DETERMINISTIC EVOLUTIONARY GAME THEORY MODELS FOR THE NONLINEAR DYNAMICS AND CONTROL OF CHEMOTHERAPEUTIC RESISTANCE

by

Yongqian Ma

A Dissertation Presented to the FACULTY OF THE USC GRADUATE SCHOOL UNIVERSITY OF SOUTHERN CALIFORNIA In Partial Fulfillment of the Requirements for the Degree DOCTOR OF PHILOSOPHY (PHYSICS & ASTRONOMY)

August 2020

Copyright 2020

Yongqian Ma

Dedication

To my dear parents.

Acknowledgments

Firstly thanks a lot for my advisor Paul Newton and my lab mates. You provided me a lot of ideas about research and my future career. Without your help I won't be able to complete all the work smoothly.

Secondly I want to thank all my friends, my roommates, my classmates, etc. They kept me accompany throughout my PhD life in the past six years in the US or from China. It wouldn't be as colorful as I had without you.

After that I received a lot of encourage and care from my relatives, my uncles, aunts and cousins. During this special time of global pandemic, it's so warm to hear from the other side of the earth.

Then as one of my numbered interest I had a lot fun watching anime and playing games during my spare time. Thanks for all the wonderful products and I hope this year won't be the terminal, but a turning point for the industry.

Last but not least, there are also many online friends and groups of netizens. We may or may not know each other, but during my struggling time in the midnight, I really got relieved by joking, talking or even trolling with each other. Best regards for everyone in the world to survive the Covid-19 pandemic and to live a happy and peaceful life.

Contents

D	edica	tion	ii			
Acknowledgments						
Li	st of	Tables	vi			
\mathbf{Li}	st of	Figures	vii			
A	bstra	let	xx			
1	Intr	oduction	1			
2	Capitalizing on Competition: An Evolutionary Model of Competitive Release					
	in N	Aetastatic Castration Resistant Prostate Cancer Treatment	6			
	2.1	Introduction	7			
		2.1.1 The likelihood of pre-existing resistance	8			
		2.1.2 Leveraging the cost of pre-existing resistance	9			
		2.1.3 Retrospective analysis of metastatic castration-resistant prostate cancer	11			
	2.2	Materials and Methods	13			
		2.2.1 The replicator equation model	16			
		2.2.2 The linearized system and the cost of resistance	18			
	2.3	Results	19			
		2.3.1 Managing competitive release	21			
	2.4	Discussion	24			
3	Nor	nlinear adaptive control of competitive release and chemotherapeutic				
	resi	stance	30			
	3.1		31			
	3.2	A three-component replicator system	34			
	3.3	The prisoner's dilemma as a cancer model	37			
	3.4	Results	41			
		3.4.1 Continuous therapy	41			
		3.4.2 Time-dependent therapy	45			
		3.4.3 Robustness	48			
	3.5	Discussion	50			

4	The	role of synergy and antagonism in designing multi-drug adaptiv	e			
	cher	notherapy schedules	56			
	4.1	Introduction	57			
	4.2	A three-component replicator system	60			
		4.2.1 Why the prisoner's dilemma game?	60			
		4.2.2 The three-component model \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots	63			
	4.3	Constant chemotherapy	65			
		4.3.1 Additive interactions $e = 0$	65			
		4.3.2 Non-additive interactions	68			
	4.4	Designing adaptive schedules	71			
	4.5	Discussion	76			
5	Opt	imal control of evolutionary games: The two-by-two case	79			
	5.1	Introduction	80			
	5.2	model	81			
		5.2.1 evolutionary games and replicator dynamics	81			
		5.2.2 Pontryagin Maximum Principle	84			
	5.3	Results	86			
	5.4	discussion	87			
		5.4.1 f_2 decrease \ldots	87			
		5.4.2 f_1 increases	90			
		5.4.3 Time re-scaling	92			
	5.5	Conclusion	92			
6	Opt	imal control of levitation by blowing	94			
-	6.1	Introduction	94			
	6.2	The model	95			
	6.3	Time-dependent blowing	99			
	6.4	Optimal blowing schedules	102			
		6.4.1 Minimum time strategy	103			
		6.4.2 Minimum air volume strategy	105			
		6.4.3 Minimum energy strategy	106			
	6.5	Pontryagin Maximum Principle	107			
		6.5.1 Singular control	110			
	6.6	Discussion	111			
Reference List						
A	repl	icator dynamics of tumor growth	129			

List of Tables

- 4.1 Total dose (D), Time period (T), and Average Dose (D/T) associated with adaptive therapies with antagonistic, additive, and synergistic drug interactions. 74

List of Figures

- 2.3Model fits of PSA measurement data for metastatic castrationresistant prostate cancer — Prostate-specific antigen (PSA) measurement data of four representative patients from three randomised clinical trials with metastatic castration-resistant prostate cancer. (A) A single patient data shown with model fit (eqn. 2.4) of growing resistant population (green), decaying sensitive population (red), where the sum represents an overall initial good response and eventual relapse (q(t), dashed line). Left column: treatment with prednisone only [75]; middle column: treatment with mitoxantrone and prednisone [112]; right column: treatment with docetaxel and prednisone [89]. PSA data is normalized by q(t = 0) (black dots). The data is fit using the exponential model (eqn. 2.1; blue curve) and the replicator model (eqns 2.2, 2.3). Each patient is fit reasonably well with both models. Data was fit by parameter sweep of cost ($\alpha - \beta$, eqn. 2.15, 2.16), initial fractional resistance f and selection pressure w. Parameters used are $\alpha - \beta = [0.02, 0.07, 0.20,$ 0.00, 0.20, 0.00, 0.20, 0.03, 0.20, 0.01, 0.15, 0.19; f = [0.20, 0.02, 0.09, 0.05, 0.09, 0.05, 0.00, 0.0(0.03, 0.07, 0.10, 0.10, 0.09, 0.10, 0.03, 0.03]; w = [0.10, 0.20, 0.20, 0.30, 0.20, 0.20, 0.30, 0.20]; w = [0.10, 0.20, 0.20, 0.30]; w = [0.10, 0.20, 0.20, 0.30]; w = [0.10, 0.20, 0.20]; w = [0.10, 0.20, 0.20]; w = [0.10, 0.20, 0.20]; w = [0.10, 0.20]; w = [0.10.10, 0.30, 0.15, 0.10, 0.15, 0.20, 0.20 for B - M respectively.

Dynamic phase portraits before and during chemotherapy — (A) 2.5Trilinear coordinate phase space representation; (B) Schematic of proposed adaptive therapy concept using the resistant nullclines to determine therapy "on" and "off" times in order to trap the tumor in the controllable region 2, and reach approximate cycle that repeats back on itself in red. The continuous therapy is also plotted in dashed blue, for comparison. Two nullclines divide the triangle into 3 regions; region 1: $\dot{x}_R < 0$ for both therapy on and off; region $2:\dot{x}_R > 0$ for the rapy off and $\dot{x}_R < 0$ for the rapy on; region 3: $\dot{x}_R > 0$ for both therapy on and off. (C) Before chemotherapy, the healthy (H), sensitive (S), and resistant (R) populations compete on a dynamical fitness landscape, with several solution trajectories shown (black) and the instantaneous relative velocity indicated by background color gradient (red to blue). All internal trajectories lead to tumor growth and eventual saturation of the sensitive population (bottom left corner). Each population nullcline (line of zero growth: $\dot{x}_i = 0$ is plotted: healthy (dashed blue), sensitive (dashed red), and resistant (dashed green). The nullclines divide the triangle into 3 regions. Region 1: $\dot{x}_H > 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Region 2: $\dot{x}_H < 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Region 3: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R > 0$; (D) Chemotherapy alters the selection pressure to the disadvantage of chemo-sensitive cancer population and advantage of the healthy population (shown for $c = 0.6, \alpha = 0.020, \beta = 0.018, w = 0.1$). In this case, the nullclines divide the triangle into 6 regions; Region 1: $\dot{x}_H > 0$ $\dot{x}_S > 0 \ \dot{x}_R < 0$; Region 2: $\dot{x}_H > 0 \ \dot{x}_S < 0 \ \dot{x}_R < 0$; Region 3: $\dot{x}_H > 0$ $\dot{x}_S < 0$ $\dot{x}_R > 0;$ Region 4: $\dot{x}_H < 0$ $\dot{x}_S < 0$ $\dot{x}_R > 0;$ Region 5: $\dot{x}_H < 0$ $\dot{x}_S > 0$ $\dot{x}_R > 0$; Region 6: $\dot{x}_H < 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Solution trajectories (black) show initial trajectory toward healthy saturation (triangle top) but eventual relapse toward resistant population (bottom right of triangle) upon passing the resistant nullcline. .

ix

- 2.6The effect of dose on tumor relapse and progression free survival under continuous and adaptive therapy — (A) PSA data is from a single patient (see figure 2.3B) under continuous treatment (Mitoxantrone and Prednisone) is replotted (black dots) along with the best model fit (eqns. 2.2, 2.3) in blue. Using identical patient parameters, continuous treatment of two "new" drugs with higher effectiveness (i.e. increased effective dose, c) is shown in red and yellow. Time to relapse significantly increases with increasing dose while the progression free survival shows marginal, but decreasing, difference. An adaptive therapy (see figure 2.5B) is also simulated (solid black line), showing an increased control over the tumor; Similar control paradigms were simulated with error in administration times: slightly shorter time on therapy (1 day; green) and slightly longer (1 day; purple). (B) The same four therapies are shown in a fish plot. Continous therapies show relapse to initial tumor size is dominated by chemo-resistant population (green). The adaptive therapy successfully suppressed the growth of the resistant population (bottom).
- 3.1 Schematic dose schedules.(a) Maximum tolerated dose (MTD) schedule;
 (b) Low dose metronomic (LDM) schedule; (c) Adaptive schedule; (d) General time-dependent schedule accuated by piecewise constant dose concentrations.
- 3.2 The competitive release mechanism. (a) The three-component phase space associated with competing populations of (H, S, R) cells. (b) With no therapy, the S corner of the triangle is a globally attracting fixed point, while the H and R corners are unstable. All initial conditions lead to a saturated tumor. Filled circles are stable, unfilled circles are unstable. (c) For continuous chemotherapy above a threshold level, the R corner of the triangle is a globally attracting fixed point, while the H and S corners are unstable. All initial conditions (except those lying on the separatrix connecting the interior balanced fitness state to the S R side) lead to a resistant tumor.

29

- 3.4 Basins of attraction and nullclines. Panel showing the separatrices and nullclines $(dx_H/dt = 0, dx_S/dt = 0, dx_R/dt = 0)$ through the balanced fitness interior fixed point that determines the basin boundaries of attraction of the S state and the R states. The interior fixed point is the one in which each of the subpopulation fitness levels exactly matches the average fitness of the entire population. Filled circles are stable, unfilled circles are unstable. . . . 42

- 3.8 Constructing a closed loop trajectory using C = 0 and C = 0.5, with $C_{avg} = 0.366$.(a) By using segments of a trajectory for C = 0 and C = 0.5, switching values at points A and B, we construct a closed periodic orbit. (b) Using the same total dose, we show the MTD and LDM trajectories starting from point A. Both eventually move to the S corner, although initially the MTD trajectory moves toward the H corner (tumor regression) before recurrance. (c) The MTD, LDM, and adaptive schedules are depicted. (d) We plot the average fitness $\langle f \rangle$ for the three different chemo-schedules. The adaptive schedule is able to maintain a higher average fitness throughout. . .

- 3.12 Robust family of nested orbits. (a) One example of a (continuous) family of nested closed orbits using bang-bang strategy toggling between C = 0 and C = 0.7, all with similar total dose ($C \sim 0.4$). A closed loop can be made arbitrarily small around a center point described in Figure 3.13. (b) Schedule associated with orbits 1 and 4; (c) Schedule associated with orbits 1 and 5; (d) Schedule associated with orbits 1 and 6; (e) Schedule associated with orbits 2 and 4; (f)Schedule associated with orbits 2 and 5; (g) Schedule associated with orbits 2 and 6; (h)Schedule associated with orbits 3 and 4; (i)Schedule associated with orbits 3 and 5; (j) Schedule associated with orbits 3 and 6. 54

3.13 All closed orbit families are centered on dashed line defined by $f_1 =$ f_3 . (a) As a family of closed orbits enclose smaller and smaller loops, they intersect at a point of tangency. The point of tangency must lie on the dashed line. (b) Examples of several closed loop orbits centered at different points 554.1(a) e > 0 showing synergistic profiles; (b) e = 0 showing additive profiles; (c) 594.2Depiction of evolutionary stable states (ESS) for different chemotherapy values. (a) With no chemotherapy, the tumor saturates to the S corner regardless of the initial make-up of the three subpopulations. (b) With $C_1 = 0$, and $C_2 = 0.8$, competitive release of the resistant population R_1 drives all trajectories to the R_1 corner. (c) With $C_1 = 0.8$ and $C_2 = 0$, competitive release of the resistant population R_2 drives all trajectories to the R_2 corner. (d) Trajectories associated with three different constant combinations of C_1 and C_2 , depicting the overlap of the trajectories at different times. 67 Tumor growth curves (log-plots) for untreated and constant therapies. Tumor 4.3recurrence (dashed line) occurs at $t \sim 75$ dimensionless time units. 684.4 Panel showing the full range of trajectories, ESS and their basins of attraction for constant additive the rapies in the range $0 \le C_1 \le 1, 0 \le C_2 \le 1, C_1 + C_2 \le C_2 \le 1$ 1. Bifurcation values along top row are $C_2 = 1/3$, $C_2 = 1/2$. Bifurcation values along left column are $C_1 = 7/18$, $C_1 = 1/2$. 69 The effect of synergy (e > 0) vs. antagonism (e < 0) on the ESS and the basins 4.5of attraction for a fixed choice of constant combination therapies $C_1 = 0.35$, $C_2 = 0.27$. (a) e = -0.4; (b) e = -0.3; (c) e = -0.2; (d) e = -0.1; (e) e = 0; (f) e = 0.1; (g) e = 0.2; (h) e = 0.3; (i) e = 0.4. 70

4.6	The opening of the basin of attraction for R_1 via a transcritical bifurcation	
	at $e = -310/1701 = -0.1822$, $C_1 = 0.35$, $C_2 = 0.27$. At the bifurcation	
	point, the areas of the basins of attraction of the respective regions are $S =$	
	76.2%, $R_2 = 13.8\%$. (a) $e = -0.21$; (b) $e = -0.19$; (c) $e = -0.175$; (d)	
	$e = -0.15.\ldots$	72
4.7	Four ways of depicting the antagonistic transcritical bifurcation. (a) Eigen-	
	value collision that takes place as the two fixed points collide at the R_1 corner.	
	The other eigenvalue remains negative for both fixed points; (b) Transcritical	
	bifurcation diagram; (c) Pre-bifurcation phase plane; (d) Bifurcation phase	
	plane; (e) Post-bifurcation phase plane	73
4.8	The closing of the basin of attraction for S via a transcritical bifurcation at	
	$e = 30/189 = 0.1587, C_1 = 0.35, C_2 = 0.27$. At the bifurcation point, the	
	areas of the basins of attraction of the respective regions are $R_1 = 72.3\%, R_2 =$	
	17.7%. (a) $e = 0$; (b) $e = 0.1$; (c) $e = 0.2$; (d) $e = 0.3$.	74
4.9	Four ways of depicting the synergistic transcritical bifurcation. (a) Eigen-	
	value collision that takes place as the two fixed points collide at the S corner.	
	The other eigenvalue remains negative for both fixed points; (b) Transcritical	
	bifurcation diagram; (c) Pre-bifurcation phase plane; (d) Bifurcation phase	
	plane; (e) Post-bifurcation phase plane	75
4.10	Areas of the three basins of attraction as a function of the parameter e . Note	
	the sensitivity of the S and R_1 regions in the antagonistic regime $-0.2 \le e \le$	
	-0.1	76
4.11	(a) Average fitness curves for $e = 0.3, 0.15, 0, -0.15, -0.3$; (b) Tumor growth	
	curves (log-plot) for $e = 0.3, 0.15, 0, -0.15, -0.3$.	76

4.12	Closed loop adaptive schedules $OABO$ using three cycles for $e = 0, 0.3, -0.3$.	
	(a) $e = 0, OA$: $C_1 = 0.5, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0$	
	0.5; (b) Corresponding adaptive schedule; (c) $e = 0.3$, OA : $C_1 = 0.5$, $C_2 = 0.2$,	
	<i>AB</i> : $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.5$; (d) Corresponding adaptive	
	schedule; (e) $e = -0.3$, OA : $C_1 = 0.5$, $C_2 = 0.2$, AB : $C_1 = 0$, $C_2 = 0$, BO :	
	$C_1 = 0.2, C_2 = 0.5$; (f) Corresponding adaptive schedule.	77
4.13	Tumor growth curves for the adaptive schedules from figure 4.12 as compared	
	with untreated growth (exponential) and constant schedule which eventually	
	leads to tumor recurrence. (a) $e = 0$; (b) $e = 0.3$; (c) $e = -0.3$	77
4.14	Rates of adaptation for $e = 0, 0.3, -0.3$: (a) R_1 rate of adaptation; (b) R_2	
	rate of adaptation; (c) x_1 rate of adaptation; (d) x_2 rate of adaptation;	78
5.1	Schema for the Maximum tolerated Dose (MTD) and Low Dose Metronomic	
	(LMD) schedule. The integral of these two schedule are the same	88
5.2	Simulation results with $D = 0.3, T = 1, x(0) = 0.1$ and payoff matrix in	
	eq.5.43 (a) the corresponding trajectory for MTD and LDM simulated by	
	standard replicator equation.(b)the corresponding trajectory for MTD and	
	LDM simulated by Adjusted replicator equation	89
5.3	Bang-bang optimal control for both standard replicator equation and adjusted	
	replicator equation	89
5.4	Simulation results with $D = 0.3, T = 1, x(0) = 0.1$ and payoff matrix in	
	eq.5.44 with Adjusted replicator equation.	91
5.5	(a) optimal control $u^*(t)$ for adjusted replicator equation (b) the corresponding	
	trajectory $x^*(t)$	91
6.1	Ping-pong ball levitated by blowing through a straw. The streaming air blows	
	upwards along the positive x direction against the gravity. x^* marks the	
	equilibrium position where the forces balance. For full video, see https:	
	//www.youtube.com/watch?v=bCRjPFhlSYk	96

6.4 Examples of piecewise constant blowing schedules that move the ball from the origin to the target point $(x, v) = (x^*, 0)$. Parameter values are given by: $b \in [0, 6], \alpha = 1m^{-1}, g = 9.8m/s^2, x^* = log(2)meter, m = 1kg$, and we choose T = 1.34s for time scale. (a) Two-step schedule: b(t) starts with $b = \frac{\alpha x^* e^{\alpha x^*}}{e^{\alpha x^*} - 1}$, then it switches to $b = e^{\alpha x^*}$. (b) The corresponding trajectories for the twostep schedule. Note that the ball will remain at equilibrium on second step, so the trajectory only has one segment. (c) Three-step schedule with blowing on the first segment, followed by free falling (b = 0) for the second segment. (d) The corresponding trajectory for (c). The trajectories will get closer and closer to vertical axis with the increase of initial b(t). (e) Examples of general three-step schedule, with b_1 for first segment and b_2 for the second. (f) The corresponding b(t) for the examples of general three step schedule. (g) A schematic of a four-step schedule, where t_1 t_2 and t_3 mark three switching points.(h) The corresponding trajectory for (g) 1156.5The schematic of the trajectory driven by a general blowing function b(t)as indicated in blue. For comparison, we show what we call the *bang-bang* strategy, starting with maximum blowing b_{max} from t = 0 to $t = t_s$, then free fall b = 0 from $t = t_s$ to $t = t_f$. Any general trajectory associated with b(t) is trapped between the two segments of the other orbit. 116The bang-bang (on-off) strategy is optimal to attain the transfer in minimum 6.6time and using minimum air volume. (a) The actual trajectory for the transfer; (b) The time-dependent blowing schedule. 116Minimum energy strategy.(a) phase space. Simulated with parameters $b_{max} =$ 6.7 $20, \alpha = 1m^{-1}, x^* = ln(5)m, g = 9.8m/s^2, m = 1kg$ (b) corresponding schedule for (a) .(c) phase space simulated with parameters: $b_{max} = 5, \alpha = 1m^{-1}, x^* =$ $ln(5)m, g = 9.8m/s^2, m = 1kg.(d)$ corresponding optimal control and it's still bang-bang control. 117

Abstract

In this thesis, we use a three by three replicator dynamical system with a payoff matrix of Prisoner's Dilemma type to develop evolutionary game theory models of adaptive chemotherapy that avoid chemotherapeutic resistance in cancer. With the primary goal of controlling resistance and delaying tumor recurrence, we introduce and develop a trajectory design method from the orbital mechanics literature in order to construct families of closed loops, called 'evolutionary cycles', using adaptive time-dependent dose schedules. With timedependent adaptive chemotherapy control, we attempt to alter the fitness landscape of the heterogeneous cell populations in order to delay tumor recurrence. We show that for a three-component population made up of healthy cells, sensitive cells, and resistant cells, this method outcompetes both maximum tolerated dose (MTD) schedules and low-dose metronomic (LDM) schedules, which are standard currently used protocols. We then consider a more complex system made up of two types of resistant cells and two drugs applied sequentially and in combination. In this case, the drugs can interact synergistically, additively, or antagonistically. We show how to develop multi-drug schedules in this context using a similar method of adaptive trajectory design. Generally speaking, we show that antagonistic drug interactions are more effective at delaying tumor recurrence than synergistic ones which effectively cause the two drugs to act as one, limiting the flexibility in schedule design. We then apply the Pontryagin Maximum Principle to two by two evolutionary games. We show examples of payoff matrices where bang-bang control is optimal, and others where singular control techniques need to be applied. As a final chapter, we use a Hamiltonian model along with the trajectory design method to show how to levitate a mass with specially designed time-dependent blowing schedules.

Chapter 1

Introduction

My Ph.D. thesis develops techniques of adaptive control applied to evolutionary game theory[108, 85, 52] models of tumor growth and chemotherapeutic resistance[58]. We use the Prisoner's Dilemma (PD) game[66, 85] and the replicator dynamical system as a paradigm for the competition (cooperation or defection) among heterogeneous cell subpopulations in a tumor as the system undergoes evolution via natural selection. We use adaptive control[40] methods to shape the fitness landscape of the tumor and steer its evolutionary outcome. An overriding goal of our adaptive (time-dependent) chemotherapy schedules is to manage the development of chemo-resistance in a tumor. Although the models used in this thesis focus on the prevention of chemotherapeutic resistance in tumors, the techniques developed may also be useful in steering the evolution of microbial populations with the goal of mitigating and reversing antibiotic resistance, which has a large and specialized separate literature[58, 1, 17, 105, 39, 44, 27].

