org/10.1038/nrc3712>
6. Aktipis CA, Boddy AM, Gatenby RA, Brown JS, Maley CC. Life history trade-offs in cancer evolution. *Nat Rev Cancer*. Nature Publishing Group; 2013 Nov 11;13(12):883–92.
7. Crawford S. Is it time for a new paradigm for systemic cancer treatment? Lessons from a century of cancer chemotherapy. 2013 Jun 24;:1–18.
8. Zeilinger AR, Olson DM. Competition between stink bug and heliothine caterpillar pests on cotton at within-plant spatial scales. ... *Experimentalis et Applicata*. 2011.
9. Kreso A, O'Brien CA, van Galen P, Gan OI, Notta F, Brown AMK, et al. Variable Clonal Repopulation Dynamics Influence Chemotherapy Response in Colorectal Cancer. *Science*. 2013 Jan 31;339(6119):543–8.
10. Chmielecki J, Foo J, Oxnard GR, Hutchinson K, Ohashi K, Somwar R, et al. Optimization of Dosing for EGFR-Mutant Non-Small Cell Lung Cancer with Evolutionary Cancer Modeling. *Science Translational Medicine*. American Association for the Advancement of Science; 2011 Jul 6;3(90):90ra59–9.
11. Silva AS, Kam Y, Khin ZP, Minton SE, Gillies RJ, Gatenby RA. Evolutionary Approaches to Prolong Progression-Free Survival in Breast Cancer. *Cancer Research*. 2012 Dec 16;72(24):6362–70.

12. Silva AS, Durand A, Ribeiro MC, Alsina M, Shain K, Baz R. An Evolutionary Approach for Personalized Therapy in Multiple Myeloma. *AM*. 2016;07(02):159–69.
13. Moore N, Houghton JM, Lyle S. Slow-cycling therapy-resistant cancer cells. *Stem cells and development*. 2011.
14. Melnyk AH, Wong A, Kassen R. The fitness costs of antibiotic resistance mutations. *Evol Appl*. 2015.
15. Gallaher JA, Anderson ARA. Evolution of intratumoral phenotypic heterogeneity: the role of trait inheritance. *Interface Focus*. 2013;3(4):20130016.
16. Gesztelyi R, Zsuga J, Kemeny-Beke A, Varga B. The Hill equation and the origin of quantitative pharmacology. *Archive for history of ...* 2012.
17. Gardner SN. A Mechanistic, Predictive Model of Dose-Response Curves for Cell Cycle Phase-specific and -nonspecific Drugs. *Cancer Research*. American Association for Cancer Research; 2000 Mar 1;60(5):1417–25.
18. Enriquez-Navas PM, Kam Y, Das T, Hassan S, Silva AS, Foroutan P, et al. Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Science Translational Medicine*. 2016 Feb 24;8(327):327ra24–4.
19. Gelles JD, Chipuk JE. Robust high-throughput kinetic analysis of apoptosis with real-time high-content live-cell imaging. *Cell Death & Disease*. 2016.
20. van Nieuwenhuijze AEM, Van Lopik T, Smeenk RJT, Aarden LA. Time between onset of apoptosis and release of nucleosomes from apoptotic cells: putative implications for systemic lupus erythematosus. *Annals of the Rheumatic Diseases*. BMJ Publishing Group Ltd; 2003 Jan 1;62(1):10–4.
21. Gerlinger M, Rowan A, Horswell S. Intratumor Heterogeneity and Branched Evolution Revealed by Multiregion Sequencing. ... *England Journal of ...* 2012.
22. Gillies RJ, Anderson AR, Gatenby RA, Morse DL. The biology underlying molecular imaging in oncology: from genome to anatome and back again. *Clinical Radiology*. Elsevier; 2010 Jul 1;65(7):517–21.
23. Gatenby RA, Grove O, Gillies RJ. Quantitative Imaging in Cancer Evolution and Ecology. *Radiology*. 2013 Oct;269(1):8–14.
24. O'Connor J, Rose CJ, Waterton JC. Imaging intratumor heterogeneity: role in therapy response, resistance, and clinical outcome. *Clinical Cancer ...* 2015.
25. Fusco D, Gralka M, Kayser J, Anderson A. Excess of mutational jackpot events in expanding populations revealed by spatial Luria-Delbruck experiments. *Nature*. 2016.
26. Fu F, Nowak MA, Bonhoeffer S. Spatial Heterogeneity in Drug Concentrations

- Can Facilitate the Emergence of Resistance to Cancer Therapy. Antia R, editor. PLoS Computational Biology. 2015 Mar 19;11(3):e1004142.
27. Shah AB, Rejniak KA, Gevertz JL. Limiting the development of anti-cancer drug resistance in a spatial model of micrometastases. *Mathematical biosciences and ....* 2016.
  28. Moffitt Cancer Center. Adaptive Abiraterone Therapy for Metastatic Castration Resistant Prostate Cancer. 2015. pp. 1–4. Available from: <https://clinicaltrials.gov/ct2/show/NCT02415621>
  29. Wallace TJ, Torre T, Grob M, Yu J, Avital I, Brücher B, et al. Current approaches, challenges and future directions for monitoring treatment response in prostate cancer. *Journal of Cancer*. 2014.
  30. Duffy MJ. Tumor markers in clinical practice: a review focusing on common solid cancers. *Medical Principles and Practice*. 2012.
  31. Lowes L, Bratman SV, Dittamore R, Done S. Circulating tumor cells (CTC) and cell-free DNA (cfDNA) workshop 2016: scientific opportunities and logistics for cancer clinical trial incorporation. *International Journal of ....* 2016.
  32. Nichol D, Jeavons P, Fletcher AG, Bonomo RA, Maini PK, Paul JL, et al. Steering Evolution with Sequential Therapy to Prevent the Emergence of Bacterial Antibiotic Resistance. Antia R, editor. PLoS Computational Biology. Public Library of Science; 2015 Sep 11;11(9):e1004493–19.
  33. Basanta D, Gatenby RA, Anderson ARA. Exploiting Evolution To Treat Drug Resistance: Combination Therapy and the Double Bind. 2012 Apr 2;9(4):914–21. Available from: <http://pubs.acs.org/doi/abs/10.1021/mp200458e>
  34. Cunningham JJ, Gatenby RA, Brown JS. Evolutionary dynamics in cancer therapy. *Mol Pharm*. 2011 Dec 5;8(6):2094–100.