Because cancer is a rising threat to human health, it is important to understand the underlying biological and physical mechanisms of tumor growth and the complex pharmacological interactions between cancer cells and the toxins used to kill them. From the view of evolution, the proliferation of cancer cells in the patient is widely recognized as a process of natural selection [28]. As cancer cells proliferate and die, they can mutate, changing their rate of reproduction and ability to survive chemotherapy. Macroscopically, a tumor consists of multiple subpopulations of cells with varying growth rates as they compete for resources. The most aggressive subgroups have enhanced ability to gain more resources, grow faster and take over the majority of the population. The weaker subpopulations, on the other hand, are unable to reproduce as fast so their population size will remain low, reduced, or even be completely eliminated. The result, after months or years of growth, is a process of evolution by natural selection in a tumor.

To quantitatively model this evolutionary process, researchers used evolutionary game theory. An important first step for a mathematical model is to quantify the ability to reproduce, which is defined as fitness of a subpopulation of cells. Evolutionary game theory constructs the fitness landscape as the system evolves to an evolutionary stable state (ESS) [108, 52] based on the interaction of the subpopulation of cells. More specifically, the fitness function will have linear dependence [85] on the population size of surrounding cells, thus a frequency dependent selection. Compared with constant selection with fixed fitnesses, evolutionary game theory models the adaptation to the changing sizes of competing populations. There are many successful models that are based on either stochastic simulations with evolutionary games, such as the Moran process[69, 85], or deterministic simulations using the replicator equations[52] or Lotka–Volterra equations[72]. In this thesis, I will develop models of adaptive chemotherapeutic response of a tumor based on the replicator equations for Chapter 2-5. We will use techniques based on adaptive, time-dependent control of nonlinear dynamical systems.

An important mechanism we need to quantify is how tumors respond to chemotherapy. Tumor recurrence for a patient undergoing treatment is one of the main clinical challenges, and the primary reason of tumor regrowth is the development of chemotherapeutic resistance[58, 1, 17]. As an evolutionary process, the drugs will alter the fitness landscapes[120] by reducing the fitness of drug-sensitive cells. Initially, due to the reduction of sensitive cells, the regression of the tumor is observed quickly. After the sensitive cells are eliminated, the drug-resistant cells (or resistant cells) are favored by the selection process and the tumor will start to regrow. This process from treatment to regrowth of the tumor is called competitive release phenomenon in the ecological literature [44, 104, 74], because the absence of drug-sensitive cells contributes to the growth of drug-resistance cells as they are released from competition. In order to prevent the recurrence of the tumor, one approach is to find some adaptive chemotherapeutic schedule [40] for determining when to give the drug to the patients and how much drug should be administered. We use our evolutionary game theory model with the altered fitness from chemotherapy, to investigate the phase space of the system governing the interacting cells with different chemo-dose levels. In the phase space, we are able to construct the desired trajectory (steering and altering the balance of cells) from one point to another, and even design closed loops which keeps the cells in perpetual competition without any of the subpopulations saturating the tumor. The adaptive chemotherapy schedule corresponding to these closed loops can keep the tumor burden low and prevent or delay resistance from emerging. Compared with the classical Maximumtolerated-Dose method (On-off piecewise constant function) or Low-Dose-Metronomic (constant value) method, which will saturate with large tumor size within finite time, the adaptive therapy outperforms both of them by trapping the tumor size below a certain level. Discussions and details for this model are provided in Chapter 2. This method is then compared with patient data of tumor relapse in Chapter 3.

In Chapter 4, we extend our evolutionary game theory model to study multi-drug combination therapies with two types of resistant populations. The goal is to achieve an adaptive schedule of two drugs to avoid competitive release of both subpopulations of resistant cells. The existence of more than one drug, however, can result in drug-drug interactions. These interactions can be additive, synergistic, or antagonistic. In addition, with a two-drug schedule, there are many new possible combinations that emerge. The discussion of multidrug interactions goes back to the 1930s[14]. These works help to determine the cocktails of multiple drugs and reach the desired outcome for treatments[50, 57, 11, 100, 3, 25, 71]. Synergistic effects are commonly assumed to be the best if the goal is to kill the maximum amount of cancer cells since synergistic interactions are analogous to adding extra drug dose. However for an evolving tumor, maximum killing of sensitive cells enhances selection for resistant cells and can actually encourage tumor resistance. Although tumor regression often occurs early in the treatment cycle, tumor recurrence inevitably occurs and our goal is to delay this event to the maximum extent possible. After we design the adaptive control method based on the finding of closed loops in phase space, we show that antagonistic effects can provide better tumor regression for an evolving system as compared with the synergistic effects, mostly because the process of selection for the resistance population is slowed.

After we design feasible schedules of controlling the tumor size as mentioned above, the next goal is to optimize the different schedules. For example, we may aim at producing a schedule that results in the smallest possible tumor (volume) at the end of the treatment. The question then becomes an optimal control[66] problem in the replicator dynamical system. Typical methods for this problem involve Pontryagin Maximum Principle (PMP) developed in 1956 by the Russian mathematician Lev Pontryagin and his students [91, 98, 21, 113]. It is common to assume that the optimal control for a replicator equation is a bang-bang control [98, 21], while for stochastic simulation with Moran process, it is reported [122] to have a high entropy schedule (Low-Dose Metronomic schedule) to be optimal for minimal tumor volume. The continuous limit of the Moran process with infinite population size is proved [114, 115] to be the adjusted replicator equation, a rescaled form of replicator equation. In Chapter 4 we apply the maximum principle to a general 2-by-2 evolutionary game along with both the standard replicator equation and the adjusted replicator equation so we can compare with the stochastic model. We find that the behaviors of these two dynamic systems are quite different depending how the drug alters the fitness landscape. The optimal schedule for the standard replicator equation to minimize tumor size is bang-bang control, while for the adjusted replicator system, may not be piecewise constant if the drug favors one type of species instead of reducing the fitness of the other one. Details are presented in Chapter 5.

Finally, in Chapter 6 we use our time-dependent control method for a different problem of levitation of a mass by non-constant blowing schedules using the optimal control techniques. This work is inspired by the time-dependent optimal drug schedules previously used to prevent tumor recurrence. We present a model of a mass (say a ping-pong ball) blown by streaming gas and floating in air. We show that constant blowing can never levitate the ball from its starting point to a floating point in mid-air. Our goal is to design a time-dependent blowing schedule (analogous to our adaptive chemotherapy schedule) to levitate the ball to a target point that finds its rest point in minimal time. We prove that bang-bang control is optimal for minimizing time and streaming gas volume, while minimizing energy requires extra segments corresponding to the singular control from optimal control theory. The results are identical to the solution from Pontryagin Maximum principle. For more details about how the problem is analyzed and discussed, see Chapter 6.

Chapter 2

Capitalizing on Competition: An Evolutionary Model of Competitive Release in Metastatic Castration Resistant Prostate Cancer Treatment

The development of chemotherapeutic resistance resulting in tumor relapse is largely the consequence of the mechanism of competitive release of pre-existing resistant tumor cells selected for regrowth after chemotherapeutic agents attack the previously dominant chemosensitive population. We introduce a prisoner's dilemma game theoretic mathematical model based on the replicator of three competing cell populations: healthy (cooperators), sensitive (defectors), and resistant (defectors) cells. The model is shown to recapitulate prostate-specific antigen measurement data from three clinical trials for metastatic castration-resistant prostate cancer patients treated with 1) prednisone, 2) mitoxantrone and prednisone and 3) docetaxel and prednisone. Continuous maximum tolerated dose schedules reduce the sensitive cell population, initially shrinking tumor burden, but subsequently "release" the resistant cells from competition to re-populate and re-grow the tumor in a resistant form. The evolutionary model allows us to quantify responses to conventional (continuous) therapeutic strategies as well as to design adaptive strategies. These novel adaptive strategies are robust to small perturbations in timing and extend simulated time to relapse from continuous therapy administration.

2.1 Introduction

In his now classic 1961 study of competition for space between two species of barnacles in the intertidal zone off the Scottish coast, Joseph Connell [27] discovered something interesting. The blue barnacles *Balanus* normally occupied the intertidal zone, while the brown barnacles *Chthamalus* occupied the coast above high tide. Despite the commonly held belief that each occupied their own niche because of different adaptations to local micro-conditions, Connell hypothesized that the colonization of the intertidal zone by *Balanus* was actually preventing *Chthamalus* from inhabiting this region. To test this, he removed the blue barnacles from the intertidal zone and tracked the subsequent penetration of *Chthamalus* into this region. He concluded that *Chthamalus* had undergone *relief from competition* with *Bal*anus which allowed it to flourish where previously it could not. The point, he emphasized, was there was nothing *inherent* about the micro-environment of the intertidal zone that was preventing *Chthamalus* from occupying this region. It was simply the competition against a more dominant species that was holding it back. Without the presence of that species, Chthamalus happily claimed both zones as fundamental niches. Thus, the important notion of *competitive release* was formulated (see Grant [45]). When two (or more) sub-species compete for the same resources, with one species dominating the other, if the dominant species is removed, this can provide the needed release from competition that can allow the less dominant species to flourish. The mirror image of competitive release is the related notion of character displacement developed by Brown and Wilson [20] in which competition can serve to displace one or more morphological, ecological, behavioral, or physiological characteristics of two closely related species that develop in close proximity. These concepts are now well established as part of the overall framework of co-evolutionary ecology theory and play an important role in the evolution of chemotherapeutic resistance in cancer [74, 6, 46, 32].

Since co-evolution among competing clones and subclones is now a well established [86] process in malignant tumors, the mechanism of competitive release should be expected to play a role and affect the chemotherapeutic strategies one might choose to eliminate or control tumor growth. Indeed, tumor relapse and the development of chemo-therapeutic

resistance is now thought largely to be a consequence of the evolutionary mechanism of competitive release of pre-existing resistant cells in the tumor which are selected for growth after chemotherapeutic agents attack the subpopulation of chemo-sensitive cells which had previously dominated the collection of competing subpopulations. This concept, perhaps most forcefully advocated by Gatenby and collaborators [38], is gaining acceptance by clinicians. A recent (2012) systematic literature analysis of cancer relapse and therapeutic research showed that while evolutionary terms rarely appeared in papers studying therapeutic relapse before 1980 (< 1%), the language usage has steadily increased more recently, due to a huge potential benefit of studying therapeutic relapse from an evolutionary perspective [1]. Anticancer therapies strongly target sensitive cells in a tumor, selecting for resistance cell types and, if total eradication of all cancer cells is not accomplished, the tumor will recur as derived from resistant cells that survived initial therapy [88]. It is argued by Gatenby [38] that eradicating most disseminated cancers may be impossible, undermining the typical goal of cancer treatment of killing as many tumor cells as possible. The underlying assumption of this approach has been that a maximum cell-kill will either lead to a cure or, at worst, maximum life extension. Taking cues from agriculturists who have long abandoned the goal of complete eradication of pests in favor of applying insecticides only when infestation exceeds a threshold in the name of "control" over "cure," there are those who advocate for a shift from the cure paradigm in cancer treatments to a control paradigm [38, 12].

2.1.1 The likelihood of pre-existing resistance

Pre-existing resistant cells should generally be present in all patients with late-stage metastatic disease (for single point mutations which confer resistance), a conclusion supported by probabilistic models [17] and from tumor samples taken prior to treatment [61, 96] which have been reported for melanoma [56], prostate cancer [97], colorectal cancer [29, 62], ovarian cancer [103], and medulloblastoma [77]. According to this view, treatment failure

would not be due to *evolution* of resistance due to therapy, but rather the pre-existing presence of resistant phenotypes that are relatively sheltered from the toxic effects of therapy [32].

The likelihood of pre-existing resistance has important therapeutic implications. If we assume no pre-existing resistance, then most models predict maximum dose-density therapy will reduce the probability of resistance largely because this treatment minimizes the number of cell-divisions, thereby minimizing the risk of a mutation leading to acquired resistance [32]. By contrast, in pre-existing resistance scenarios, the maximum dose-density therapy strategy lends itself to competitive release due to the evolutionary nature of tumor progression. Most pre-clinical efforts that aim to maximize the short-term effect of the drug on sensitive cells does not significantly affect the long-term control of cancer [17]. This is because the phenomenon of competitive release can occur via the harsh selective pressure imposed by the tumor microenvironment after cancer therapies diminish the presence of the dominant (i.e. the chemo-sensitive) subpopulation. Additionally, the process of metastasis may allow a resistant subpopulation in the primary tumor to emerge elsewhere [116].

2.1.2 Leveraging the cost of pre-existing resistance

Pre-existing mutations that are responsible for conferring resistance may be associated with a phenotypic cost, or a reduced fitness, compared to the average fitness of the sensitive cell population [40, 41]. Even factoring in this fitness cost, deleterious mutations are still expected to be present in late-stage metastatic cancers [16]. This cost can come in many ways, such as an increased rate of DNA repair, or an active pumping out of the toxic drug across cell membranes. All of these strategies use up a finite energy supply that would otherwise be available for invasion into non-cancerous tissues or proliferation. The rapid removal of chemosensitive cells during therapy releases the resistant population from unwanted competition and thereby permits unopposed proliferation of the resistant cell population. In contrast, the goal of an adaptive therapy is to maintain a stable tumor burden that permits a significant population of chemo-sensitive cells for the purpose of suppressing the less fit but chemoresistant populations, consistent with the philosophy that it takes an evolutionary strategy to combat an evolving tumor [116]. There is some evidence that adaptive therapies may be beneficial for a range of fitness cost (or lack thereof) and therapeutic sensitivity [36, 49].

A theoretical framework for these adaptive therapies first developed by Gatenby [40], leverages the notion that pre-existing resistance is typically present only in small population numbers due to a cost of resistance. This less fit phenotype is suppressed in the Darwinian environment of the *untreated* tumor but treatments that are designed to kill maximum numbers of cells remove the competition for the resistant population and ultimately select for that population during tumor relapse¹. While the goal of an adaptive therapy (to capitalize on the competition resistant subpopulations incur through maintaining a stable sensitive cell population) has gained some level of acceptance, the ideal adaptive therapy schedule in practice, for different tumor types and growth rates is far from settled. Gatenby's paper modulates dose density only, while stating that an ideal adaptive therapy would also modulate the drug, as well as the dose schedule (both dose and density) [40]. Some of these evolutionary ideas were tested experimentally using mouse models with modulated dose strength and dose vacations designed to maintain a stable, controllable tumor volume [31, 105]. This twophase adaptive therapy involved an initial high-dose phase to treat the exponential growth of the tumor and a second phase designed for stable tumor control using a variety of strategies (such as decreasing doses or skipping doses when stability is achieved). Several spatial, agent-based computational models have modulated dose strength with respect to a threshold value of tumor size (a fraction of the original tumor burden) [36, 8]. Findings suggest that adaptive therapies based on evolutionary treatment strategies that maintain a residual population of chemo-sensitive cells may be clinically viable, and is currently extended to an on-going clinical trial (NCT02415621) which adaptively controls the on/off cycling of

¹It is important to note that both high-dose, maximum tolerated dose schedules and low-dose, metronomic dose schedules have this cumulative goal of achieving maximum cell-kill over the course of many cycles of treatment.

abiraterone [134]. We refer the reader to similar approaches used in mitigating antibiotic resistance [4, 48].

With these advances in mind, the goal of this manuscript is to use the evolutionary framework introduced and advocated over the past 10 years [40, 38, 36, 8, 134] to mathematically model the important parameters of competitive release and use that framework to better understand therapeutic implications of the cost of developing resistance and to learn how to exploit competition between subpopulations. Here, we ignore clonal dynamics, mutations between subpopulations, and non-genetic (rapid, bidirectional) state transitions, to focus on the inter-subpopulation competition [106, 51, 90]. Specifically, we introduce a three-component replicator system with a prisoner's dilemma payoff matrix [119] to model the three relevant subclonal populations: healthy cells (H), sensitive cells (S), and resistant cells (R). We consider healthy cells as the surrounding non-neoplastic cells that do not contribute to the overall growth of the tumor, but nevertheless compete with neoplastic cells for space and resources. Using the nullcline information in a triangular phase plane representation of the nonlinear dynamics of the system, we first show the essential ingredients that render competitive release possible. Then, using the parameters that control selection pressure (hence relative growth rates) on the three subclonal populations, we attempt to maintain the tumor burden at low levels so that the resistant population does not reach fixation.

2.1.3 Retrospective analysis of metastatic castration-resistant prostate cancer

A recent retrospective analysis of tumor measurement data (PSA levels) from eight randomized clinical trials with metastatic castration-resistant prostate cancer (mCRPC) used a simple linear combination of exponentials model to estimate the growth and regression rates of disease burden over time [127]. In total, over 67% of patients were fit to models with a positive regrowth rate, indicating failure due to resistance. Prostate-specific antigen (PSA) measurement data for patients in each treatment silo (prednisone only, mitoxantrone



Figure 2.1: Schematic of competitive release in a tumor — (a) Prior to treatment, a tumor consists of a large population of sensitive cells (red) and a small population of less fit resistant cells (green) competing for resources with the surrounding healthy cells (blue); (b) Chemotherapy targets the sensitive population (middle), selecting for the less fit resistant population that thrives in the absence of competition from the sensitive population; (c) Upon regrowth, the tumor composition has larger numbers of resistant cells, rendering the subsequent rounds of treatment less effective.

and prednisone, docetaxel and prednisone) were obtained through the Project Data Sphere open data portal (https://www.projectdatasphere.org), and we show that this model is able to adequately fit data for each treatment type with the additional capability of allowing us to track responses to conventional therapeutic strategies and design new adaptive strategies as the tumor evolves. The model can be used to test previously proposed adaptive therapies, but we propose a novel schedule utilizing quantitative tools from nonlinear dynamical systems theory which use the current global state of the nonlinear replicator system with respect to the nullcline curves of the equations as well as parameters controlling relative fitness levels of the competing sub-populations. The simulated chemotherapeutic strategies that we implement, based on tracking the phase-space structure of of replicator system, are ones that can adapt on the same timescale as the inherent timescale of evolution of the subpopulations comprising the tumor, i.e. are as dynamic as the tumor. While this approach cannot be preplanned by the oncologist at the beginning of therapy like classical strategies, we provide discussion to explain how an evolutionary game theory model describing the fitness landscape (described below) is useful to understand the underlying features of a dynamical fitness landscape associated with a cost to resistance: a three-way prisoner's dilemma. Specifically, the model indicates a boundary over which an adaptive therapy will cease to be effective. Our model focuses on cost to resistance, as opposed to specific mechanisms of resistance (i.e. hormone resistance, for example). This has the advantage in some ways as being agnostic to resistance mechanisms, but on the other hand, in this manuscript, we can not distinguish between different mechanisms: the analysis explores the effect of a fitness competition between subpopulations on the competitive release phenomenon. Three common therapies used in mCRPC do show that a fit of our model to prostate data lead to the *emergence* of a cost of resistance in a majority of patients, without predetermining a value to the cost in the model *a priori*. Similarly, ref. [127] found that continuous administration of different therapies (i.e. different resistance mechanisms) is associated with slow tumor regrowth rates (i.e. a cost of resistance conferred to growth potential). We now describe the model.

2.2 Materials and Methods

A schematic of a finite-population three component model (healthy, sensitive, resistant) of competitive release is shown in figure 2.1, where the tumor consisting of sensitive and resistant cells is competing with the surrounding healthy tissue. At diagnosis (see figure 2.1, left), the tumor is dominated by sensitive cells (red) which out compete the surrounding healthy population (blue) during unhindered tumor progression. A small portion of resistant cells (green) remains in small numbers, suppressed by the larger sensitive population. After several rounds of chemotherapy, the tumor shrinks, leaving the resistant population largely unaffected (figure 2.1, middle). Inevitably, the tumor relapses due to the small number of



begin continuous chemotherapy treatment

Figure 2.2: Clonal evolution of competitive release — A fishplot (sometimes known as a Müller plot), showing the tumor quantity (vertical axis) and composition (sensitive: red; resistant: green) over time (horizontal axis, left to right) with important events annotated. After first driver mutation (left), initial exponential growth of sensitive population occurs until diagnosis (dashed line). Continuous therapy targeting the chemo-sensitive population responds well with a decrease in tumor burden. In the absence of sensitive cells, the resistant population (existing in small numbers before the start of therapy) grows to become the dominant subpopulation at relapse, albeit typically with lower exponential growth rate due to the cost of resistance.

sensitive cancer cells remaining after therapy (figure 2.1, right). In the absence of competition from the dominant sensitive population, the resistant cells grow unhindered, rendering subsequent rounds of chemotherapy less effective. Subsequent application of identical therapies will have a diminished effect. Figure 2.2 shows the process in a 'Müller fishplot', which we will use later to track the subclonal populations. This representation was first utilized in cancer to compare modes of clonal evolution in acute myeloid leukemia (see [30]) and is particularly useful in computational models where every cell type can be tracked in time. A fishplot shows the tumor burden (vertical axis) over time (horizontal axis). The first resistant cell is assumed to arise from a single sensitive cell, undetectable by current methods of measurement. The schematic depicts unhindered tumor growth after the first driver mutation (figure 2.2, left) where the tumor grows exponentially before diagnosis, during which time a resistant mutation occurs (figure 2.2, middle). After diagnosis (dashed line), a regimen of continuous chemotherapy shows initial good response and tumor regression, but the resistant population grows back (although at a slower growth rate) unhindered by competition, leading to relapse (figure 2.2, right). Previously, a linear combination of exponentials model has been proposed to track the tumor quantity via prostate-specific antigen (PSA) measurement data [23]. The tumor quantity q(t) is a function of the exponential death rate of the sensitive cells, d, the exponential growth rate of the resistant cells, g, and the initial fraction of resistant cells, f [17]. In this equation, q(t) is the tumour burden at time t in days, normalized to the PSA level at t=0, (see figure 2.3, blue lines). The model can be written as follows:

$$q(t) = (1 - f)e^{-dt} + fe^{gt}.$$
(2.1)

This model, shown to be a reasonably good description of the changing tumor size during therapy for colorectal, prostate, and multiple myeloma cancers, identifies the important parameters in competitive release: initial fractional resistance (f), and birth/death rates (g,d) for the resistant and sensitive populations, respectively. The model is used to fit prostate-specific antigen (PSA) measurement data from retrospective analysis of three randomized clinical trials with metastatic castration-resistant prostate cancer to estimate the growth (g) and regression (d) rates of disease burden over time. Four representative patients are chosen from the control arms of each randomized trial and shown in figure 2.3: treatment with prednisone only [75] (left column: figures 2.3B,E,H,K); treatment with mitoxantrone and prednisone [112] (middle column: figures 2.3C,F,I,L) and treatment with docetaxel and prednisone [89] (right column: figures 2.3D,G,J,M). PSA data and model fits are normalized by q(t = 0) (black dots) and exponential fits are shown in blue.

Despite the fact that the model (2.1) curve-fits data reasonably well (labeled "exp." in figure 2.3), it contains no evolutionary information or concepts, a keystone principle behind competitive release. We also include in the figure 2.3 fits of our model presented here (eqns. 2.2, 2.3) in red (labeled "rep." in figure 2.3, where $\tilde{q}(t)$ is given by eqn. 2.4, below). Our evolutionary model is able to capture similar trends as the exponential model of equations (eqn. 2.1) but has the important property of allowing us to calculate cost of resistance associated with model fit. While each drug (columns) may have very different resistance *mechanisms*, ten of the twelve patients shown have a cost of resistance ($\alpha - \beta$, described below). In the last section we will describe the dynamical phase portrait and implement adaptive strategies to capitalize on competition between resistant and sensitive subpopulations in order to maintain tumor control.

2.2.1 The replicator equation model

The dynamics of the fitness landscape of three competing cell types are described by the replicator equation (see [114]), which is a deterministic birth-death process in which birth and death rates are functions of cell fitness, and cell fitness is a function of prevalence in the population. Each *i*th cell type (i = 1, 2, 3) competes according to equation 2.2, where x_1 , x_2 , x_3 are the corresponding frequency of healthy (H), sensitive (S) and resistant (R) cells, respectively, such that $\sum_i x_i = 1$.

$$\dot{x}_i = (f_i - \langle f \rangle) x_i \tag{2.2}$$

$$f_i = 1 - w_i + w_i (A\vec{x})_i \tag{2.3}$$

Here, \vec{x} is the vector $\vec{x} = (x_1, x_2, x_3)^T$ and $(Ax)_i$ is the *i*th element of vector Ax. The prevalence of each sub-population, x_i , changes over time according to the changing population fitness, f_i , as compared to the average fitness of all three populations $\langle f \rangle = f_1 x_1 + f_2 x_2 + f_3 x_3$. If the fitness of the sub-population is greater than the average fitness $(f_i - \langle f \rangle > 0)$, that sub-population grows exponentially, whereas if it is less $(f_i - \langle f \rangle < 0)$, it decays.

In order to directly compare results to previously published models of tumor quantity, q(t), during competitive release (i.e. equation 2.1), the tumor quantity can be written as the sum of the resistant and sensitive tumor populations (x_2, x_3 respectively), $\tilde{q}(t)$,
$$\tilde{q}(t) = x_2(t) + x_3(t),$$
(2.4)

normalized by it's initial condition, $\tilde{q}(0)$. This is shown in a schematic in figure 2.3A. These represent the same subpopulations shown in green (resistant) and red (sensitive) in figure 2.2.

Before therapy, each subpopulation (healthy, chemo-sensitive, and resistant cells) the selection pressure is constant across all cell types (i.e. $w_i \equiv w, i = 1, 2, 3$) at a level that represents the natural selection pressure the tumor environment imposes on the different subpopulations. These values discussed in the literature are typically small, in the range $w_i \equiv w \approx 0.1 - 0.3$. This is a weighting parameter between neutral dynamics (w = 0) and fitness-dependent dynamics (w = 1). Typically in replicator systems, this parameter scales out with time. In our model, we treat each subpopulation selection parameter independently. We implement chemotherapy in our model by changing the selection pressure parameters on each of the subpopulations of cells. Therapy can be administered at different doses (i.e. values of the drug concentration: $c; 0 \le c \le 1$). A higher value of c indicates a stronger dose of chemotherapy drug (described in more detail in [122]). This follows the schematic in figure 2.4 which depicts the change in the fitness landscape before and after therapy. In figure 2.3, dose concentration is assumed to be constant for a specific drug while patient-specific parameters are the selection pressure (w), cost of resistance (discussed below), and initial fraction of resistant cells (f). Values are altered as follows (see figure 2.4 for explanation of changing fitness landscape):

$$w_1 = (1+c)w \quad \text{(healthy)} \tag{2.5}$$

$$w_2 = (1-c)w \quad \text{(sensitive)} \tag{2.6}$$

$$w_3 = w$$
 (resistant) (2.7)

The fitness landscape (eqn. 2.3) is described in detail by the entries of the payoff matrix A,

$$A = \begin{array}{ccc} H & S & R \\ A = \begin{array}{ccc} H & \begin{pmatrix} a & b & o \\ \\ S & \\ R & \begin{pmatrix} h & j & k \\ \\ l & m & n \end{pmatrix} \end{array}$$
(2.8)

where each pairwise cell-cell interaction is described by the row and column values, which are parameters in the fitness equation (2.3). The 3×3 payoff matrix is constructed as a standard prisoner's dilemma matrix where healthy cells are the cooperators and sensitive/resistant cells are the defectors. In the absence of resistant cells, we describe in more detail the dynamics of the system and its relevance to tumor growth in the appendix. The key features are the Gompertzian growth of the cancer cell population which saturates at a lower overall population fitness level than the initial all healthy cell population. This necessitates the following inequalities of the payoff matrix below (eqn. 2.8): h > a > j > b, l > a >n > o, and k > n > j > m. Note that these payoff entries remain constant before and during therapy: chemotherapy is viewed as affecting the selection balance among the three subpopulations. This is implemented in our model by a change in the the selection pressure parameters only (see eqns. 2.5, 2.6, 2.7). More discussion of why the prisoner's dilemma matrix, which models the evolution of defection, is a useful paradigm for cancer can be found in [119, 118] and the appendix.

2.2.2 The linearized system and the cost of resistance

An additional important feature of the payoff matrix is the notion of the *cost of resistance* which is highlighted in figure 2.4A. With no therapy, the sensitive cells exhibit fastest growth due to their higher fitness value relative to both the resistant population and the healthy population. The difference between the baseline fitness values of the sensitive cells and the resistant cells can be thought of as the 'price paid' by the resistant population to retain their resistance to toxins. This cost, in our model, is quantified as the difference in the (linearized)

growth rates of the two populations (type 2: sensitive; type 3: resistant). Linearizing eqn (2.2), (2.3), (which form a cubic nonlinear system if expanded out) gives rise to the sensitive-resistant uncoupled system:

$$\dot{x}_2 = \alpha x_2 \tag{2.9}$$

$$\dot{x}_3 = \beta x_3, \tag{2.10}$$

with the growth parameters:

$$\alpha = w_1(1-a) + w_2(h-1) \tag{2.11}$$

$$\beta = w_1(1-a) + w_3(l-1). \tag{2.12}$$

Using eqns (2.5), (2.6), (2.7) gives:

$$\alpha = w(h-a) - cw(h+a-2)$$
(2.13)

$$\beta = w(l-a) + cw(1-a).$$
(2.14)

With no therapy, c = 0, we have:

$$\alpha = w(h-a) \tag{2.15}$$

$$\beta = w(l-a). \tag{2.16}$$

We call the fitness cost of resistance the difference between these growth rates with no therapy, hence $(\alpha - \beta) = w(h - l)$, where we require h > l in the payoff matrix. However, the *w* parameter is patient-specific, allowing each patient a varied cost.

2.3 Results

It is useful to view the nonlinear dynamical trajectories of the system using the trilinear coordinates shown in figure 2.5A, which gives a representation of the clonal phase space for every possible value of \vec{x} [102]. The corners represent saturation of a single cell type (e.g. the top corner represents $\vec{x} = [1, 0, 0]$, or all healthy cells). The healthy (H; top corner), sensitive (S; bottom left corner), and resistant (R; bottom right corner) populations compete according to equation (2.2) and follow trajectories shown (black) in figure 2.5C.

Figure 2.5C shows the complete dynamical information for untreated patient dynamics $(c = 0, \alpha = 0.020, \beta = 0.018, w = 0.1)$. All internal trajectories (black lines) lead to tumor growth and eventual saturation of the sensitive population (bottom left corner, "S"). Instantaneous relative velocity is indicated by background color gradient (blue to red). Subpopulation "nullclines," the curves for which $\dot{x}_i = 0$ are shown for healthy (dashed blue), resistant (dashed green), and sensitive (dashed red). On one side there is positive growth ($\dot{x}_i > 0$); on the opposite side negative growth ($\dot{x}_i < 0$). These nullclines delineate the phase space into three regions.

- Region 1: $\dot{x}_H > 0 \ \dot{x}_S > 0 \ \dot{x}_R < 0;$
- Region 2: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R < 0;$
- Region 3: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R > 0$.

Likewise, figure 2.5D shows the complete dynamical information for the patient treated with chemotherapy ($c = 0.6, \alpha = 0.020, \beta = 0.018, w = 0.1$). Trajectories initially move away from the sensitive/resistant corners toward healthy, but subsequently relapse to the saturation of the resistant subpopulation. The phase space for treated dynamics is divided into six regions by the three nullclines.

- Region 1: $\dot{x}_H > 0 \ \dot{x}_S > 0 \ \dot{x}_R < 0;$
- Region 2: $\dot{x}_H > 0 \ \dot{x}_S < 0 \ \dot{x}_R < 0;$
- Region 3: $\dot{x}_H > 0 \ \dot{x}_S < 0 \ \dot{x}_R > 0;$
- Region 4: $\dot{x}_H < 0 \ \dot{x}_S < 0 \ \dot{x}_R > 0;$

- Region 5: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R > 0;$
- Region 6: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R < 0.$

An important detail emerges from the model: during treatment, the resistant nullcline is reached before the healthy nullcline (figure 2.5D). For example, a tumor with an initial diagnosis in region 2 (see figure 2.5D) can be expected to respond to treatment. The trajectory will follow along the black line towards the healthy corner (i.e. decreasing tumor quantity) until the trajectory passes over the nullcline for the resistant subpopulation (dashed green line; $\dot{x}_R = 0$) reaching region 3. Subsequently, the trajectory passes the healthy nullcline (dashed blue line; $\dot{x}_H = 0$) and the tumor relapses ($\dot{x}_H < 0$), this time toward saturation of the resistant subpopulation.

If the desire is to maintain treatment until the point of positive progression, this occurs when the trajectory is far past the resistant nullcline. Appearances (based on tumor burden) can be deceiving – while the tumor may appear to be responding, the overall state may be well past the point of no return, and the resistant population is preparing to re-populate.

2.3.1 Managing competitive release

Figure 2.5B shows a schematic of the two resistant nullclines from 2.5C and 2.5D (untreated: solid green line; treated: dashed green line) overlaid on the same phase portrait. From the initial condition (purple dot) treatment begins as the trajectory (solid blue line) approaches the nullcline, positive resistant growth can be avoided by a well-timed drug holiday (solid red line). Here, we propose an adaptive therapy with the goal of trapping the trajectory between the treated resistant nullcline and the untreated resistant nullcline, creating an orbit in a closed loop for a finite period of time to 'steer' tumor evolution in region 2 of figure 2.5B. We do this by altering the dose concentration parameter c (a parameter that can be accessed clinically) in eqns (2.5) in an off-on (bang-bang) fashion, (eqn. 2.6, 2.7) between the two resistant nullclines: the untreated nullcline from figure 2.5C and the treated nullcine from 2.5D. While all details of the 'tumor phase space' may not yet be directly measurable in the clinic, we propose that all successful adaptive therapies will operate in regions 1 and 2 in figure 2.5D; otherwise they ultimately will not be successful. In this way, meaningful insight is gained into the dynamics behind the cost to resistance, regardless of the mechanism of that resistance. Previous adaptive therapy schedules (described in the introduction; see [36, 8, 31, 134]) have benefited from not crossing this "hidden" resistant nullcline. We now propose one example of an adaptive therapy schedule which actively captures and uses information from the dynamic phase space.

The simple control paradigm proposed *indirectly* controls the resistant population by systematically choosing when to administer therapy and when to give drug holidays. These holidays allow a sufficient number of sensitive cells to remain in order to suppress the resistant population. A continuous dose of therapy is administered until the treated nullcline ($\dot{x}_R = 0$) is reached (see figure 2.5D, green dashed line). This is the starting point of positive growth for the resistant population (further therapy would result in $\dot{x}_R > 0$). A drug holiday is then imposed until the second nullcline is reached (see figure 2.5C, green dashed line). The sensitive population is allowed to regrow until it is large enough to suppress the resistant population once again (and when $\dot{x}_R = 0$). Therapy is administered to allow the tumor to cycle back and forth between the two nullclines. This bang-bang (on-off) strategy allows an extension of relapse times. We emphasize that the specific times we turn the therapy on and off in this bang-bang strategy cannot be pre-planned, but depend on the position of the trajectory in the tumor phase space as the disease evolves.

Next, we measure the success of this control paradigm compared to continuous treatment via two important measures of therapy effectiveness: progression free survival (PFS) and time to relapse. PFS is the time to which the patient has a measured response, before relapse of the tumor: shown in figure 2.6A, bottom. A more useful metric is "time to relapse," where disease reaches initial state. Measuring the effectiveness of a chemotherapy schedule based on the killing rate or progression free survival alone are not sufficient predictive measures of long-term cancer control [17]. PSA data from a single patient who relapsed due to treatment resistance (figure 2.3B) is replotted in figure 2.6A (black dots) along with the model best fit of continuous treatment (blue; therapy 1). Also shown are simulated new drugs with increased effectiveness of killing sensitive cells (red, yellow) simulated by increasing the effective dose concentration with identical patient-specific parameters (from the patient in figure 2.3B). Each increased dose corresponds to a slightly shorter PFS, but an increased time to relapse to the initial tumor burden. However, despite the increase in relapse times, none of these doses optimizes tumor control, as seen in the fishplots (figure 2.6, B,C,D). At the point of relapse to the initial tumor burden, the tumor is dominated by the presence of resistant subpopulations (green), rendering future treatments ineffective. Oftentimes, the effectiveness of a new chemotherapy drug is be determined by PFS times when drugs that have high killing rates of sensitive cells may have *shorter* times to progression and lower total tumor burden at all times (everything else equal). The figure clearly shows that all treatments have similar progression free times but with a greater range of relapse times (even though continuous treatment always eventually leads to relapse).

Next, compare continuous treatments to the proposed control paradigm (solid black line) for identical initial conditions and identical drug dose. As seen in the fishplot (figure 2.6E), the resistant population (green) is suppressed during drug holidays, leading to an extended time without relapse. This adaptive technique is successful for two reasons. First, the drug holidays allow an adequate sensitive population size to suppress the growth of the lower-fitness resistant population. Second, the resistant population is never allowed to reach a positive growth under treatment ($\dot{x}_R > 0$), therefore cannot take over the tumor cell population. In order to test the robustness of this method, similar control paradigms were simulated with error in administration times: slightly shorter time on therapy (1 day; figure 2.6A; green) and slightly longer (1 day; figure 2.6A, purple). All else equal, it is better to err on the side of longer treatment on periods without decreasing off treatment time (purple).

2.4 Discussion

The chemotherapeutic scheduling strategies outlined in this paper cannot be pre-planned by the oncologist at the beginning of therapy like classical strategies [84], as they rely on significant decision making and continuous monitoring of the different subpopulations of cells that co-evolve as the tumor progresses. This means that the quality of the cell population monitoring system is crucial to the entire strategy, as has been pointed out in [33]. There can be no adaptive tumor control strategy without continuous monitoring of the subpopulations as it is not the tumor burden that is of primary interest, but the heterogeneous balance of the subpopulations comprising the tumor. In addition, the information gleaned from a detailed monitoring system cannot be acted upon unless the various administered drugs are sufficiently targeted to act efficiently and exclusively on specific subpopulations. These two systems must be in place (sensing and actuating) in order to successfully shape the fitness landscape and steer a growth trajectory in a desired direction. We also want to emphasize a separate point, which is that it is not enough to know in detail the *current* state of the system in order to steer it successfully. One must also have a description of all possible *nearby* states of the system, both under therapeutic pressure and without therapy. Better yet is to have a global picture of *all* possible states of the system, with nonlinear nullcline information, as one would obtain by analyzing the full phase space of the entire system. With this information, one would know *where* to steer the system to get to a desired state, even if one does not know how to achieve this (clinically). In current state-of-the-art medical practice, such sophisticated sensor-actuator capability is not yet sufficiently developed as it is in many engineering contexts where versions of adaptive control theory are routinely used. Many similar challenges, and the necessary steps towards their implementation, present themselves in the ecology and pest control communities, and we point to Gould's article [44] for a nice early overview. More recently, connections between the approaches developed in the past by ecologists and possible future strategies for oncologists have been discussed by Gatenby and collaborators [39]. Other groups [64, 18, 65] have also developed highly mathematical approaches to tumor control from different points of view. Clearly not all of the clinical steps are in place to effectively test and implement many of the strategies that have been explored theoretically. Yet it is still important to continue to develop the kinds of mathematical models and computer simulations that would serve to identify the many possible schemes, parameter ranges, and sensitivities that could be tested via clinical trials that focus on adaptive therapies with the goal of suppression of potential evolution of resistance.



Figure 2.3: Model fits of PSA measurement data for metastatic castrationresistant prostate cancer — Prostate-specific antigen (PSA) measurement data of four representative patients from three randomised clinical trials with metastatic castrationresistant prostate cancer. (A) A single patient data shown with model fit (eqn. 2.4) of growing resistant population (green), decaying sensitive population (red), where the sum represents an overall initial good response and eventual relapse (q(t), dashed line). Left column: treatment with prednisone only [75]; middle column: treatment with mitoxantrone and prednisone [112]; right column: treatment with docetaxel and prednisone [89]. PSA data is normalized by q(t = 0) (black dots). The data is fit using the exponential model (eqn. 2.1; blue curve) and the replicator model (eqns 2.2, 2.3). Each patient is fit reasonably well with both models. Data was fit by parameter sweep of cost ($\alpha - \beta$, eqn. 2.15, 2.16), initial fractional resistance f and selection pressure w. Parameters used are $\alpha - \beta = [0.02, 0.07, 0.20, 0.00, 0.20, 0.00, 0.20, 0.03, 0.20, 0.01, 0.15, 0.19]; <math>f = [0.20, 0.02, 0.09, 0.05, 0.03, 0.07, 0.10, 0.10, 0.09, 0.10, 0.03, 0.03]; <math>w = [0.10, 0.20, 0.20, 0.30, 0.20, 0.01, 0.30, 0.15, 0.10, 0.15, 0.20]$ for B - M respectively.



Figure 2.4: Fitness landscape before and during therapy — (A) Before therapy, a driver mutation leads to a fitness advantage of the cancer cell (red) and a subsequent resistant-conferring mutation comes at a fitness cost (green). Parameters determined by the prisoner's dilemma payoff matrix reflect these relative fitness differences. (B) Rather than change the elements of the game payoff matrix directly, the selection parameter is manipulated for healthy (increased) and cancer (decreased) subpopulations such that $w_1 - w_2 = 2c$. (C) Changes in w_i results in a new relative fitness of each subpopulation during therapy. The fitness of the resistant population is unaffected by therapy's selective pressure, but the healthy population is given an advantage over the chemo-sensitive population.



Figure 2.5: Dynamic phase portraits before and during chemotherapy - (A) Trilinear coordinate phase space representation; (B) Schematic of proposed adaptive therapy concept using the resistant nullclines to determine therapy "on" and "off" times in order to trap the tumor in the controllable region 2, and reach approximate cycle that repeats back on itself in red. The continuous therapy is also plotted in dashed blue, for comparison. Two nullclines divide the triangle into 3 regions; region 1: $\dot{x}_R < 0$ for both therapy on and off; region $2:\dot{x}_R > 0$ for the rapy off and $\dot{x}_R < 0$ for the rapy on; region 3: $\dot{x}_R > 0$ for both the rapy on and off. (C) Before chemotherapy, the healthy (H), sensitive (S), and resistant (R) populations compete on a dynamical fitness landscape, with several solution trajectories shown (black) and the instantaneous relative velocity indicated by background color gradient (red to blue). All internal trajectories lead to tumor growth and eventual saturation of the sensitive population (bottom left corner). Each population nullcline (line of zero growth: $\dot{x}_i = 0$) is plotted: healthy (dashed blue), sensitive (dashed red), and resistant (dashed green). The nullclines divide the triangle into 3 regions. Region 1: $\dot{x}_H > 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Region 2: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R < 0$; Region 3: $\dot{x}_H < 0 \ \dot{x}_S > 0 \ \dot{x}_R > 0$; (D) Chemotherapy alters the selection pressure to the disadvantage of chemo-sensitive cancer population and advantage of the healthy population (shown for $c = 0.6, \alpha = 0.020, \beta = 0.018, w = 0.1$). In this case, the nullclines divide the triangle into 6 regions; Region 1: $\dot{x}_H > 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Region 2: $\dot{x}_H > 0 \ \dot{x}_S < 0 \ \dot{x}_R < 0$; Region 3: $\dot{x}_H > 0 \ \dot{x}_S < 0 \ \dot{x}_R > 0$; Region 4: $\dot{x}_H < 0 \ \dot{x}_S < 0 \ \dot{x}_R > 0$; Region 5: $\dot{x}_H < 0$ $\dot{x}_S > 0$ $\dot{x}_R > 0$; Region 6: $\dot{x}_H < 0$ $\dot{x}_S > 0$ $\dot{x}_R < 0$; Solution trajectories (black) show initial trajectory toward healthy saturation (triangle top) but eventual relapse toward resistant population (bottom right of triangle) upon passing the resistant nullcline.



Figure 2.6: The effect of dose on tumor relapse and progression free survival under continuous and adaptive therapy — (A) PSA data is from a single patient (see figure 2.3B) under continuous treatment (Mitoxantrone and Prednisone) is replotted (black dots) along with the best model fit (eqns. 2.2, 2.3) in blue. Using identical patient parameters, continuous treatment of two "new" drugs with higher effectiveness (i.e. increased effective dose, c) is shown in red and yellow. Time to relapse significantly increases with increasing dose while the progression free survival shows marginal, but decreasing, difference. An adaptive therapy (see figure 2.5B) is also simulated (solid black line), showing an increased control over the tumor; Similar control paradigms were simulated with error in administration times: slightly shorter time on therapy (1 day; green) and slightly longer (1 day; purple). (B) The same four therapies are shown in a fish plot. Continous therapies show relapse to initial tumor size is dominated by chemo-resistant population (green). The adaptive therapy successfully suppressed the growth of the resistant population (bottom). Parameters used: $\alpha - \beta = 0.02$; w = 0.10; f = 0.2

Chapter 3

Nonlinear adaptive control of competitive release and chemotherapeutic resistance

We use a three-component replicator system with healthy cells, sensitive cells, and resistant cells, with a prisoner's dilemma payoff matrix from evolutionary game theory to model and control the nonlinear dynamical system governing the ecological mechanism of competitive release by which tumors develop chemotherapeutic resistance. The control method we describe is based on nonlinear trajectory design and energy transfer methods first introduced in the orbital mechanics literature for Hamiltonian systems. For continuous therapy, the basin boundaries of attraction associated with the chemo-sensitive population and the chemo-resistant population for increasing values of chemo-concentrations have an intertwined spiral structure with extreme sensitivity to changes in chemo-concentration level as well as sensitivity with respect to resistant mutations. For time-dependent therapies, we introduce an orbit transfer method to construct continuous families of periodic (closed) orbits by switching the chemo-dose at carefully chosen times and appropriate levels to design schedules that are superior to both maximum tolerated dose (MTD) schedules and low-dose metronomic (LDM) schedules, both of which ultimately lead to fixation of sensitive cells or resistant cells. By keeping the three subpopulations of cells in competition with each other indefinitely, we avoid fixation of the cancer cell population and re-growth of a resistant tumor. The method can be viewed as a way to dynamically shape the average population fitness landscape of a tumor to steer the chemotherapeutic response curve. We show that the method is remarkably insensitive to initial conditions and small changes in chemo-dosages, an important criterion for turning the method into an actionable strategy.

3.1 Introduction

It is widely appreciated that the development of chemotherapeutic resistance is the primary reason for recurrance of cancer in patients undergoing treatment, and remains one of the primary challenges in the field of oncology [58, 1, 17]. As a tumor grows, and even as tumor cells spread throughout the system and metastasis ensues, standard prescheduled chemotherapeutic protocols such as maximum tolerated dose (MTD) and low-dose metronomic schedules (LDM) often show early success as the tumor regresses temporarily. Schematics for several kinds of chemotherapeutic schedules are shown in Figure 3.1. A typical chemotherapeutic cycle might involve one strong dose every three weeks, or a dose for 5 consecutive days, followed by a 28 day rest period [88]. After months of fixed periodic cycles, the cancer often recurs and the tumor begins to regrow. Because of the genetic and cellular heterogeneity of a typical tumor [116], instead of killing all of the cancer cells and thereby eliminating the tumor, the chemotherapeutic regimen actually selects for a resistant phenotype [105, 39, 44, 27]. The diversity of cells within a tumor effectively protects the tumor from single-line or pre-scheduled chemotherapeutic assaults by allowing for elimination of the chemo-sensitive population in order to accomplish the subsequent release of the chemoresistant population. By reducing the relative fitness of the sensitive cells, chemotherapy acts as the primary mechanism of natural selection that selects specifically against rapidly dividing cells [68].

The characterization of a typical tumor as an adaptive landscape made up of competing cells of varying degrees of fitness, which determine growth rates of the various subpopulations is a more accurate characterization of a tumor and suggests an ecological or evolutionary approach [74, 6, 10, 9, 31, 107, 12, 118, 122, 119]. If one had access to time-resolved information [33] on the relative balance and growth rates of the subpopulations of cells making up the tumor, then one could use chemotherapy as a control device (accuator) to keep the

subpopulations in balance, competing with each other indefinitely, without any one of the cancerous subpopulations dominating the landscape [105]. Chemotherapy would then be regarded more as a maintenance mechanism than a cure [38, 54] and one would be imposing no more selection than is necessary as has been advocated by Read *et al.* and others [94].

We introduce an evolutionary game theory model of chemotherapeutic resistance [58] along with a method of adaptive control to design adventageous chemotherapeutic schedules that are able to overcome, or at least manage resistance. The mechanism of resistance that we model is based on the ecological notion of competitive release [44, 104, 74] of the resistant cell population when the sensitive cell population is reduced below a certain threshold. Above the threshold, the sensitive cells are able to outcompete (on average) the resistant cells due to the inherent *cost-of-resistance* [5] which tilts the fitness landscape of the system in favor of the sensitive cell population allowing the tumor to grow. Under sufficient chemotherapeutic pressure, the sensitive cell population is reduced enough to allow the resistant population to begin to flourish and eventually regrow the tumor in a form that is much harder to treat. A quantitative understanding of this phenomena is necessary in order to develop chemotherapeutic strategies (i.e. adaptive therapies) to combat it, a point of view adopted and developed in [38, 40, 107, 17, 10, 9]. In this paper, we frame the problem as one in nonlinear dynamical systems theory and use trajectory transfer methods developed relatively recently in the orbital mechanics literature [60, 43, 47, 99, 93] to design chemotherapeutic schedules that have the potential to outperform more standard approaches [84].

The mathematical model we use is based on a three-component replicator dynamical system with a frequency dependent fitness function based on a prisoner's dilemma (PD) payoff matrix [85, 52]. Cell interactions occur between three cell types that form our ecosystem: healthy cells (H); chemo-sensitive cells (S); and chemo-resistant cells (R). The healthy cell subpopulation can be thought of as non-neoplastic cells that have a lower fitness than the chemo-sensitive subpopulation. The independent variables represent relative frequencies in the total population, and the tumor alone would be made up of S + R cancer cell subpopulations. In a PD scenario, the healthy cells act as cooperators while the cancer cells act



Figure 3.1: Schematic dose schedules.(a) Maximum tolerated dose (MTD) schedule; (b) Low dose metronomic (LDM) schedule; (c) Adaptive schedule; (d) General time-dependent schedule accuated by piecewise constant dose concentrations.

as defectors [7]. Unchecked, the defectors saturate the population as the average fitness decreases to a sub-optimal outcome. The goal of chemotherapeutics in this framework is to coax the defectors to cooperate, leading to a higher fitness Nash equilibrium [85, 52]. Growth rates are determined by cell fitness functionals which, in turn, depend on subpopulation sizes, i.e. they are frequency dependent. There is a time-dependent controller in our model which determines chemotherapeutic dose schedule and concentrations. Although the model system is nonlinear and the control parameter enters as a coefficient in most all of the terms of the three-component cubic nonlinear system (making classic control schemes [18]

non-applicable), we use nonlinear trajectory design techniques introduced and developed in an orbital mechanics context [60, 43, 47, 99, 93]. In that body of literature, time-dependent controllers are used to design orbit transfers in a Hamiltonian mechanics setting, piecing together partial orbits at different energy levels and switching energies at carefully chosen times, much like classic Homann transfers for satellite control [92]. While the replicator system we describe is not a Hamiltonian system in which orbits transfer from one *energy* level to another, the time-dependent chemotherapeutic parameter, C(t), can be used to design adventageous orbits in the replicator dynamics tri-linear phase space in a similar manner where piecewise constant dose concentrations are used, $C_i = const$. for carefully chosen time intervals $t_i < t < t_{i+1}$, (i = 0, ..., n) with switching times t_i chosen in such a way as to produce a periodic (closed), continuous, piecewise differentiable orbit that stays trapped in a desirable region of the phase space. Orbits designed this way are shown to maintain a higher average level of fitness for the full population and avoid tumor recurrance. The existence of such orbits for appropriately chosen chemotherapeutic schedules in our model system suggests the possibility that similar orbits may also exist in a more complex tumor environment with carefully designed adaptive schedules [134, 40]. The general technique of designing an orbit with a time-dependent controller should also work in other contexts such as microbial drug resistance [81] or pest management [44, 19], although with perhaps different numbers of competing subpopulations and different payoff matrices that determine the fitness landscapes.

3.2 A three-component replicator system

The model we develop is based on a three-component replicator nonlinear dynamical system governing three competing subpopulations of cells: $\vec{x} = (x_1, x_2, x_3)^T = (x_H, x_S, x_R)^T$, where x_1 represents the proportion of healthy cells (H), x_2 represents the proportion of sensitive cells (S), and x_3 represents the proportion of resistant cells (R), with $x_1+x_2+x_3 = 1$. In this context, the model makes the *well-mixed* assumption on the cell population [85, 52], i.e. no spatial dependence is modeled in order to highlight the method in as clean a setting as possible. For discussions and comparisons of well-mixed models vs. those with spatial dependence, see [132]. The equations which describe the subpopulation interactions are:

$$\dot{x}_1 = (f_1 - \langle f \rangle) x_1, \tag{3.1}$$

$$\dot{x}_2 = (f_2 - \langle f \rangle) x_2, \tag{3.2}$$

$$\dot{x}_3 = (f_3 - \langle f \rangle) x_3. \tag{3.3}$$

with f_i representing the fitness of each of the subpopulations (i = 1, 2, 3) as their relative populations change, and $\langle f \rangle$ representing the average fitness of the entire population. The exponential growth-decay rates of each of the subpopulations are then determined by $(f_i - \langle f \rangle)$ which dictates whether the subpopulation fitness is above or below the average population fitness, hence whether the subpopulation decays or grows.

The fitness of each of the three subpopulations are defined by the linear functionals:

$$f_1 = 1 - w_1 + w_1 \left(A \vec{x} \right)_1, \qquad (3.4)$$

$$f_2 = 1 - w_2 + w_2 (A\vec{x})_2, \qquad (3.5)$$

$$f_3 = 1 - w_3 + w_3 (A\vec{x})_3, \qquad (3.6)$$

where $0 \le w_i(t) \le 1$, (i = 1, 2, 3) are time-dependent selection parameters (that serve as our controllers) we use to shape the fitness landscape of the system. A is the payoff matrix associated with the cell-cell interactions, which we describe in §3.3. The time dependence in our model enters through a chemo-concentration parameter C(t):

$$w_1 = w_0,$$
 (3.7)

$$w_2 = w_0(1 - C(t)), (3.8)$$

$$w_3 = w_0,$$
 (3.9)

where w_0 scales time (we typically take $w_0 = 1$). Note that the chemotherapy parameter acts linearly on the sensitive cell population lowering its fitness, although the three populations are coupled nonlinearly through [3.1]-[3.3]. The average population fitness is defined by the quadratic functional:

$$\langle f \rangle = x_1 f_1 + x_2 f_2 + x_3 f_3 \tag{3.10}$$

in the usual way. It is straightforward to see (for fixed values of the chemo-concentration parameter C) that the fixed points of the system [3.1]-[3.3] are of three basic types: (i) There are the three fixed points at each of the corners of the triangular phase space diagram shown in Figure 3.2, when two of the subpopulation values are zero, and the third is saturated: $(x_1, x_2, x_3) = (x_H, x_S, x_R) = (1, 0, 0); (0, 1, 0); (0, 0, 1);$ (ii) There are three possible fixed points on the triangle sides, which correspond to one of the subpopulation values equaling zero, with the other two having fitness values equal to the population average: $x_1 = 0, f_2 = \langle f \rangle, f_3 = \langle f \rangle; x_2 = 0, f_1 = \langle f \rangle, f_3 = \langle f \rangle; x_3 = 0, f_1 = \langle f \rangle, f_2 = \langle f \rangle;$ (iii) There is the balanced fitness state, when none of the subpopulation values is zero, but each of the subpopulation fitness values equals the population average: $f_1 = \langle f \rangle; f_2 = \langle f \rangle; f_3 = \langle f \rangle$. Which of these fixed points lies on or inside the triangular phase space, and their stability properties, depend in detail on the parametric values, which we describe in §3.4.

By using a prisoner's dilemma payoff matrix (see section 3.3) we ensure: (i) Gompertzian growth of the cancer cells [83]; (ii) a reduction in overall fitness of the population as the tumor grows; and (iii) a fitness cost associated with resistance. We first study the details of the nonlinear dynamics associated with eqns [3.1]-[3.3] for *constant* values of the chemotherapy parameter $0 \le C \le 1$ to demonstrate the mechanism of competitive release for threshold values $C \ge 1/3$. Then we investigate piecewise constant time-dependent functions C(t) to show how to avoid the evolution of resistance of the tumor. Figure 3.1 shows several examples of the chemotherapeutic schedules we consider. These include maximum tolerated dose (MTD) schedules (Figure 3.1(a)), low-dose metronomic schedules (LDM) (Figure 3.1(b)), adaptive



Figure 3.2: The competitive release mechanism. (a) The three-component phase space associated with competing populations of (H, S, R) cells. (b) With no therapy, the S corner of the triangle is a globally attracting fixed point, while the H and R corners are unstable. All initial conditions lead to a saturated tumor. Filled circles are stable, unfilled circles are unstable. (c) For continuous chemotherapy above a threshold level, the R corner of the triangle is a globally attracting fixed point, while the H and S corners are unstable. All initial conditions (except those lying on the separatrix connecting the interior balanced fitness state to the S - R side) lead to a resistant tumor.

schedules (Figure 3.1(c)), and more general time-dependent schedules (Figure 3.1(d)) which we break-up into piecewise constant doses as done in forming the Riemann sum approximation to an area under a curve. In all cases, we compare outcomes of the different schedules holding the total dosage, D, fixed over time period τ ,

$$D = \int_0^\tau C(t)dt = const.$$
(3.11)

3.3 The prisoner's dilemma as a cancer model

We first describe the standard version of the PD payoff matrix [85] in a 2×2 setting to make clear why the PD evolutionary game gives rise to a Gompertzian (i.e. 'S-shaped') growth curve when two sub-populations compete. To be specific, we first describe what happens when healthy cells compete with cancer cells using the 2×2 payoff matrix:

$$A_{PD} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 3 & 0 \\ 5 & 1 \end{bmatrix},$$
(3.12)

with c > a > d > b. The first row and column correspond to the payoffs associated with the *cooperator* (C) in the PD evolutionary game, and the second row and column correspond to the payoffs associated with the *defector* (D). In the simplest tumor growth paradigm in which a population of healthy cells competes with a population of cancer cells, the healthy cells are the cooperators, while the cancer cells are the defectors. In any mixed population $\vec{x} = (x_C, x_D)^T$, $0 \le x_C \le 1$; $0 \le x_D \le 1$; $x_C + x_D = 1$, the fitness functions, $\vec{f} = (f_C, f_D)^T$, associated with the two subpopulations are:

$$\vec{f} = A_{PD}\vec{x}, \tag{3.13}$$

which in component form yields:

$$f_C = (A_{PD}\vec{x})_1 = 3 \cdot x_C + 0 \cdot x_D, \qquad (3.14)$$

$$f_D = (A_{PD}\vec{x})_2 = 5 \cdot x_C + 1 \cdot x_D, \qquad (3.15)$$

while the average fitness of the total population is given by the quadratic form:

$$\langle f \rangle = \vec{x}^T A_{PD} \vec{x} = 3x_C^2 + 5x_C x_D + x_D^2 \ge 1.$$
 (3.16)

Note that the average fitness of the healthy state $(x_C, x_D) = (1, 0)$ is given by $\langle f \rangle|_{(x_C=1)} = 3$, while that of the cancerous state $(x_C, x_D) = (0, 1)$ is given by $\langle f \rangle|_{(x_D=1)} = 1$, which minimizes the average fitness. Tumor growth is then modeled as a 2×2 evolutionary game governed by the replicator dynamical system:

$$\dot{x}_C = (f_C - \langle f \rangle) x_C, \tag{3.17}$$

$$\dot{x}_D = (f_D - \langle f \rangle) x_D. \tag{3.18}$$

It is straightforward to show:

$$\dot{x}_D = [(c-a) - (d-b)]x_D(1-x_D) \left(\frac{1}{1 - \left(\frac{d-b}{c-a}\right)} - x_D\right),$$

with fixed points at $x_D = 0, 1, \frac{(c-a)}{(c-a)-(d-b)}$. From this, we can conclude that for any initial condition containing at least one cancer cell: $0 < x_D(0) \le 1$, we have:

(i) $x_D \to 1, x_C \to 0 \text{ as } t \to \infty$

(ii)
$$\langle f \rangle \to 1 \text{ as } t \to \infty$$
.

Condition (i) guarantees that the cancer cell population will saturate, while condition (ii) guarantees that the saturated state is sub-optimal, since $\langle f \rangle|_{(x_D=1)} < \langle f \rangle|_{(x_C=1)}$. For these two reasons, the prisoner's dilemma evolutionary game serves as a simple model for tumor growth both in finite population models, as well as replicator system (infinite population) models [118, 119, 122, 114].

For our purposes, we now generalize to a three-component system where the fitness functions [3.4]-[3.6] are defined via a payoff matrix A which is a 3×3 matrix defining the evolutionary game played by the cell population. For this, we take every 2×2 sub-matrix to be a PD game, i.e. we take A to be the 3×3 PD matrix:

$$A = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} = \begin{bmatrix} 3 & 1.5 & 1.5 \\ 4 & 2 & 2.8 \\ 3.9 & -2 & 2.2 \end{bmatrix},$$
(3.19)

with the PD inequalities [85]:

$$g > a > i > c, \tag{3.20}$$

$$d > a > e > b, \tag{3.21}$$

$$f > i > e > h. \tag{3.22}$$

The numerical values in [3.19] are chosen for convenience and satisfy the constraints [3.20]-[3.22]. In each cell-cell interaction, the healthy cells x_1 (healthy H) are cooperators, and the two-species of cancer cells, x_2 (sensitive S), and x_3 (resistant R), are the defectors. In any interaction between a chemo-sensitive cell (S) and a chemo-resistant cell (R), the sensitive cell is the defector, while the resistant cell is the cooperator. The payoff matrix [3.19] guarantees that for any interaction between two cells, the system retains features (i),(ii) detailed previously (i.e. tumor growth leading to sub-optimal fitness). In addition, the payoff matrix also imposes a *cost to resistance* if we add the extra constraint d > g which guarantees:

$$f_{S} = dx_{H} + ex_{S} + fx_{R}$$

$$(3.23)$$

$$> f_{R} = gx_{H} + hx_{S} + ix_{R}$$

It is worth pointing out that shaping the fitness landscape by adjusting the selection parameters (w_1, w_2, w_3) is equivalent to choosing

$$f_i = (\tilde{A}\vec{x})_i, \tag{3.24}$$

where:

$$\tilde{A} = \begin{bmatrix} 1 - w_1 + aw_1 & 1 - w_1 + bw_1 & 1 - w_1 + cw_1 \\ 1 - w_2 + dw_2 & 1 - w_2 + ew_2 & 1 - w_2 + fw_2 \\ 1 - w_3 + gw_3 & 1 - w_3 + hw_3 & 1 - w_3 + iw_3 \end{bmatrix}.$$

To see this, use the fact that $(1 - w_i) \equiv (1 - w_i)(x_1 + x_2 + x_3)$ in the fitness equations [3.4]-[3.6].



Figure 3.3: Fixed point locations. Tracking the location of the fixed points as a function of chemo-concentration parameter C. The fractional values indicate that the values are analytically obtained.

3.4 Results

3.4.1 Continuous therapy

For the purposes of understanding how to implement time-dependent therapies in our model, we first describe the phase space dynamics for constant values $0 \leq C \leq 1$, so



Figure 3.4: Basins of attraction and nullclines. Panel showing the separatrices and nullclines $(dx_H/dt = 0, dx_S/dt = 0, dx_R/dt = 0)$ through the balanced fitness interior fixed point that determines the basin boundaries of attraction of the S state and the R states. The interior fixed point is the one in which each of the subpopulation fitness levels exactly matches the average fitness of the entire population. Filled circles are stable, unfilled circles are unstable.

total dose delivered over time period t is $D = C \cdot t$. Figure 3.2 shows the mechanism of chemotherapeutic resistance via competitive release determined by the system [3.1]-[3.3] (for constant C) as depicted in the triangular phase space diagram shown in Fig 3.2(a). With no therapy (Fig.3.2(b), C = 0), the sensitive corner S is globally attracting, while the H and R corners are unstable. When the continuous therapy parameter C = 0.6 (Figure 3.2(c)), the resistant corner is globally attracting, while the H and S corners are unstable. Filled corners are stable, unfilled corners are unstable. This is the basic mechanism of competitive release induced by sufficiently strong chemotherapeutic dose [105].



Figure 3.5: **Basin of attraction areas.** Areas of basin of attraction of S fixed point and R fixed point as a function of chemo-concentration C.

Intermediate values of C reveal a much more complex picture. We show in Fig 3.3 the location of the fixed points as a function of C. For values 79/228 = 0.34649... < C < 0.7, the balanced fixed point state is an interior fixed point, which forms the central spiral associated with the basin boundaries between the asymptotically stable S corner and R corner, shown in the panel in Figure 3.4. Also shown in the figure are the nullclines defined by the curves $dx_H/dt = 0$, $dx_S/dt = 0$, $dx_R/dt = 0$. Opposite sides of each of the nullclines mark a change in whether the particular subpopulation decays or grows. The mixed population state exists between values 1/3 < C < 1/2 where the basin of attraction sizes (areas), shown in Figure 3.5, change sensitively (as C varies), one at the expense of the other, in an intertwined spiral



Figure 3.6: Fitness landscapes as a function of time. With no therapy (C = 0), the fitness curves are contunuously decreasing functions as the tumor saturates with the sensitive cell population in a sigmoidal shaped growth curve. With continuous therapy (C = 0.6), the fitnesses initially increase indicating tumor regression, but eventually decrease. The healthy subpopulation initially increases before the resistant population eventually saturates the tumor. Initial condition (0.9, 0.09, 0.01) with $w_0 = 0.1$

structure centered at the balanced fixed point state. The steep transition curves between the two states occuring for small changes in the chemo-concentration parameter highlights the sensitivity of the system to chemotherapeutic dosing levels. Also worth highlighting in Fig. 3.4(b),(c) is the sensitivity to resistant mutations of the final asymptotic state of the system,

even if no pre-existing mutations exist in the population. As the tumor grows, assuming no pre-existing resistant mutations, the dynamics would traverse down the left side of the triangle from the H corner to the S corner. As it does, the very thin passageway (on the H-S side of the triangle in Figure 3.4(b),(c)) associated with the basin of attraction of the resistant corner indicates that one single resistant mutation would push the trajectory off the side into the basin of attraction associated with the R corner instead of converging to the S corner.

The fitness landscapes are shown in Figure 3.6 both for C = 0 and C = 0.6. With no therapy, the fitness curves are monotonically decreasing functions as the sensitive population saturates the tumor in a sigmoidal shaped [83, 118] growth curve. With continuous therapy above threshold, the fitness curves initially increase (i.e. tumor regression), but eventually decrease monotonically as the tumor relapses. The healthy cell population initially increases, but eventually the resistant population saturates the tumor.

3.4.2 Time-dependent therapy

Figure 3.7 shows the key idea behind the method we use to design trajectories for eqns [3.1]-[3.3] for time-dependent chemotherapeutic schedules C(t). With no chemotherapy, C = 0, since the sensitive corner S is a globally asymptotically attracting fixed point whose basin of attraction is the full region, all trajectories that start inside the triangle eventually get trapped in the left S corner. By contrast, with chemotherapeutic levels at C = 0.7, competitive release acts to create a basin of attraction for the right resistant corner R for all initial conditions inside the triangle. Using these families of solution trajectories, we overlay the solution curves in Figure 3.7(c) to show the underlying curvilinear grid that spans the full tri-linear phase plane. By switching between the two values C = 0 and C = 0.7 at times when two curves intersect, it is possible to transition from a trajectory associated with the C = 0 family to one associated with the C = 0.7 family. This creates multiple possibilities for designing complex orbits using piecewise constant values of C, with switching at appropriately chosen times.

One such trajectory is shown in Figure 3.8(a) achieved by switching between the two values C = 0 and C = 0.5. The trajectory starts at point A((H, S, R) = (0.5, 0.4, 0.1))using a trajectory with C = 0.5. When the trajectory reaches point B, the chemotherapy is switched off C = 0, and the system then travels back down to point A in a closed loop. With switching times labeled T_1 and T_2 (which are computed by monitoring the relative balances of the three subpopulations), this closed orbit can be maintained indefinitely. By contrast, Figure 3.8(b) shows the corresponding MTD and LDM trajectories with the same initial condition (point A) using the same total dosages over the same time-periods. The MTD trajectory eventually gets trapped in the left S corner, as does the LDM trajectory. Figure 3.8(c) shows the schedules for all three cases. In all cases, we first design the adaptive schedule, then we create the MTD and LDM schedules using the same total dosage D. Neither the MTD nor the LDM standard chemo-schedules is able to prevent the system from saturating with a full grown tumor, whereas the adaptive schedule keeps the system trapped indefinitely near the top H corner of the system. Figure 3.8(d) shows the average fitness of the system for the MTD, LDM, and adaptive schedules. It is clear that the adaptive schedule is able to maintain a higher average fitness throughout the full course of chemotherapy.

Figure 3.9 shows the sensitivity of the system to the chemo-concentration levels chosen. Here we switch between C = 0 and C = 0.6 (higher average dose than in figure 3.8) to construct the closed loop (figure 3.9(a)), which looks very similar to that in Figure 3.8. But notice that for these values, the LDM schedule creates an orbit that saturates at the R corner, whereas the MTD schedule saturates at the S corner. The actual schedules are shown in Figure 3.9(c). Figure 3.9(d) compares the average population fitness of the three schedules, showing the adaptive schedule maintaines a higher average throughout. Notice that LDM initially achieves a higher average fitness before tumor regression occurs.

In Figure 3.10 we show the result of toggeling between values C = 0.3 and C = 0.6 to maintain the periodic loop (Figure 3.10(a)). For this case, both the MTD and LDM schedules send the trajectory to the *R* corner (Figure 3.10(b)). Figure 3.10(d) shows the

initial benefit of the MTD and LDM schedules in terms of higher initial average fitness, but eventually the adaptive schedule shows its superiority over both.

Figure 3.11(a) shows that we can actually design an orbit that starts at an arbitrary point A inside the triangle, and send it to an arbitrary point B. We accomplish this by constructing the incoming and outgoing orbits from point A for two different C values, and those associated with point B for those same two C values, showing that they must intersect at some point which we label O. By sending the orbit out from point A to point O, then switching values of C at point O until we arrive at point B, we complete the transfer. One can immediately see the potential richness in the possible design of different orbits one can construct by switching values of C among two, three, or more values, at appropriately chosen times. Figure 3.11(b) shows the richness of the curvilinear grid that can be created with three values of C = 0, 0.3, 0.7 and the multitude of possible paths from one point to another inside the triangle if one allows for switching among the three values shown.



Figure 3.7: Dynamical trajectories for constant C.(a) With no chemotherapy (C = 0), the sensitive corner S is a globally attracting fixed point. All initial conditions inside the triangle move to S along the sample trajectories shown. (b) Above the chemotherapy threshold C > 0.5, all initial conditions inside the triangle move to the resistant corner R. Shown are sample trajectories for C = 0.7. (c) Overlay of the solution trajectories for C = 0 and C = 0.7 create a curvilinear grid throughout the triangle. By switching between these two values of C, we construct a global trajectory made up of segments of the two families of trajectories.



Figure 3.8: Constructing a closed loop trajectory using C = 0 and C = 0.5, with $C_{avg} = 0.366.$ (a) By using segments of a trajectory for C = 0 and C = 0.5, switching values at points A and B, we construct a closed periodic orbit. (b) Using the same total dose, we show the MTD and LDM trajectories starting from point A. Both eventually move to the S corner, although initially the MTD trajectory moves toward the H corner (tumor regression) before recurrance. (c) The MTD, LDM, and adaptive schedules are depicted. (d) We plot the average fitness $\langle f \rangle$ for the three different chemo-schedules. The adaptive schedule is able to maintain a higher average fitness throughout.

3.4.3 Robustness

To be clinically actionable, at the very least it would be important that the method not be sensitive to specific initial conditions and parameter values. We describe several remarkably robust features of the strategy in this section. We show in Figure 3.12(a) an example of a



Figure 3.9: Constructing a closed loop trajectory using C = 0 and C = 0.6, with $C_{avg} = 0.375.$ (a) By using segments of a trajectory for C = 0 and C = 0.6, switching values at points A and B, we construct a closed periodic orbit. (b) Using the same total dose, we show the MTD and LDM trajectories starting from point A. The MTD trajectory eventually moves to the S corner, while the LDM trajectory moves to the R corner. Both trajectories initially move toward the H corner (tumor regression) before recurrance. (c) The MTD, LDM, and adaptive schedules are depicted. (d) We plot the average fitness $\langle f \rangle$ for the three different chemo-schedules. The adaptive schedule is able to maintain a higher average fitness throughout, although LDM initially achieves higher average fitness before declining.

continuous family of closed orbits that are easily achievable by using a bang-bang strategy, with blue orbits (no therapy) and red orbits (C = 0.7) forming a continuous family of closed, and nearly closed loops. The associated dose schedules are shown in Figure 3.12(b)-(j), with average doses all very close to 0.4. Because of the continual interlacing of these orbits, and the robustness of their associated schedules, it is clear there is nothing particularly special about any of the values shown, or the average dose - any of the orbits and schedules would work. Interestingly, Figure 3.13 shows that the entire region on the right side of the triangle centered on the dashed curve can be populated by continous families of closed orbits each centered at any point on the line. Figure 3.13(a) depicts the defining feature of the line which requires that two orbits with different C values have a point of tangency lying on it. A necessary and sufficient condition for this is $f_1 = f_3$, which defines a line, a condition obtained by requiring that tangent vectors associated with two different values of C point in opposite directions, as shown in Figure 3.13(a). Figure 3.13(b) shows examples of closed orbits obtained at 5 representative points on the dashed line.



Figure 3.10: Constructing a closed loop trajectory using C = 0.3 and C = 0.6, with $C_{avg} = 0.396$.(a) By using segments of a trajectory for C = 0.3 and C = 0.6, switching values at points A and B, we construct a closed periodic orbit. (b) Using the same total dose, we show the MTD and LDM trajectories starting from point A. Both trajectories move toward the R corner after initially moving toward H. (c) The MTD, LDM, and adaptive schedules are depicted. (d) We plot the average fitness $\langle f \rangle$ for the three different chemo-schedules. The adaptive schedule is able to maintain a higher average fitness throughout. Both MTD and LDM initially show higher average fitness before declining.

3.5 Discussion

The model shows that if the chemo-dose exceeds a threshold value of C > 0.5, the tumor may regress for a period of time, but eventually regrows to form a resistant tumor (Figure 3.2(c)) as long as there is at least one resistant cell in the population (i.e. a pre-existing



Figure 3.11: Constructing orbit transfers. (a) We take two arbitrary points labeled A and B and construct the incoming and outgoing solution trajectories from each, using two values C = 0 and C = 0.7. There must be a crossing point, which we label point O. The two segments AO followed by OB is the two-switch trajectory that takes us from A to B. (b) Shown is the curvilinear grid constructed by families of solution curves for the three values C = 0, C = 0.3, C = 0.7. Any point of intersection can be used as a starting point and an ending point to construct a three-switch orbit.

resistant mutation). This process of competitive release is very robust and occurs for all initial distributions of the three subpopulations. For chemo-doses that are not as large, the results are much more sensitive to small changes in C since the phase space diagram generally has two basins of attraction, one associated with a saturated sensitive state, the other with a saturated resistant state, and these basins of attraction are spirally intertwined. Small differences in the balances among the three types of cells comprising our model, and a single resistant mutation can determine whether the long-time dynamics converges to the S state or the R state. The relative areas of the basins of attraction also sharply change with small changes in chemo-dose values between 1/3 < C < 1/2 reflecting the sensitivity in choosing the optimal chemo-dose levels. There is no constant value of C for which the system converges to the H state, as this state is always unstable.

By contrast, for a time-dependent chemotherapy parameter C(t), one can design a rich, robust, and potentially endless array of trajectories (i.e. tumor responses) that remain in desirable regions of the triangular phase space indefinitely, depending on how many timeswitches and different dose levels one is willing to accept. One can, in principle, steer the system trajectory along any path (in the model). Identifying a given trajectory in our model as well as the chemo-therapeutic schedule that produces it is fairly straightforward and amounts to superposing the response curves associated with different *constant* therapies, and piecing together a global trajectory made up of sections generated from each of the constant therapy curves. Of course, in a clinical setting this procedure will be much more complex but suggests the possibility of using actual chemo-therapeutic responses with different chemo-schedules used on different patients as a design template to combine in new ways to predict possible responses to various time-dependent switching strategies. By using one dosing schedule to create a time-limited response, then a different dosing schedule to create a different (time-limited) response, one could think of piecing together finite-time limited responses to create new outcomes, superior to what would have been achieved by sticking to one single chemo-schedule.

One of the challenges of testing and designing new schedules via clinical trials is that short term gain in average fitness with LDM or MTD does not always result in more long-term increased averaged fitness levels, i.e. recurrance sets in if the schedule does not completely eradicate the tumor. This is clearly shown in our simulations. The possibility of designing complex orbits with various potentially adventageous features is virtually endless if one allows for enough switches among many different levels of the chemo-concentration parameter C. In this manuscript, we have focused on comparisons of average fitness levels to evaluate the quality of the schedule. One could, of course, imagine designing orbits that maximize a (time-averaged) population fitness while minimizing total dose, or perhaps delay recurrance as long as possible while avoiding regions of the phase space (i.e. relative balances of the subpopulation levels) where the total tumor burden becomes unsustainable. The fact that closed loop trajectories can be designed in our model three-component replicator system by chemo-scheduling alone suggests the possibility that similar orbits could potentially be created in an actual tumor environment, microbial environment, or pest management setting
with the right accuation (chemo-schedule/antibiotic regimen/pest control schedule), where fixation of the sensitive population or fixation of the resistant population are both avoided indefinitely as the balance is managed with proper adaptive therapy. This, of course, all hinges on our ability to carefully and frequently monitor the tumor environment [33]. Clinical trials that make use of adaptive scheduling ideas are currently being carried out at the Moffitt Cancer Center, and results show great promise [134]. More possibilities exist with the use of combination multi-drug therapies [10, 81, 130], or immunotherapies [95], but these are outside the scope of our model which focuses only on the use of dosing schedules to shape the fitness landscape of a tumor to steer its evolutionary response.



Figure 3.12: Robust family of nested orbits. (a) One example of a (continuous) family of nested closed orbits using bang-bang strategy toggling between C = 0 and C = 0.7, all with similar total dose ($C \sim 0.4$). A closed loop can be made arbitrarily small around a center point described in Figure 3.13. (b) Schedule associated with orbits 1 and 4; (c) Schedule associated with orbits 1 and 5; (d) Schedule associated with orbits 1 and 6; (e) Schedule associated with orbits 2 and 4; (f)Schedule associated with orbits 2 and 5; (g) Schedule associated with orbits 3 and 4; (i)Schedule associated with orbits 3 and 4; (i)Schedule associated with orbits 3 and 4; (j)Schedule associated with orbits 3 and 6.



Figure 3.13: All closed orbit families are centered on dashed line defined by $f_1 = f_3$. (a) As a family of closed orbits enclose smaller and smaller loops, they intersect at a point of tangency. The point of tangency must lie on the dashed line. (b) Examples of several closed loop orbits centered at different points down the line.

Chapter 4

The role of synergy and antagonism in designing multi-drug adaptive chemotherapy schedules

Chemotherapeutic resistance via the mechanism of competitive release of resistant tumor cell subpopulations is a major problem associated with cancer treatments and one of the main causes of tumor recurrence. Often, chemoresistance is mitigated by using multidrug schedules (two or more combination therapies) that can act synergistically, additively, or antagonistically on the heterogeneous population of cells as they evolve. In this paper, we develop a three-component evolutionary game theory model to design two-drug adaptive schedules (timing and dose levels associated with $C_1(t)$ and $C_2(t)$) that mitigate chemoresistance and delay tumor recurrence in an evolving collection of tumor cells with two resistant subpopulations: R_1 (sensitive to drug 1, resistant to drug 2), and R_2 (sensitive to drug 2, resistant to drug 1). A key parameter, e, takes us from synergistic (e > 0), to additive (e = 0), to antagonistic (e < 0) drug interactions. In addition to the two resistant populations, the model includes a population of chemosensitive cells, S that have higher baseline fitness but are not resistant to either drug. Using the nonlinear replicator dynamical system with a payoff matrix of Prisoner's Dilemma (PD) type (enforcing a cost to resistance), we investigate the nonlinear dynamics of the three-component system (S, R_1, R_2) , along with an additional tumor growth model whose growth rate is a function of the fitness landscape of the tumor cell populations. We show that *antagonistic* drug interactions generally result in slower rates of adaptation of the resistant cells than *synergistic* ones, making them more effective in combating the evolution of resistance. We then design closed loops in the three-component phase space by shaping the fitness landscape of the cell populations (i.e. altering the evolutionary stable states of the game) using appropriately designed time-dependent schedules (adaptive therapy), altering the dosages and timing of the two drugs using information gleaned from constant dosing schedules. We show that the bifurcations associated with the evolutionary stable states are transcritical, and we detail a typical antagonistic bifurcation that takes place between the sensitive cell population S and the R_1 population, and a synergistic bifurcation that takes place between the sensitive cell population S and the R_2 population for fixed values of C_1 and C_2 . These bifurcations help us further understand why antagonistic interactions are more effective at controlling competitive release of the resistant population than synergistic interactions in the context of an evolving tumor.

4.1 Introduction

We study a mathematical model to explore the role of synergisitic vs. antagonistic multidrug interactions on an evolving population of cancer cells in a tumor. Our model builds on the (single drug) adaptive therapy model developed in [78] and is based on a replicator dynamical system of three (well-mixed) populations of cells: (i) sensitive cancer cells, S, that are sensitive to both drug 1 and drug 2, (ii) a resistant population, R_1 , that is sensitive to drug 1 but resistant to drug 2, and (iii) resistant population, R_2 , that is sensitive to drug 2 but resistant to drug 1. The replicator dynamical system governing the relative frequencies of (S, R_1, R_2) makes use of a 3×3 payoff matrix A of Prisoner's Dilemma type [52, 85, 7, 119]. We introduce chemotherapeutic schedules using time-dependent dosing functions $C_1(t)$, $C_2(t)$ that alter the fitness levels of the sensitive cell population and the two resistant populations independently allowing us to shape the fitness landscape of the population of cells adaptively as the tumor evolves [40, 39, 107, 32, 37, 134, 34, 35]. Our general goal in this context is to design schedules that delay tumor recurrence (re-growth) due to competitive release of the resistant cell population [42, 17, 63, 58] and to quantify the role of synergistic and antagonistic drug interactions in this process. To quantify the synergistic vs. antagonistic effects of the two toxins on the evolving populations, we use a parameter e which we alter. When e > 0, the drugs interact synergistically, when e = 0, the drugs are additive, and when e < 0, they interact antagonistically. Level curves of the fitness profiles as a function of C_1 and C_2 are shown in figure 4.1 which can be compared with figure 1 of [50]. Details of how we define our fitness functions in the context of our replicator model will be described in §II. In the neutral case, additive drug interactions (figure 4.1(b)) accomplish exactly the kill rate that the sum of each of the two would accomplish acting independently. By contrast, synergistic drug interactions (figure 4.1(a)) bend the curves inward, indicating that the growth rates (fitness) are lowered more than the two dosages would accomplish independently, while antagonistic interactions (figure 4.1(c)) bend the curves outward, indicating that the growth rates are lowered less than the two dosages would accomplish independently.

There is a large and dedicated literature on characterizing the interactions of many different combinations of toxins on static cell populations. As far as we are aware, the first comprehensive study of synergistic vs antagonistic effects was carried out by Bliss [14] in 1939, using joint probabilities, leading to a formula that is commonly called the Bliss index for drug interactions. A similar but slightly modified criterion was introduced by Loewe [71] and more recently developed further by Chou and collaborators [25]. These indices have all been used to help quantify the many different types of interactions that can occur with two or more toxins applied jointly in a static population of cells. In this context, it is common to assume that synergistic interactions are desirable in most circumstances, as a lower total dosage accomplishes the same kill rate as a higher dose would accomplish if the drugs acted independently. These kinds of studies have been used effectively to choose appropriate drug cocktails to individual patients by testing wide ranges of combinations on tissue samples obtained from patient tumors [110].

When the interacting population of cells are evolving, however, the relevant criteria become more complex. This is due to the fact that the subpopulations of cells respond differently to the different toxins applied, and as they respond, an ever-changing (adaptive)



Figure 4.1: (a) e > 0 showing synergistic profiles; (b) e = 0 showing additive profiles; (c) e < 0 showing antagonistic profiles.

combination of toxins might be required to accomplish a given goal. Instead of necessarily killing the maximum number of cancer cells with the least amount of toxin, it is often the case that the goal becomes avoiding chemo-resistance and delaying unwelcome tumor recurrence to the maximum extent possible. A strategy called *resistance management* [94] is often advocated and occasionally implemented [10]. These kinds of strategies have been advocated and implemented in chemotherapy settings [42, 39, 134, 37, 122, 121, 9], but perhaps have been most elegantly and thoroughly carried out in a bacterial setting (since experiments are more practical) by Kishony and collaborators [50, 76, 3, 100, 11, 57] who even discuss strategies that might reverse antibiotic resistance [11]. See also [128] for recent work discussing both microbial populations and cancer cells and [82, 131] for novel sequential therapy methods. In the context of evolving microbial populations [70], mutations occurr frequently and it is important to consider not only pre-existing mutated subpopulations, but also mutations that occur as a result of the application of antibiotic agents. In the case of chemotherapeutic resistance in tumors, it is often assumed that resistant mutations occured before the application of treatment, hence it is common to separate the subpopulations into

sensitive and resistant subpopulations as we do in our deterministic model which does not include further mutations during treatment. See [16] for further discussions of these and related issues, and [53, 15] for discussions of the general approach of using evolutionary game theory in biology.

In section §II we introduce the details of the three-component replicator dynamical system that we use. Section §III describes the effects of constant chemotherapy schedules on the evolving populations, using the range of values $0 \le C_1 \le 1$, $0 \le C_2 \le 1$ (with a total dose upper threshold $C_1+C_2 \le 1$), along with our parameter e over a range of positive (synergistic) to negative (antagonistic) values. We describe in detail the transcritical bifurcations that occur and the tumor growth in response to the chemotoxins. In §IV we introduce adaptive time-dependent schedules ($C_1(t), C_2(t)$) along with the parameter e with the goal of delaying tumor recurrence and we discuss the rate of adaptation of the cell populations in this context. Finally in §V we discuss the relevance of our model to the design of adaptive-therapy clinical trials.

4.2 A three-component replicator system

4.2.1 Why the prisoner's dilemma game?

To understand why the prisoner's dilemma (PD) game is a useful paradigm for tumor growth, consider a two-component system, with a population of healthy cells H, and a population of sensitive cancer cells S, each competing for the best payoffs. The standard version of the prisoner's dilemma payoff matrix is:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 3 & 0 \\ 5 & 1 \end{bmatrix},$$
(4.1)

where the first row corresponds to my payoffs if I choose to cooperate (associated with the healthy cell population (H)) and the second row corresponds to my payoffs if I choose to defect (associated with chemo-sensitive cells (S)). I am better off defecting than cooperating

if my opponent cooperates (first column), since my payoff is 5 instead of 3. I am also better off defecting than cooperating if my opponent defects (second column), since my payoff is 1 rather than 0. No matter what my opponent does, I am better off defecting. But my opponent uses the same logic and comes to the same conclusion - that she is also better off defecting, no matter what I do. The result is that we both defect and get a payoff of 1 (second row, second column) which is less than we would have gotten had we both cooperated (payoff of 3). The fact that defect-defect is a sub-optimal Nash equilibrium (neither player is better off by making a different choice) is the essence of the prisoner's dilemma game [52, 85] and the reason it is so widely used as the basis for modeling the evolution of cooperation in many different contexts. For cancer cell modeling, the healthy cells are the cooperators and the cancer cells are the defectors. Unlike other contexts in which game theory is used, cells are not making strategic decisions, instead, their strategy is encoded in their reproductive prowess, and selection is frequency dependent. In any mixed population $\vec{x} = (x_H, x_S)^T$, $0 \le x_H \le 1$; $0 \le x_S \le 1$; $x_H + x_S = 1$, the fitness functions, $\vec{f} = (f_H, f_S)^T$, associated with the two subpopulations are:

$$\vec{f} = A\vec{x}, \tag{4.2}$$

which in component form yields:

$$f_H = (A\vec{x})_1 = 3 \cdot x_H + 0 \cdot x_S, \tag{4.3}$$

$$f_S = (A\vec{x})_2 = 5 \cdot x_H + 1 \cdot x_S, \tag{4.4}$$

while the average fitness of the total population is given by the quadratic form:

$$\langle f \rangle = \vec{x}^T A \vec{x} = 3x_H^2 + 5x_H x_S + x_S^2 \ge 1.$$
 (4.5)

The average fitness of the healthy state $(x_H, x_S) = (1, 0)$ is given by $\langle f \rangle|_{(x_H=1)} = 3$, while that of the cancerous state $(x_H, x_S) = (0, 1)$ is given by $\langle f \rangle|_{(x_S=1)} = 1$, which minimizes the average fitness. For the static game, the cancerous state $(x_H, x_S) = (0, 1) \equiv \mathbf{p}^{*T}$ is a strict Nash equilibrium since $\mathbf{p}^{*T} A \mathbf{p}^* > \mathbf{p}^T A \mathbf{p}^*$, for all \mathbf{p} [52]. We can then embed this static game into an evolutionary context using the replicator dynamical system [85]:

$$\dot{x}_H = (f_H - \langle f \rangle) x_H, \tag{4.6}$$

$$\dot{x}_S = (f_S - \langle f \rangle) x_S, \tag{4.7}$$

from which (using values from eqn (4.1)) it is straightforward to show:

$$\dot{x}_S = x_S(1 - x_S)(2 - x_S),\tag{4.8}$$

with fixed points at $x_S = 0, 1, 2$. The cancerous state $(x_s = 1)$ then becomes an asymptotically stable fixed point of the dynamical system (4.8) and an evolutionary stable state (ESS) of the system (4.6), (4.7) which serves to drive the system to the strict Nash equilibrium under the flow. The fact that this ESS also corresponds to the one with the lowest average fitness is an extra feature of the PD game. For any initial condition containing at least one cancer cell: $0 < x_S(0) \le 1$, we have:

(i) $x_S \to 1, x_H \to 0 \text{ as } t \to \infty$

(ii)
$$\langle f \rangle \to 1 \text{ as } t \to \infty$$

The first condition (and the structure of the nonlinear equations) guarantees that the cancer cell population will saturate at the carrying capacity of 1 in an S-shaped (logistic) growth curve, while the second guarantees that this asmptotically stable carrying capacity is sub-optimal, since $\langle f \rangle|_{(x_S=1)} < \langle f \rangle|_{(x_H=1)}$. For these two reasons, the prisoner's dilemma evolutionary game serves as a simple paradigm for tumor growth both in finite population models, as well as replicator system (infinite population) models [118, 119, 122, 114]. This two-component system alone, however, is not able to account for the evolution of resistance.

4.2.2 The three-component model

The model we employ is a three-component replicator dynamical system for the three subpopulations of cells: $(S, R_1, R_2) \equiv (x_1, x_2, x_3)$

$$\dot{x}_1 = (f_1 - \langle f \rangle) x_1, \tag{4.9}$$

$$\dot{x}_2 = (f_2 - \langle f \rangle) x_2, \tag{4.10}$$

$$\dot{x}_3 = (f_3 - \langle f \rangle) x_3, \tag{4.11}$$

where each dependent variable represents the relative frequencies of cells, with $x_1 + x_2 + x_3 = 1$. In these equations, f_i represents the fitness of subpopulation i = 1, 2, 3, while $\langle f \rangle$ represents the average fitness of all three subpopulations. These equations then give rise to the obvious interpretation that if a given subpopulation's fitness is above(below) the average, it grows(decays) exponentially - reproductive provess is directly associated with the deviation of the fitness of a subpopulation from the average fitness of the entire population.

The fitness functions are frequency-dependent (i.e. non-constant), which couples eqns (4.9)-(4.11) nonlinearly. The subpopulation fitness f_i (i = 1, 2, 3) function is given by:

$$f_i = 1 - w_i + w_i (A\vec{x})_i \tag{4.12}$$

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$
(4.13)

where A is a 3×3 payoff matrix which introduces the evolutionary game being played by x_1 , x_2 , and x_3 . In this paradigm, the sensitive population (x_1) are the defectors (higher fitness) and both groups of resistant cells $(x_2 \text{ and } x_3)$ are the cooperators (lower, but not equal fitness). We use $0 \le w_i \le 1$ as a parameter to determine the strength of selection in the

system. When $w_i \sim 0$, selection is relatively weak and the evolutionary game does not play a big role in the balance of the three subpopulations. When $w_i \sim 1$, selection is strong, and the game plays a bigger role. Both of those limiting cases have been discussed in the literature [129, 117], but the tumor response through the full range of values is poorly understood. We use the selection parameters w_i as our mechanism to introduce chemotherapy, by way of altering the relative fitness of the three subpopulations:

$$w_1 = w_0(1 - C_1 - C_2 - eC_1C_2) (4.14)$$

$$w_2 = w_0(1 - C_1) \tag{4.15}$$

$$w_3 = w_0(1 - C_2). (4.16)$$

Here, C_1 is the chemotherapy parameter associated with drug 1, C_2 is that for drug 2, and *e* is our synergy (e > 0) vs. antagonism (e < 0) parameter, and we take $w_0 = 0.1$ which sets the timescale in our simulations. Level curves of the effective fitness functions defined by eqns (4.14)-(4.16) are shown in figure 4.1. The average fitness of the population is given by:

$$\langle f \rangle = f_1 x_1 + f_2 x_2 + f_3 x_3. \tag{4.17}$$

The condition for the payoff matrix A to be of (PD) type is:

$$a_{21} < a_{11} < a_{22} < a_{12} \tag{4.18}$$

$$a_{31} < a_{11} < a_{33} < a_{13} \tag{4.19}$$

$$a_{32} < a_{22} < a_{33} < a_{23} \tag{4.20}$$

and for definiteness, we choose the specific values:

$$A = \begin{bmatrix} 2 & 2.8 & 2.8 \\ 1.5 & 2.1 & 2.3 \\ 1.5 & 1.8 & 2.2 \end{bmatrix}$$
(4.21)

For the static game, it is easy to show that the cancerous state $(S, R_1, R_2) = (x_1, x_2, x_3) = (1, 0, 0) \equiv \mathbf{p}^{*T}$ is a strict Nash equilibrium in the absence of chemotherapy $(C_1 = 0, C_2 = 0)$ and an ESS for the replicator system, as shown in the diagram of figure 4.2(a) where the entire triangular region is the basin of attraction for the *S* population.

A seperate important ingredient in our model is our tumor-growth equation, which we take as:

$$\dot{x}_{tumor} = (\langle f \rangle - g) x_{tumor} \tag{4.22}$$

where the growth(decay) of the tumor is a function of the average fitness associated with the tumor minus a constant background fitness level g, associated with the surrounding tissue (say healthy cells) and micro-environment [123]. When the average fitness level of the population of cancer cells is higher than g, the tumor grows, and when it is lower, it regresses. Our chemotherapy functions C_1 and C_2 largely control this complex dynamic by modifying the Nash equilibria and ESS's of the system via the fitness function (4.12).

4.3 Constant chemotherapy

First, we examine the dynamical system for different constant levels of the chemoparameters C_1 , C_2 ($0 \le C_1 \le 1$; $0 \le C_2 \le 1$; $C_1 + C_2 \le 1$) in the case of additive e = 0 interactions, synergistic (e > 0) interactions, and antagonistic (e < 0) interactions.

4.3.1 Additive interactions e = 0

Figure 4.2 shows the panel of trajectories for three different scenarios. To start, figure 4.2(a) shows the trajectories with no chemotherapy – all trajectories lead to the S corner

which saturates the tumor. In these figures, solid circles denote the ESS states, while open circles denote the unstable states. Figure 4.2(b) shows trajectories with $C_1 = 0, C_2 = 0.8$. In this case, competitive release of the R_1 population allows it to take over the tumor, with all trajectories leading to the R_1 corner. For these parameter values, R_1 is the ESS (and strict Nash equilibrium) of the system. Figure 4.2(c) with $C_1 = 0.8, C_2 = 0$ depicts competitive release of the R_2 population, with all trajectories leading to the R_2 corner, showing that R_2 is the ESS of the system.

In figure 4.2(d) we plot three different constant therapy schedules together, showing how they can intersect at different times. This gives the possibility of switching the therapies off and on at the intersection times in order to create a trajectory that stays in a closed loop and never reaches any of the corners – these closed loops represent scenarios in which the three subpopulations stay in perpetual competition, driven by a time-dependent schedule that shapes the fitness landscapes in such a way as to manage chemotherapeutic resistance. We will describe the systematic construction of these closed loops in §4.4. Figure 4.3 shows the tumor size (using eqn (4.22)) plotted logarithmically. The untreated tumor ($C_1 = 0, C_2 = 0$) grows exponentially, while each of the treated tumors initially show tumor regression up until $t \sim 75$, then tumor recurrence due to chemotherapeutic resistance (either of the R_1 or R_2 populations), an unwelcome common scenario. This scenario is clearly depicted in prostate cancer data sets shown in [120] and widely discussed in the clinical literature. Our goal is to show how specially designed adaptive multi-chemotherapies can push the tumor recurrence point further to the right on these plots.

The full range of possible profiles are shown in figure 4.4. For certain ranges of chemodosing, there are mixed basins of attraction to each of the corners, hence multiple evolutionary stable states of the system. In the top row, with $C_1 = 0$, the two critical values where the ESS's bifurcate are $C_2 = 1/3$ (R_1 changes from unstable to stable via a transcritical bifurcation) and $C_2 = 1/2$ (S changes from stable to unstable via a transcritical bifurcation). In the left column, with $C_2 = 0$, the two critical values where the ESS's bifurcate are $C_1 = 7/18$ (R_2 changes from unstable to stable via a transcritical bifurcation) and $C_2 = 1/2$ (S changes from stable to unstable via a transcritical bifurcation). This panel shows the complexity associated with the stability and basins of attraction of the three tumor subpopulations even in such a simple model, but also gives us the opportunity to exploit the inherent underlying dynamics associated with the different trajectories using piecewise constant chemo-dosing protocols.



Figure 4.2: Depiction of evolutionary stable states (ESS) for different chemotherapy values. (a) With no chemotherapy, the tumor saturates to the S corner regardless of the initial make-up of the three subpopulations. (b) With $C_1 = 0$, and $C_2 = 0.8$, competitive release of the resistant population R_1 drives all trajectories to the R_1 corner. (c) With $C_1 = 0.8$ and $C_2 = 0$, competitive release of the resistant population R_2 drives all trajectories to the R_2 corner. (d) Trajectories associated with three different constant combinations of C_1 and C_2 , depicting the overlap of the trajectories at different times.



Figure 4.3: Tumor growth curves (log-plots) for untreated and constant therapies. Tumor recurrence (dashed line) occurs at $t \sim 75$ dimensionless time units.

4.3.2 Non-additive interactions

Using the panel in figure 4.2 as our guide, we focus in more detail on the values $C_1 = 0.35$, $C_2 = 0.27$, comparing the effects of synergistic interactions and antagonistic interactions with values e = -0.4, -0.3, -0.2, -01., 0, 0.1, 0.2, 0.3, 0.4 in figure 4.5. For strongly antagonistic interactions (e = -0.4, -0.3, -0.2), only the two subpopulations S and R_2 compete for dominance (both ESS), with the R_1 population being an unstable state. For strongly synergistic interactions (e = 0.2, 0.3, 0.4), the two subpopulations R_1 and R_2 compete for dominance (both ESS), with the population of cells S sensitive to both drugs being an unstable state. It is the intermediate regime e = -0.1, 0, 0.1 that is the most interesting, and has all three evolutionary stable populations competing for dominance (all three ESS) with intertwined basins of attraction for each. We examine the details of the bifurcations that occur between the two-species co-existence and three-species co-existing states next.



Figure 4.4: Panel showing the full range of trajectories, ESS and their basins of attraction for constant additive therapies in the range $0 \le C_1 \le 1$, $0 \le C_2 \le 1$, $C_1 + C_2 \le 1$. Bifurcation values along top row are $C_2 = 1/3$, $C_2 = 1/2$. Bifurcation values along left column are $C_1 = 7/18$, $C_1 = 1/2$.

Transcritical bifurcations

In figure 4.6, we focus on the range of antagonistic values e = -0.21, -0.19, -0.175, -0.15for fixed values of C_1 and C_2 . The relevant bifurcation that occurs at the critical value $e_c^{(a)} = -310/1701 = -0.1822$ takes place at the R_1 corner, when the fixed point $R_1 = 1$ goes from unstable ($e < e_c^{(a)}$) to asymptotically stable ($e > e_c^{(a)}$) in a transcritical bifurcation (exchange of stability). Through the bifurcation point, a stable fixed point (shown in figure 4.6(a) outside the triangle below R_1) moves up the $R_1 - S$ side of the triangle ($R_2 = 0$),



Figure 4.5: The effect of synergy (e > 0) vs. antagonism (e < 0) on the ESS and the basins of attraction for a fixed choice of constant combination therapies $C_1 = 0.35$, $C_2 = 0.27$. (a) e = -0.4; (b) e = -0.3; (c) e = -0.2; (d) e = -0.1; (e) e = 0; (f) e = 0.1; (g) e = 0.2; (h) e = 0.3; (i) e = 0.4.

and exchanges stability with the fixed point at the R_1 corner. In figure 4.7 we show the details of the collision of eigenvalues that takes place (figure 4.7(a)) and the process in the dS/dt vs. S plane (figure 4.7(b)-(e)). Figure 4.7(b) shows the classic transcritical bifurcation diagram (see [109]). When the level of antagonism is sufficiently large, there are only the two evolutionary stable states S and R_2 .

In figure 4.8 we highlight the bifurcation that takes place in the synergistic regime around values e = 0, 0.1, 0.2, 0.3 for fixed values of C_1 and C_2 . The transcritical bifurcation occurs

at the critical value $e_c^{(s)} = 30/189 = 0.1587$ takes place at the *S* corner, when the fixed point S = 1 goes from unstable $(e > e_c^{(s)})$ to asymptotically stable $(e < e_c^{(s)})$. In figure 4.9 we show the details of the bifurcation diagram that governs this synergistic transcritical bifurcation. When the level of synergism is sufficiently large, there are only the two evolutionary stable states R_1 and R_2 .

In figure 4.10 we show the areas of three basins of attraction through the full range of values $-0.3 \le e \le 0.3$. The basin areas begin to rapidly change in the antagonistic regime at e = -0.2 and, in general, show much more sensitivity to changes in e in the antagonistic regime than the synergistic regime.

Figure 4.11 shows the average fitness curves (figure 4.11(a)) and tumor growth curves (figure 4.11(b)) through a range of values of e. Notice, in the beginning, the average fitness of the antagonistic case e = -0.3 is higher, but ends up lower over time. The tumor growth curve in this case follows the same trend - initially tumor growth is highest, then ends up lowest. Tumor recurrence time for the antagonistic case is also pushed later in time ($t \sim 300$) showing the effectiveness of antagonistic drug interactions in terms of managing resistance.

4.4 Designing adaptive schedules

Many new possibilities can be created with time-dependent chemoschedules $C_1(t)$, $C_2(t)$, if we monitor the balance of the subpopulations and adaptively make changes at judiciously chosen time points. The basic idea is shown in figure 4.2(d) where we see how the trajectories associated with different constant chemotherapy schedules cross. At any of the crossing times, it is possible to switch from one trajectory to another by switching the values of C_1 or C_2 at those crossing times (typically termed bang-bang control [64]). This basic procedure allows us to design schedules that take us from any point A in the triangle to any other point B along the legs of a path that are separated by the switching times. Using multiple time-switching, we can also design trajectories that form closed loops and never converge to any of the corners.



Figure 4.6: The opening of the basin of attraction for R_1 via a transcritical bifurcation at e = -310/1701 = -0.1822, $C_1 = 0.35$, $C_2 = 0.27$. At the bifurcation point, the areas of the basins of attraction of the respective regions are S = 76.2%, $R_2 = 13.8\%$. (a) e = -0.21; (b) e = -0.19; (c) e = -0.175; (d) e = -0.15.

Figure 4.12 shows examples of how closed (piecewise differentiable) orbits are designed in practice. In each of the figures, point O is fixed, as is the untreated (blue) curve with $C_1 = 0, C_2 = 0$. Consider the loop OABO created figure 4.12(a), with additive interaction parameter e = 0. In traversing the OA leg, we use $C_1 = 0.5, C_2 = 0.2$. When the trajectory reaches point A, we switch to $C_1 = 0, C_2 = 0$, i.e. no therapy. When we reach point B, we switch to $C_1 = 0.2, C_2 = 0.5$ until we reach point O, and then we start the schedule again to traverse the same loop as many times as we desire. The dose density plot is shown in figure



Figure 4.7: Four ways of depicting the antagonistic transcritical bifurcation. (a) Eigenvalue collision that takes place as the two fixed points collide at the R_1 corner. The other eigenvalue remains negative for both fixed points; (b) Transcritical bifurcation diagram; (c) Pre-bifurcation phase plane; (d) Bifurcation phase plane; (e) Post-bifurcation phase plane.

4.12(b). Figures 4.12(c),(d) uses the same dosing values, but with e = 0.3 (synergistic). The loop in this case is larger (encloses more area), so there is a larger deviation in the subpopulations throughout the loop than there was for the additive case. Figure 4.12(e),(f) shows an example of an antagonistic e = -0.3 adaptive therapy loop. Table I shows the total dose D, time period T, and average dose D/T for each case. The antagonistic adaptive loop delivers the smallest total dose over the shortest time-period to achieve one closed loop, whereas the synergistic adaptive loop delivers the smallest average dose over the loop, since the time-period is longest.

Figure 4.13 shows the tumor growth curves for each of the adaptive schedules as compared with the untreated growth curve (exponential growth), and constant therapy curve (each



Figure 4.8: The closing of the basin of attraction for S via a transcritical bifurcation at e = 30/189 = 0.1587, $C_1 = 0.35$, $C_2 = 0.27$. At the bifurcation point, the areas of the basins of attraction of the respective regions are $R_1 = 72.3\%$, $R_2 = 17.7\%$. (a) e = 0; (b) e = 0.1; (c) e = 0.2; (d) e = 0.3.

e	-0.3	0	0.3
D	118.23	119.42	120.47
Т	178.7	192.5	206.7
D/T	0.6616	0.6204	05828

Table 4.1: Total dose (D), Time period (T), and Average Dose (D/T) associated with adaptive therapies with antagonistic, additive, and synergistic drug interactions.

shows tumor regression followed by recurrence). In each case, the adaptive loop overcomes recurrence, with the antagonistic schedule minimizing the tumor re-growth leg (AB) of the schedule.



Figure 4.9: Four ways of depicting the synergistic transcritical bifurcation. (a) Eigenvalue collision that takes place as the two fixed points collide at the S corner. The other eigenvalue remains negative for both fixed points; (b) Transcritical bifurcation diagram; (c) Prebifurcation phase plane; (d) Bifurcation phase plane; (e) Post-bifurcation phase plane.

A final metric that we use for comparisons is the rate of adaptation of the subpopulation i, defined by:

$$\alpha_i(t) = \frac{1}{t} \int_0^t x_i(s) ds, \qquad (4.23)$$

which is the time-average of subpopulation level $x_i(t)$.

Figure 4.14 shows the rates of adaptation associated with the two resistant subpopulations R_1 and R_2 during the course of the adaptive schedules. Notice the rate of adaptation is lowest for the antagonistic interaction, which is the main reason the tumor growth curve in figure 4.13(c) is most effective at controlling and delaying tumor recurrence as the cell population evolves.



Figure 4.10: Areas of the three basins of attraction as a function of the parameter e. Note the sensitivity of the S and R_1 regions in the antagonistic regime $-0.2 \le e \le -0.1$.



Figure 4.11: (a) Average fitness curves for e = 0.3, 0.15, 0, -0.15, -0.3; (b) Tumor growth curves (log-plot) for e = 0.3, 0.15, 0, -0.15, -0.3.

4.5 Discussion

The role of synergistic vs. antagonistic combination drug interactions on the dynamical balance of evolving subpopulations of cells is not simple to characterize. In general terms,



Figure 4.12: Closed loop adaptive schedules *OABO* using three cycles for e = 0, 0.3, -0.3. (a)e = 0, OA: $C_1 = 0.5, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.5$; (b) Corresponding adaptive schedule; (c) e = 0.3, OA: $C_1 = 0.5, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.5$; (d) Corresponding adaptive schedule; (e) e = -0.3, OA: $C_1 = 0.5, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.2, AB$: $C_1 = 0, C_2 = 0, BO$: $C_1 = 0.2, C_2 = 0.5$; (f) Corresponding adaptive schedule.



Figure 4.13: Tumor growth curves for the adaptive schedules from figure 4.12 as compared with untreated growth (exponential) and constant schedule which eventually leads to tumor recurrence. (a) e = 0; (b) e = 0.3; (c) e = -0.3.

antagonistic interactions are able to exert a more targeted and subtle range of influences on an evolving population than synergistic ones which, roughly speaking, cause the two drug combination to effectively act as one. This limits their flexibility for designing effective strategies to manage resistance, although might be useful in producing a bigger killing effect



Figure 4.14: Rates of adaptation for e = 0, 0.3, -0.3: (a) R_1 rate of adaptation; (b) R_2 rate of adaptation; (c) x_1 rate of adaptation; (d) x_2 rate of adaptation;

with a smaller dose. As a by-product of the ability to deliver a more nuanced effect on the balance of cells, the sensitivity to small changes in the relative doses of the two drugs in an antagonistic setting seems to be higher than in a synergistic setting (as shown best in figure 4.10). This, perhaps, makes these interactions harder to control. Tumor recurrence times for the antagonistic drug interactions are delayed more effectively than for synergistic interactions, consistent with the fact that adaptation rates are slower for antagonistic interactions. Anatagonistic drug interactions provide more flexibility in designing adaptive resistance management schedules. Additional features to our model could be the introduction of mutations that might occur in response to the chemotherapy, which potentially could be handled using a mutational replicator system [85], or a finite-cell Moran process based model [118]. It might also be possible to analyze existing individual patient data and tumor response curves (as was done in [120]) to design and optimize better multi-drug strategies retrospectively, which is ongoing by our group. There are also ongoing adaptive multi-drug clinical trials at the Moffitt Cancer Center that show promise in prostate cancer patients.

Chapter 5

Optimal control of evolutionary games: The two-by-two case

In this chapter, we focus on the optimal control problem in the replicator system with two species competing based on evolutionary games. The control u(t) in this problem represents an external/manual factor that alters the fitness landscapes, thus it alters the entries of the payoff matrix and changes the game type the species play. Additionally, we place an extra constraint on the integral of the control u(t), which indicates the external effect could only function for finite time. After the Pontryagin Maximum Principle is applied to the dynamical system according to an appropriate payoff function, we find that the standard replicator equations and the adjusted replicator equation may have different behaviors. For example, if we want to minimize the final frequency of the defector in the Prisoner's dilemma game, the standard replicator equation always has bang-bang control as optimal, while the adjusted replicator equation may have a very complex time-dependent function as the optimal solution. However, the adjusted replicator equation is a nonlinear re-scaling of the replicator equation. The comparison between two replicator systems highlights that adjusted replicator equations differs from the standard replicator equation when the timing is sensitive, especially in optimal control problems. The solution to this optimal control problem, therefore, helps to build the proper models for each individual problem in evolutionary games and it could contribute to the drug-scheduling, adaptive therapy, or resistance prevention when these models are applied to clinical data.

5.1 Introduction

In an evolutionary process, the final outcome may not be ideal. It is desirable to develop a control theory framework for the evolutionary process in order to select for particular phenotypes or genotypes [28]. Then modeling and optimizing this whole process requires control and optimal control techniques, and there are a lot of applications reported, such as prey-predator system[101], drug dosage design and resistance prevention[55, 13], or cooperation networks[2]. Particularly, a lot of literature have discussed about optimal chemotherapy scheduling for cancer with differential equations[73, 111, 125, 79, 22], or stochastic birth death process[126, 26].

This paper is also inspired by the optimization for chemotherapy, but we keep the analysis general, without including much detail associated with cancer modeling. In section 5.2 we use an evolutionary game model[108] to determine the fitness landscape of the two competing species. The control u(t) alters the fitness, thus it actually change the game type and the corresponding evolutionary stable strategy[108, 85, 52]. The system is then evolved with standard replicator equations or adjusted replicator equations. After that, we use the Pontryagin Maximum principle[67, 66, 98, 113] to find out the optimal control of the dynamic system in section 5.2.2.

In section 5.3, we analyze the properties of the optimal control, and we find that standard replicator equations and adjusted replicator equations may have different behavior. With appropriate target function for optimization, the standard replicator will have bang-bang control to be optimal, while the adjusted replicator equations may have more complex time-dependent optimal control. Numerical simulations of a prisoner's dilemma game[87, 85, 52] are provided in section 5.4. Because the Prisoner's dilemma game has the defector as the Nash equilibrium[87] while the defector is not optimal for maximizing payoff, we use the Prisoner's dilemma game as a cancer model[124] and controlling the defector is analog to controlling the tumor. Numerical results shows that standard replicator equation always has bang-bang optimal control to minimize the final defector. However, when the average fitness increases with the control value, the adjusted replicator equation will have a slowly increasing

function as the optimal control. The adjusted replicator equation has been proved to be the continuous limit of the Moran process [114, 115], and similar results has been reported [126] that high entropy performs the best with Moran process simulation. The reason of the difference is explained in section 5.45. High average fitness will weaken the effects of the control, and the bang-bang control in this case has the highest fitness when u(t) = 1, so it won't be optimal any more.

5.2 model

5.2.1 evolutionary games and replicator dynamics

Now consider there exist two competing species with x_1 and x_2 as the individual frequency, so the sum always equal to one $x_1 + x_2 = 1$. The selection is based on the fitness landscape of the system, and the individual fitness is determined by the 2-by-2 payoff matrix A:

$$f_i = (A\vec{x})_i \tag{5.1}$$

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$
(5.2)

where $\vec{x} = [x_1, x_2]^T$ and i = 1, 2. Suppose u(t) is the control for the system, and it could be any external influence such as drugs for some disease, disinfection for the bacterial, foods for one single animal, etc. The value of the control u(t) is always positives but smaller than 1 at any time.

$$0 \le u(t) \le 1 \tag{5.3}$$

The fitness will be altered when the control u(t) is none zero. Without losing generosity, we assume the altered selection tends to be in favor of x_1 , which means it will either reduce the fitness of x_2 or increase the fitness of x_1 . Suppose the change of fitness is proportional to the control and the coefficient is the linear combinations of the frequencies as shown in eq.5.4. α_1 and α_2 are constant. There's no constant in the coefficient because adding any constant is equivalent to adding it to both α_1 and α_2 .

$$\Delta f = (\alpha_1 x_1 + \alpha_2 x_2) u(t) \tag{5.4}$$

First if f_2 is reduced by Δf , then let's denote the altered fitness as f_2 . Because of the change of the fitness landscape, the game type, indicated by the new payoff matrix A', is varied as well. Reducing the fitness of the x_2 is equivalent to reducing the entries in the second raw. Similarly, if f_1 is increased, then entries in the first row of the payoff matrix will be increased as well.

$$f_{2}' = f_{2} - \Delta f \Leftrightarrow A' = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} - \alpha_{1}u & a_{22} - \alpha_{2}u \end{bmatrix}$$
 (5.5)

$$f_{1}' = f_{1} + \Delta f \Leftrightarrow A' = \begin{bmatrix} a_{11} + \alpha_{1}u & a_{12} + \alpha_{2}u \\ a_{21} & a_{22} \end{bmatrix}$$
(5.6)

Once we construct the fitness landscape, the system is simulated with either the standard replicator equation or the adjusted replicator equation. It is worth pointing out that both of them can be transformed into an single variable ODE by eliminating x_1 . We will use this result to simplify our system later.

$$\dot{x}_i = (f_i - \langle f \rangle) x_i \tag{5.7}$$

First let's consider the standard replicator equation in eq.5.7 where $\langle f \rangle = f_1 x_1 + f_2 x_2$ is the average fitness. It's quite simple to combine the formula of the fitness with the replicator equations. In eq.5.8 the standard replicator equation is rewritten with single variable x_2 . Here we only put one single equation because increasing f_2 and reducing f_1 by Δf lead to the same ordinary differential equations, thus the same trajectory.

$$\dot{x}_2 = (a_{21} - a_{11} - \alpha_1 u(t))(1 - x_2)^2 x_2$$
 (5.8)

$$+(a_{22}-a_{12}-\alpha_2 u(t))(1-x_2){x_2}^2$$
(5.9)

$$f_{2}' = f_{2} - \Delta f \Leftrightarrow f_{1}' = f_{1} + \Delta f$$
(5.10)

However, if we take the adjusted replicator equation in eq.5.11 for the simulation, we can find that increase f_2 or reducing f_1 are quite different because of the average fitness $\langle f \rangle$ in the denominator. In the next section we will discuss both for the adjusted replicator equation.

$$\dot{x}_i = \frac{(f_i - \langle f \rangle)}{\langle f \rangle} x \tag{5.11}$$

$$f_2' = f_2 - \Delta f \not\Leftrightarrow f_1' = f_1 + \Delta f \tag{5.12}$$

Since the control u(t) goes against x_2 , if we aim at minimizing the population of x_2 and there's no extra constraint on the control u(t), the optimal control will simply be u(t) = 1for all time. In reality, for example about drug scheduling, the total dosage the patient can take with a period is limited, and the task for the doctor is to distribute drugs over time for the patient. Therefore the extra constraint in eq.5.13 is put on the integral of the control u(t).

$$\int_0^T u(t)dt = D \tag{5.13}$$

Where D is a fixed constant and T is the terminal time. For simplicity, we denote x_2 as x as the single variable in the remaining sections.

$$x \equiv x_2 \tag{5.14}$$

$$x_1 = 1 - x \tag{5.15}$$

5.2.2 Pontryagin Maximum Principle

The Pontriagin Maximum(Minimum) Principle[67, 66, 98, 113] is a widely used method to solve optimal control problem. Assuming that our target is to minimize a general cost function in eq.5.16

$$minimize: \int_0^T L(x(t), u(t), t)dt + \varphi(x(T))$$
(5.16)

The standard form for the Maximum Principle is applied for dynamic system with boundary value constraint, so we need to convert the integral constraint in eq.5.13 into the standard form. Eq.5.17 defines a new variable d(t) so that the constraint is transformed into an ODE with two boundary values. Together with the standard replicator equation the new dynamic system is constructed in eq.5.21.

$$d (t) = \int_0^t u(s) ds$$
 (5.17)

$$d \quad (t) = u(t) \tag{5.18}$$

$$d (0) = 0 (5.19)$$

$$d \quad (T) = D \tag{5.20}$$

$$\dot{X} = F(X) \tag{5.21}$$

$$X = [x(t), d(t)]^{T}$$
(5.22)

$$F(X) = [(f_2 - \langle f \rangle)x, u(t)]^T$$
(5.23)

If we use the adjusted replicator equation rather than the standard ones, we only need to update F(X):

$$F(X) = \left[\left(\frac{(f_2 - \langle f \rangle)}{\langle f \rangle} x, u(t)\right]^T$$
(5.24)

Then following the notations from [113], we construct the Control Theory Hamiltonian as eq.5.25. λ_1 and λ_2 are the co-state function, or in Physics they correspond to momentum.

$$H(x, d, \lambda_1, \lambda_2, u) = \lambda^T F(x) + L(x, u, t)$$
(5.25)

$$\lambda = \left[\lambda_1, \lambda_2\right]^T \tag{5.26}$$

Assuming that $u^*(t)$ is the optimal control for this problem, with corresponding trajectory $x^*(t), d^*(t)$, there exist co-state functions $\lambda_1^*(t)$ and $\lambda_2^*(t)$ that satisfy the canonical equation 5.27

$$\dot{x}^{*}(t) = \frac{\partial H(x^{*}(t), d^{*}(t), \lambda_{1}^{*}(t), \lambda_{2}^{*}(t), u(t))}{\partial \lambda_{1}}$$
(5.27)

$$\dot{d}^{*}(t) = \frac{\partial H(x^{*}(t), d^{*}(t), \lambda_{1}^{*}(t), \lambda_{2}^{*}(t), u(t))}{\partial \lambda_{2}}$$
(5.28)

$$\dot{\lambda_1^*}(t) = -\frac{\partial H(x^*(t), d^*(t), \lambda_1^*(t), \lambda_2^*(t), u(t))}{\partial x}$$
(5.29)

$$\dot{\lambda_2^*}(t) = -\frac{\partial H(x^*(t), d^*(t), \lambda_1^*(t), \lambda_2^*(t), u(t)))}{\partial d}$$
(5.30)

with the corresponding boundary condition:

$$x(0) = x_0 (5.31)$$

$$d(0) = 0, d(T) = D \tag{5.32}$$

$$\lambda_1(T) = \frac{\partial \varphi(x(T))}{\partial x(T)}$$
(5.33)

At any time point, the optimal control $u^*(t)$ will minimize the Control Theory Hamiltonian as shown in eq.5.34

$$u^{*}(t) = \underset{u(t)}{\arg\min} H(x^{*}(t), d^{*}(t), \lambda_{1}^{*}(t), \lambda_{2}^{*}(t), u(t))$$
(5.34)

Now the optimization problem turns into the two-point boundary value problem, and the solution will be the optimal control and the corresponding trajectory.

5.3 Results

Now consider the case where the function L only has explicit dependence on the trajectory L = L(x(t)). We can use eq.5.34 to find some useful properties of the optimal control function $u^*(t)$. For standard replicator equation and adjusted replicator equation, we denote the Control Theory Hamiltonian as H_{std} and H_{adj} respectively in eq.5.35 and 5.36

$$H_{std} = \lambda_1 (f_2 - \langle f \rangle) x + \lambda_2 u + L(x(t))$$
(5.35)

$$H_{adj} = \lambda_1 \frac{(f_2 - \langle f \rangle)}{\langle f \rangle} x + \lambda_2 u + L(x(t))$$
(5.36)

In addition, the f_1 and f_2 are affine functions of the control u(t) no matter how you alter the fitness, thus the average fitness $\langle f \rangle$ is also an affine function of u(t). As a result, so does the Control theory Hamiltonian H_{std} for the standard replicator equation. To minimize H_{std} , the optimal control $u^*(t)$ depends on the sign of $\frac{\partial H_{std}}{\partial u}$. If $\frac{\partial H_{std}}{\partial u} \neq 0$ holds almost everywhere, the then optimal control will be the bang-bang control.

$$u^{*}(t) = \begin{cases} 0 & \frac{\partial H_{std}}{\partial u} > 0\\ 1 & \frac{\partial H_{std}}{\partial u} < 0 \end{cases}$$
(5.37)

For the case of the adjusted replicator equation, things are more complex and hard to see the minimal values of H_{adj} . Finding the minima is first step by taking the partial derivative of u(t) and setting it to 0 in eq.5.38:

$$\frac{\partial H_{adj}}{\partial u} = \frac{\lambda_1 x \left(\langle f \rangle \frac{\partial f_2}{\partial u} - f_2 \frac{\partial \langle f \rangle}{\partial u}\right) - \lambda_2 \langle f \rangle^2}{\langle f \rangle^2}$$
(5.38)

$$= 0$$
 (5.39)

As mentioned in section 5.2 $f_1, f_2, \langle f \rangle$ are all affine functions of u(t), so the solution to the minima in eq.5.38 is actually solving a quadratic equation. Also the minima may not be the global minimal. The control is bounded between 0 and 1 so even if the minima falls within

the range, we still have to compare the minima with the boundary value at 0 and 1. In conclusion, it's possible that the optimal control $u^*(t)$ is still bang-bang control, but for some cases it may be some value between 0 and 1, different from the standard replicator equation. We will give an example in section 5.4

5.4 discussion

Now we consider the case of chemotherapy dosage scheduling for cancer. The game between the cancer cells and surroundings healthy cells is a Prisoner's dilemma game[124] in eq.5.40. The cancer cells play as a defector and the healthy cells are the cooperator. In a Prisoner's dilemma game the defector, thus the cancer cell, is an evolutionary stable strategy (ESS) and Nash Equilibrium (NE)[87, 85, 52], but the average fitness for all cancer cell will be lower than the one with all healthy cells. We want to find the optimal scheduling of the chemotherapy $u^*(t)$ so that the frequency of cancer cells is minimized at the end of the therapy. The corresponding cost function is updated as eq.5.41

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 1 \end{bmatrix} \tag{5.40}$$

$$L = 0 \tag{5.41}$$

$$\varphi(x(T)) = x(T) \tag{5.42}$$

5.4.1 f_2 decrease

Now we consider the altered game in eq.5.43, and the fitness of the cancer cells is reduced with $\alpha_1 = 5$ and $\alpha_2 = 0$. When u > 1/4, the game type changes from prisoner's dilemma game to coordinate game, and x = 1 becomes an evolutionary stable strategy for the system. Before we find the optimal solution for this problem, I want to introduce two pre-determined therapy schedule. One is the Maximum tolerated Dose (MTD), which provides the patient as high dose as possible in the beginning then stops the drug, and the Low dose Metronomic therapy, which gives the patient constant dosage throughout the treatment.

$$A' = \begin{bmatrix} 3 & 0\\ 4 - 4u(t) & 1 \end{bmatrix}$$
(5.43)



Figure 5.1: Schema for the Maximum tolerated Dose (MTD) and Low Dose Metronomic (LMD) schedule. The integral of these two schedule are the same.

Fig.5.1 is an example of the MTD versus LDM with T = 1, D = 0.3. The corresponding numerical simulations are in Fig.5.2. We can see that in this case the MTD works better than the LDM for both standard replicator equation and adjusted replicator equation. We know that unless singular control exists, bang-bang control will be optimal for standard replicator equation. After solving the two point boundary value problem from eq.5.27, the optimal control for both ODE system is the bang-bang control, and the schedule is exactly the same as the MTD therapy in Fig.5.1. The corresponding trajectory $x^*(t)$, therefore, is just the same as the MTD curves in Fig.5.2(a) and Fig.5.2(b).


Figure 5.2: Simulation results with D = 0.3, T = 1, x(0) = 0.1 and payoff matrix in eq.5.43 (a) the corresponding trajectory for MTD and LDM simulated by standard replicator equation.(b)the corresponding trajectory for MTD and LDM simulated by Adjusted replicator equation.



Figure 5.3: Bang-bang optimal control for both standard replicator equation and adjusted replicator equation

5.4.2 f_1 increases

However, as we discussed in section 5.3, the bang-bang control may not be optimal for adjusted replicator equation. Here we consider increasing the fitness of x_1 by the same amount as in section 5.4.1, and the new payoff matrix is shown in eq.5.44 with increased entries in the first row.

$$A' = \begin{bmatrix} 3+4u(t) & 0\\ 4 & 1 \end{bmatrix}$$
(5.44)

Because we have proved in section 5.2 that the increasing the first row or decreasing the second row are the same for standard replicator equation, the corresponding trajectories for MTD and LDM schedule should be identical to Fig.5.2(a). The optimal control for the standard replicator equation will just be bang-bang control, the same as Fig.5.3, so we will focus on the Adjusted replicator equation for this section.

Fig.5.4 is the trajectories simulated with adjusted replicator equation with D = 0.3, T = 1, x(0) = 0.1. Contradictory to what we get in section 5.4.1, the LDM method works better than MTD schedule so the MTD is no longer the optimal control. Similar results simulated by Moran Process has been reported that high entropy schedule, correlated with LDM methods, outperform other schedules. After we numerically solve the two point boundary value problem in eq.5.27 based on Pontryagin Maximum principle, we have the optimal control $u^*(t)$ and trajectory $x^*(t)$ in Fig.5.5

The optimal control $u^*(t)$ is non-trivial in this case as an slowly decreasing function, but the value is always close to 0.3 so it is more similar to the LDM schedule and that's the reason why the LDM has better performance than MTD method, or the regular bang-bang control.



Figure 5.4: Simulation results with D = 0.3, T = 1, x(0) = 0.1 and payoff matrix in eq.5.44 with Adjusted replicator equation.



Figure 5.5: (a)optimal control $u^*(t)$ for adjusted replicator equation (b) the corresponding trajectory $x^*(t)$

5.4.3 Time re-scaling

To convert the adjusted replicator equation to the standard replicator equation, we need the transformation[114] for the timing in eq.5.45

$$ds = \frac{dt}{\langle f \rangle} \tag{5.45}$$

$$\frac{dx_i}{ds} = (f_i - \langle f \rangle)x_i \tag{5.46}$$

This transformation actually varies the time structure, so after that the terminal time T and the integral value D are varied as well. Eq.5.47 provides the new terminal time as Ts and Eq.5.48 gives the new integral under the new time coordinate s. We can see that if the average fitness $\langle f \rangle$ increase a lot when the control u(t) increases, then after the transformation the time range of high u(s) value will be shorten, the integral of dose value u(s) is also reduced. For example, for bang-bang control, the new time range for u(s) = 1 will be relatively shorter but the new time range for u(s) = 0 will become longer. That's why in section 5.4.2 the bang-bang control doesn't work very well because increasing the control u(t) dramatically increases the average fitness $\langle f \rangle$. In the other hand, if the average fitness $\langle f \rangle$ decreases with the control u(t), the bang-bang control could be enhanced.

$$T^{s} = \int_{0}^{T^{s}} ds = \int_{0}^{T} \frac{dt}{\langle f \rangle}$$
(5.47)

$$D^s = \int_0^{T^s} u(s)ds = \int_0^T \frac{u(t)dt}{\langle f \rangle}$$
(5.48)

5.5 Conclusion

In this paper we solve the optimal control problem for 2-by-2 evolutionary game evolving with standard replicator equation or adjusted replicator equation. After we simplify the problem into the standard form, the optimal control is found by the Pontryagin Maximum principle. We take the prisoner's dilemma game as the numerical example and the standard replicator equation always has bang-bang control as optimal to reduce the defector, while the adjusted replicator equation, possibly Moran process as well, may have a no-trivial time dependent function as optimal control, especially when the average fitness of the system increases with the control value. This results take a deeper insight into the evolutionary game models, and it will potentially help to find the optimal drug scheduling as a medical problem.

Chapter 6

Optimal control of levitation by blowing

Is it possible to levitate an object with constant blowing? We show that it is not possible to levitate the object from the origin to a pre-specified height x^* using constant blowing, because the point at which the downward gravitational force balances the upward blowing force is an elliptic fixed point of the governing equations, so there is no transfer trajectory that connects the origin to x^* . To overcome this, we design time-dependent blowing schedules that achieve the transfer borrowing on orbital transfer ideas developed in the orbital mechanics literature. Then we ask which of these schedules are optimal. We show that, generally, it is bang-bang (on-off) blowing schedules that achieve the transfer in minimal time, using minimal energy, and minimal air volume. For certain parameter values, however, there are more complicated blowing schedules that are optimal (with respect to energy) which can be designed using the Pontryagin Maximum Principle (PMP) and singular control. The model we describe is frictionless and treats the object (i.e. a ping-pong ball) as a point mass.

6.1 Introduction

Is it possible to levitate a ping-pong ball from rest to a fixed height by constant blowing? At first glance, one might imagine the answer is as simple as supplying a blowing force from below that exactly balances the gravitational force pushing down at the height x^* where you want the ball to levitate as shown in figure 6.1. An experiment, such as the one shown here: https://www.youtube.com/watch?v=bCRjPFhlSYk might even confirm that it is achievable, until one examines the problem more closely. Let the ball start from rest, at

x = 0 where the blower is located. If one supplies a *constant* blowing force that is required to balance the ball at height $x^* > 0$, it turns out, surprisingly, that it is not possible to drive the ball to the position required for the forces to balance. We will show why, and describe how to overcome the problem by using *non-constant* blowing schedules designed appropriately. We formulate the problem as one in nonlinear dynamics and optimal control and show how some blowing schedules are better than others. The methods are borrowed from orbit transfer ideas developed initially in the orbital mechanics literature [60] in which intersecting phase space trajectories at different fixed parameter values are pieced together by switching from one parameter value to the other at precisely the time at which the trajectories intersect, creating a time-dependent control protocol that drives the system to a desirable state not achievable without the switching time. This idea is also currently being developed to design time-dependent chemotherapeutic schedules that achieve better results than standard protocols, making use of the replicator dynamical system as a tumor growth model in the context of evolution. In this paper, our goal is to use the simplest possible model to elucidate how the method can work in settings more commonly developed using more standard control theory ideas, and we show that the methods yield equivalent results.

Using air to levitate objects (as opposed to acoustic or magnetic levitation) is a common method of conveyance and has been explored in micro-manufacturing settings as discussed in [59], but to our knowledge, the observation that non-constant blowing is required to levitate to a fixed height has not been discussed in the literature. Reference [24] discusses the design of a low-cost air levitation system for teaching control engineering.

6.2 The model

The model we consider is the simple dimensionless nonlinear system:

$$m\ddot{x} = -mg + mgb(t)\exp(-\alpha x) \tag{6.1}$$



Figure 6.1: Ping-pong ball levitated by blowing through a straw. The streaming air blows upwards along the positive x direction against the gravity. x^* marks the equilibrium position where the forces balance. For full video, see https://www.youtube.com/watch?v=bCRjPFhlSYk

Here, b(t) is a dimensionless blowing parameter (proportion to amount of air per unit of time) that we will use to levitate the ball from rest (v(0) = 0), x(t) is the vertical height, m is mass, and $\alpha > 0$ is a constant that controls the distance over which the blower's influence is felt. The ball starts at position x(0) = 0 where the blower is located, as shown in figure 6.1 and the goal is to levitate it to height marked x^* . Our assumption in (6.1) is that the influence of the blowing decreases exponentially from the location of the blower (x = 0), but our results are not sensitive to that choice, as long as the influence decreases monotonically with height. We also assume there is no frictional loss of energy as the ball levitates and that the ball does not rotate, hence we treat it as a point mass.We emphasize that our model is simple in that it does not take into account the detailed fluid mechanical forces and dynamics at the jet-ball contact point (that would create such effects as dissipation and rotation of the ball). Our model is just complex enough to show that (i) it is not possible to drive the ball from its resting state at the base to the levitated state at height x^* with constant blowing, and (ii) it is possible to drive the ball from its resting state to the levitated state by piecing together time-dependent blowing schedules with precise switching times determined from crossing phase-space trajectories. To levitate the ball, the upward blowing force must exceed the downward gravitational force at x = 0, which implies that the right hand side of eqn (6.1) be greater than zero (at x = 0). The condition for this is that b > 1.

The governing Hamiltonian for the system is given by:

$$H(x,v) = \frac{1}{2}mv^2 + mgx + \frac{mgb\exp\left(-\alpha x\right)}{\alpha} \equiv E_b$$
(6.2)

where we highlight the fact that b is a parameter we will vary in order to change energy values. Notice for fixed x and v, it is evident that $E_{b_1} > E_{b_2}$ when $b_1 > b_2$ since eqn (6.2) is linear in b. The corresponding Hamiltonian equations of motion are:

$$m\dot{x} = \frac{\partial H}{\partial v} = mv \tag{6.3}$$

$$m\dot{v} = -\frac{\partial H}{\partial x} = -mg + mgb \exp(-\alpha x).$$
 (6.4)

For any given b, the forces balance (i.e. acceleration is zero) when:

$$x \equiv x^* = \frac{\ln\left(b\right)}{\alpha} \tag{6.5}$$

where the right hand side of eqn (6.1) is zero. With the condition b > 1, the equilibrium position is positive $x^* > 0$.

The phase space diagram in the (x, v) plane, obtained from plotting level curves $H(x, v) = E_b = constant$, is shown in figure 6.2(a) for different energy values. The elliptic fixed point where the forces balance correspond to:

$$E_b^* = H(x^*, 0) = \frac{mg}{\alpha} (\ln(b) + 1).$$
(6.6)

At the origin where the blower is located, the energy value is:

$$E_b^0 = H(0,0) = \frac{mgb}{\alpha} > E_b^*.$$
(6.7)

We show in figure 6.2(b) curves E_b^0 for increasing values of b. We also show in figure 6.2(c) the family of free-fall trajectories (b = 0) associated with energy values $E_0 = \frac{1}{2}mv^2 + mgx$ which start at $(x, v) = (0, v^*)$.

To compute the point x_{max} shown in figure 6.2(a), we follow the phase curve that passes through the origin to the point x_{max} , using the fact that the Hamiltonian is constant, which gives rise to the transcendental equation for x_{max} :

$$\alpha x_{max} = b(1 - \exp(-\alpha x_{max})). \tag{6.8}$$

The left and right sides of eqn (6.8) are plotted in figure 6.2(d) where the intersection defines x_{max} . Because x^* is an elliptic fixed point, it is not possible to traverse from the origin to this point on a phase curve, i.e. there is no constant energy trajectory that passes both through the origin and x^* . Instead, the levitating ball starting at the origin will reach height $x_{max} > x^*$, then drop back down to the origin in a periodic orbit, always overshooting height x^* where we want to position the ball.

For future use, notice the velocity formula from eqn (6.2) is:

$$v = \pm \sqrt{\frac{2}{m} \left[E_b - mgx - \frac{mgb}{\alpha} \exp(-\alpha x) \right]}.$$
(6.9)

6.3 Time-dependent blowing

To levitate the ball from the origin to x^* we need to drop the energy from its initial energy value E_b^0 to the final energy E_b^* by altering our blowing parameter b. Consider the family of phase curves for two different choices of the blowing parameter $b = b_1 > b_2 > 0$, as shown in figure 6.3. The diagram depicts two different paths to traverse from point O to point x^* . The blue curve corresponds to energy value $E_b = E_b^0$ for value $b = b_1$, giving rise to the elliptic balance point $x^*|_{b_1} = \frac{\ln(b_1)}{\alpha}$. The red curve that intersects $x^*|_{b_1}$ corresponds to energy value $E_b = E_b^0$ obtained by choosing blowing parameter $b = b_2 < b_1$, where $x_{max}|_{b_2} = x^*|_{b_1}$. To traverse the red elliptical path that connects O to the resting point x^* , we first use blowing parameter value $b_2 > 1$ until the ball arrives at x^* , say at time $t = t_1$. Then we can balance it there if we instantaneously change to value $b_1 > b_2$ at t_1 , which traps the ball at the elliptic fixed point. We call this a *two-step* schedule since it requires choosing two different value of b with one switch-time t_1 . To calculate t_1 , we need to integrate eqn (6.9) from O to x^* :

$$t_{1} = \int_{0}^{\ln(b_{1})/\alpha} \frac{dx}{\sqrt{\frac{2}{m} \left[E_{b_{2}}^{0} - mgx - \frac{mgb_{2}}{\alpha} \exp(-\alpha x) \right]}}$$
(6.10)

A useful quantity to track is the energy:

$$\langle E \rangle = \int_0^{t_f} E(t) dt \tag{6.11}$$

required to achieve a given transfer, where t_f is the time the ball reaches x^* . In this case, we have:

$$\langle E \rangle = E_{b_2}^0 \cdot t_1 \tag{6.12}$$

$$= \frac{mgb_2}{\alpha} \cdot t_1. \tag{6.13}$$

A second quantity we track is the area under the b(t) curve:

$$\langle b(t) \rangle = \int_0^{t_f} b(t) dt \tag{6.14}$$

which we think of as the total volume of air used to transfer the ball through time t_f . In this case, we have:

$$\langle b(t) \rangle = b_2 \cdot t_1. \tag{6.15}$$

A second path that takes the ball from O to x^* traverses through point A. Consider the piecewise differentiable path $O \longrightarrow A \longrightarrow x^*$. On the first piece from $O \longrightarrow A$ we launch the ball using blowing parameter value b_1 until time t_1 when we reach point A. At that time, we set b = 0 and let the ball free-fall along path $A \longrightarrow x^*$ until time t_2 when we arrive at x^* . For this, we use the free-fall trajectory given by:

$$x = -\frac{1}{2}gt^2 + v_A t + x_A (6.16)$$

$$v = -gt + v_A \tag{6.17}$$

where (x_A, v_A) mark the position and velocity coordinates at point A in figure 6.3. We know from eqn (6.17) that v = 0 (maximium height) at time $t = v_A/g$. Therefore for the total time to traverse the free-fall trajectory, we have:

$$t_2 - t_1 = v_A/g. (6.18)$$

Here, we use:

$$t_{1} = \int_{0}^{x_{A}} \frac{dx}{\sqrt{\frac{2}{m} \left[E_{b_{1}}^{0} - mgx - \frac{mgb_{1}}{\alpha} \exp(-\alpha x) \right]}}.$$
 (6.19)

Eliminating t in eqns (6.16),(6.17) gives the parabolic curve:

$$x = -\frac{v^2}{2g} + \left(\frac{v_A^2}{2g} + x_A\right)$$
(6.20)

beginning at (x_A, v_A) and ending at the point $(x^*, 0)$, where:

$$x^* = \left(\frac{v_A^2}{2g} + x_A\right) = \frac{\ln(b_1)}{\alpha}.$$
 (6.21)

At time t_2 we then let:

$$b = b_1 = \exp\left[\alpha \left(\frac{v_A^2}{2g} + x_A\right)\right] \tag{6.22}$$

to trap the ball at this resting point. We call this a *three-step* sequence since it requires three sequential values of b ($b = b_1$; b = 0; $b = b_1$) and two switching times t_1 and t_2 . The total energy along the trajectory is:

$$\langle E \rangle = E_{b_1}^0 \cdot t_1 + E_0 \cdot (t_2 - t_1)$$
 (6.23)

$$= \frac{mg}{\alpha} \left[b_1 \cdot t_1 + \ln(b_1)(t_2 - t_1) \right], \tag{6.24}$$

where E_0 is the free-fall energy with b = 0, and $(x, v) = (x^*, 0)$, i.e. $E_0 = mgx^*$. For this trajectory, we have:

$$\langle b(t) \rangle = b_1 \cdot t_1 + 0 \cdot (t_2 - t_1) = b_1 \cdot t_1. \tag{6.25}$$

Path	Time t_f	Energy $\langle E \rangle$	Air volume $\langle b(t) \rangle$
$O \longrightarrow x*$	1.036	21.83	2.227
$O \longrightarrow A \longrightarrow x*$	0.6662	16.84	0.7482

Table 6.1: Time, Energy and Air volume values for two trajectories shown in figure 6.3. The results are simulated using variables: $b_1 = 6, \alpha = 1m^{-1}, g = 9.8m/s^2, m = 1kg$.

It is clear that once we allow for a time-dependent blowing function b(t), there exist many schedules that allow us to achieve the goal of moving the ball from the origin to the resting point x^* , each taking a different total time to achieve the transfer and each expending a different amount of energy and air. We show in Table 6.1 numerically computed values for the two different paths described in figure 6.3 which shows in all cases, the three-step path: $O \longrightarrow A \longrightarrow x^*$ (which includes a free-fall segment) is more efficient than the more direct two-step path: $O \longrightarrow x^*$. A more comprehensive panel of schedules, using different fixed values of the parameter b, and the trajectories that achieve the transfer to the levitation point, is shown in figure 6.4.

We are now is a position to formulate the problem as one in optimal control theory to ask, of all possible schedules, which one achieves the transfer from O to x^* in minimal time? Which transfer uses minimal energy? Which transfer uses minimal amounts of air?

6.4 Optimal blowing schedules

Consider the problem of finding time-dependent blowing schedules b(t) and the corresponding trajectories (x(t), v(t)) in the phase plane that accomplish the task of moving the ball from (0, 0) to $(x^*, 0)$ in total time t_f , subject to the equations:

$$m\ddot{x} = -mg + mgb(t)\exp(-\alpha x) \tag{6.26}$$

$$x(0) = 0; \dot{x}(0) = 0 \tag{6.27}$$

$$x(t_f) = x^*; \dot{x}(t_f) = 0.$$
 (6.28)

We use the fact that $x \ge 0$, and we assume the blowing function is bounded from below and above: $0 \le b(t) \le b_{max}$.

6.4.1 Minimum time strategy

The transfer time t_f is given by:

$$t_f = \int_0^{t_f} dt = \int_0^{x^*} \frac{dx}{dx/dt} = \int_0^{x^*} \frac{dx}{v}.$$
 (6.29)

To minimize t_f , we need to minimize the area under 1/v(x) from $0 < x < x^*$, as shown in eqn (6.29). We claim that a *bang-bang* strategy, in which we choose $b = b_{max}$ initially, until switching time t_s at the point $(x, v) = (x_s, v_s)$, when we choose a free-fall trajectory b = 0, accomplishes the minimization. To see this, consider the trajectories associated with the two extreme (constant) values b = 0 and $b = b_{max}$ as shown in figure 6.5(a). The lower (red) trajectory is given by the curve $E_{b_{max}}^0$. At the switching time $t = t_s$, we switch values of the blowing parameter to b = 0, so that the ball travels on a free-fall trajectory E_0 to the final resting point $(x^*, 0)$ at time t_f . These curves define a sector in the first quadrant of the (x, v) plane as shown in the figure. We show in figure 6.5(b) the same curves in the 1/v vs. x plane and consider the areas under the curve using $E_{b_{max}}^0$ from $0 < x < x_s$, and E_0 from $x_s < x < x^*$. In the interval $0 < x < x_s$, for any (v(t), x(t)) associated with the schedule b(t), the following function will be used to derive the maximum velocity v_{max} for each fixed x:

$$\frac{d}{dt}\left(\frac{1}{2}v(t)^{2} + gx(t) + \frac{b_{max}g}{\alpha}e^{-\alpha x(t)}\right)$$
(6.30)

$$= (b(t) - b_{max}) vg e^{-\alpha x(t)}$$
(6.31)

$$\leq 0$$
 (6.32)

Thus the initial value gives the upper bound:

$$\frac{1}{2}mv(t)^{2} + mgx(t) + \frac{mb_{max}g}{\alpha}e^{-\alpha x(t)}$$
(6.33)

$$\leq \frac{1}{2}mv(0)^{2} + mgx(0) + \frac{mb_{max}g}{\alpha}e^{-\alpha x(0)}$$
(6.34)

$$= \frac{mb_{max}g}{\alpha} = E^0_{b_{max}} \tag{6.35}$$

103

From this, we know that v(t) is bounded by

$$v \le \sqrt{\frac{2}{m} \left[E_{b_{max}}^0 - gx - \frac{b_{max}g}{\alpha} e^{-\alpha x} \right]} = v_{max}$$
(6.36)

Hence:

$$\int_{0}^{x_{s}} \frac{dx}{v_{max}} < \int_{0}^{x_{s}} \frac{dx}{v}.$$
(6.37)

Similarly, in the interval $x_s < x < x^*$:

$$\frac{d}{dt}\left(\frac{1}{2}v(t)^2 + gx\left(t\right)\right) \tag{6.38}$$

$$= b(t) vg \tag{6.39}$$

$$\geq 0$$
 (6.40)

Thus the final value gives the upper bound:

$$\frac{1}{2}v(t)^2 + gx(t) \tag{6.41}$$

$$\leq \frac{1}{2}v(t_f)^2 + gx(t_f)$$
 (6.42)

$$= gx^* \tag{6.43}$$

From this, we know that v(t) is bounded by

$$v \le \sqrt{2\left[gx^* - gx\right]} = v_{max} \tag{6.44}$$

Hence:

$$\int_{x_s}^{x^*} \frac{dx}{v_{max}} < \int_{x_s}^{x^*} \frac{dx}{v}.$$
 (6.45)

Therefore the minimum-time transfer is achieved using the bang-bang strategy. An example of an optimal time trajectory and schedule is shown in figure 6.6.

6.4.2 Minimum air volume strategy

To minimize eqn (6.14), first consider that:

$$\dot{v} = \frac{dv(x)}{dx}\frac{dx}{dt} = v'(x)v \tag{6.46}$$

$$= -g + gb(t)\exp(-\alpha x), \tag{6.47}$$

which gives:

$$b = \frac{v'(x)v+g}{g\exp(-\alpha x)}.$$
(6.48)

From this, we have:

$$\int_0^{t_f} b(t)dt \tag{6.49}$$

$$= \int_0^{x_f} b(t) \frac{dx}{v} \tag{6.50}$$

$$= \int_{0}^{x_{f}} \frac{v'(x)v + g}{g \exp(-\alpha x)} \frac{dx}{v}$$
(6.51)

$$= \int_{0}^{x_{f}} \frac{v'(x)}{g} \exp(\alpha x) + \int_{0}^{x_{f}} \frac{\exp(\alpha x)}{v} dx$$
(6.52)

The first term can be integrated by parts, and using the fact that $v(0) = v(x_f) = 0$ gives:

$$\int_{0}^{t_{f}} b(t)dt = \int_{0}^{x_{f}} \frac{\exp(\alpha x)}{v} dx \qquad (6.53)$$
$$- \int_{0}^{x_{f}} \frac{\alpha v}{g} \exp(\alpha x) dx.$$

It is easy to see that maximizing v(x) will minimize both integrals, which gives the same result as the bang-bang strategy to minimize time (shown in figure 6.6).

6.4.3 Minimum energy strategy

The problem of finding a minimum energy strategy is more complicated. We start with:

$$\int_{0}^{t_{f}} Hdt = \int_{0}^{t_{f}} (\frac{1}{2}v^{2} + gx + \frac{bg\exp(-\alpha x)}{\alpha})dt$$
(6.54)

$$= \int_{0}^{t_{f}} (\frac{1}{2}v^{2} + gx + \frac{\dot{v} + g}{\alpha})dt$$
(6.55)

$$= \int_{0}^{x^{*}} \left(\frac{1}{2}v^{2} + gx + \frac{\dot{v} + g}{\alpha}\right) \frac{dx}{v}$$
(6.56)

$$= \int_{0}^{x^{*}} \left(\frac{1}{2}v^{2} + gx + \frac{\frac{dv}{dx}v + g}{\alpha}\right) \frac{dx}{v}$$
(6.57)

$$= \int_{0}^{x^{*}} (\frac{1}{2}v^{2} + gx + \frac{g}{\alpha})\frac{dx}{v} + \int_{0}^{x^{*}} \frac{dv}{\alpha}$$
(6.58)

The last integal is simply $v(x^*) - v(0) = 0$, using the boundary conditions. So,

$$\int_{0}^{t_{f}} Hdt = \int_{0}^{x^{*}} f(v(x))dx, \qquad (6.59)$$

$$f(v) = (\frac{1}{2}v^2 + gx + \frac{g}{\alpha})/v.$$
 (6.60)

To minimize this integral, we consider each fixed, x, and minimize with respect to v(x), i.e.

$$\frac{\partial f}{\partial v} = \frac{1}{2} - \frac{gx + \frac{g}{\alpha}}{v^2} = 0 \implies v(x) = \sqrt{2gx + \frac{2g}{\alpha}}.$$
(6.61)

However, the value above may not be achievable for all x, e.g we have v(x = 0) = 0, so

$$v(x) = \begin{cases} \sqrt{2gx + 2g/\alpha} & \sqrt{2gx + 2g/\alpha} \le v_{max} \\ v_{max} & \sqrt{2gx + 2g/\alpha} \ge v_{max} \end{cases}$$
(6.62)

According to figure 6.7, if the switching point (v_s, x_s) is on the right side of the curve $v(x) = \sqrt{2gx + \frac{2g}{\alpha}}$, then the extra transfer segment will exist connecting the b_{max} trajectory with the b = 0 trajectory. This gives a sufficient condition:

$$v_s > \sqrt{2gx_s + \frac{2g}{\alpha}} \tag{6.63}$$

106

where

$$x_s = -\frac{\ln(1 - \frac{\alpha x^*}{b_{max}})}{\alpha} \tag{6.64}$$

$$v_s = \sqrt{2gx^* - 2gx_s} \tag{6.65}$$

After we simplify the equations:

$$x_s < \frac{x^* - \frac{1}{\alpha}}{2} \tag{6.66}$$

Meanwhile, the inequality is the necessary condition as well because the equations:

$$x_s \geq \frac{x^* - \frac{1}{\alpha}}{2} \tag{6.67}$$

$$v_{max} \geq \sqrt{2gx + 2g/\alpha} \tag{6.68}$$

have no solution in the interval $0 < x < x_s$. In summary eqn (6.66) is both a necessary and sufficient condition for the existence of the extra transfer segment shown in figure 6.7(a).

6.5 Pontryagin Maximum Principle

The Pontryagin maximum (or minimum) principle (PMP) [91] is the most widely used method to find the optimal control scheme for certain dynamical systems. We will first give a statement of the PMP specifically for our problem and for proofs or more general versions, see reference [91, 66, 98, 113]. Following the notation of [113], consider a general dynamical system with fixed initial condition \mathbf{X}_0 , fixed final condition \mathbf{X}_f , and control b(t):

$$\dot{\mathbf{X}}(t) = \mathbf{f}(\mathbf{X}(t), b(t)) \tag{6.69}$$

$$\mathbf{X}(0) = \mathbf{X}_{\mathbf{0}} \tag{6.70}$$

$$\mathbf{X}(t_f) = \mathbf{X}_{\mathbf{f}} \tag{6.71}$$

$$0 \le b(t) \le b_{max} \tag{6.72}$$

And the target is to minimize some objective function:

$$\min_{b(t)\in[0,b_{max}]} \int_0^{t_f} L(\mathbf{X}(t), b(t)) dt.$$
(6.73)

First, we construct the control theory Hamiltonian function \mathbf{H} by introducing the Lagrange multiplier vector $\boldsymbol{\lambda}$:

$$\mathbf{H}\left(\mathbf{X}(t),\lambda(t),b\left(t\right)\right) = L + \boldsymbol{\lambda}^{T}\mathbf{f}$$
(6.74)

The dynamical system for this optimal control problem be same as eqn (6.3) and eqn (6.4) so we have:

$$\mathbf{X} = (x(t), v(t))^{T}$$
(6.75)

$$\mathbf{f} = (v(t), -g + gb(t) \exp(-\alpha x(t)))^{T}$$
(6.76)

Suppose $\boldsymbol{\lambda} = (\lambda_x, \lambda_v)^T$, the control theory Hamiltonian function for this problem will be:

$$\mathbf{H} = L + \lambda_x v + \lambda_v (-g + gb(t) \exp(-\alpha x))$$
(6.77)

Assuming that $b^*(t)$ is our optimal control function, with the corresponding optimal trajectory $x^*(t), v^*(t)$, there exist functions λ_x^*, λ_v^* , which satisfy the canonical equations of Hamilton:

$$\dot{x}^{*} = \frac{\partial \mathbf{H} \left(x^{*}, v^{*}, \lambda_{x}^{*}, \lambda_{v}^{*}, b^{*}\right)}{\partial \lambda_{x}}$$

$$\dot{v}^{*} = \frac{\partial \mathbf{H} \left(x^{*}, v^{*}, \lambda_{x}^{*}, \lambda_{v}^{*}, b^{*}\right)}{\partial \lambda_{v}}$$

$$\dot{\lambda}_{x}^{*} = -\frac{\partial \mathbf{H} \left(x^{*}, v^{*}, \lambda_{x}^{*}, \lambda_{v}^{*}, b^{*}\right)}{\partial x}$$

$$\dot{\lambda}_{v}^{*} = -\frac{\partial \mathbf{H} \left(x^{*}, v^{*}, \lambda_{x}^{*}, \lambda_{v}^{*}, b^{*}\right)}{\partial v}.$$
(6.78)

The optimal control $b^*(t)$ will minimize the control theory Hamitonian **H** at any time point:

$$b^{*}(t) = \underset{b(t)\in[0,b_{max}]}{\arg\min} \mathbf{H}\left(x^{*}, v^{*}, \lambda_{x}^{*}, \lambda_{v}^{*}, b(t)\right)$$
(6.79)

If we leave the final time t_f as a free parameter, we need to impose one extra constraint on the final state:

$$\mathbf{H}\left(x^{*}(t_{f}), v^{*}(t_{f}), \lambda_{x}^{*}(t_{f}), \lambda_{v}^{*}(t_{f}), b^{*}(t_{f})\right) = 0.$$
(6.80)

Together with the initial condition in eqn (6.27) and final condition in eqn (6.28), the optimal control problem becomes two-point boundary value problem. For the three cases we considered in section 6.4, L = 1 minimize the time, L = b(t) minimizes the air volume and L = H minimize the energy respectively. It's not hard to observe that for all the cases the control theory Hamitonian **H** is linear with the control b, so eqn (6.79) can be simplified based on the sign of $\frac{\partial \mathbf{H}}{\partial b}$:

$$b^{*}(t) = \begin{cases} 0 & \frac{\partial \mathbf{H}}{\partial b} > 0\\ b_{\max} & \frac{\partial \mathbf{H}}{\partial b} < 0 \end{cases}$$
(6.81)

This is the bang-bang schedule as calculated previously. The special case $\frac{\partial \mathbf{H}}{\partial b} = 0$ will be dealt with in the next section. If $\frac{\partial \mathbf{H}}{\partial b}$ is non-zero almost everywhere, the optimal control **b** will be switching between 0 and b_{max} , thus a bang-bang control. After solving the two-point

boundary value problem for L = 1 and L = b(t), the optimal bang-bang control is identical to section 6.4.

6.5.1 Singular control

For the minimum energy case L = H, bang-bang control may not provide a solution to the two-point boundary value problem for the reason that $\frac{\partial \mathbf{H}}{\partial b} = 0$ will hold for at least some time interval in which case the Pontryagin minimum principle fails to yield the complete solution. For this, we need to implement what is called singular control [21].

The most simple method to solve this problem is to repeatly differentiate $\frac{\partial \mathbf{H}}{\partial b}$ and set it to zero [21, 133]. First we know that

$$\frac{\partial \mathbf{H}}{\partial b} = 0 \tag{6.82}$$

for some time interval, so set the first order and second order time derivatives of $\frac{\partial \mathbf{H}}{\partial b}$ to 0 as well:

$$\frac{d}{dt} \left(\frac{\partial \mathbf{H}}{\partial b} \right) = 0 \tag{6.83}$$

$$\frac{d^2}{dt^2} \left(\frac{\partial \mathbf{H}}{\partial b} \right) = 0. \tag{6.84}$$

Since the canonical equations (6.78) give the formulas of $\dot{x}^*, \dot{v}^*, \dot{\lambda}^*_x, \dot{\lambda}^*_v$, the time derivatives can be easily calculated with the chain rule.

After solving the equations (6.82), (6.83) and (6.84), we have the following:

$$\lambda^*_v(t) = -\frac{1}{\alpha} \tag{6.85}$$

$$\lambda_{x}^{*}(t) = -v^{*}(t) \tag{6.86}$$

$$b^*(t) = 2e^{\alpha x^*(t)}. \tag{6.87}$$

Now we achieve the formula b(t). We will show that eqn (6.61) in section 6.4 (the transfer segment) is actually a singular arc. Take the time derivative of equation (6.61):

$$\frac{d}{dt}v(x) = \frac{d}{dt}\sqrt{2gx + 2g/\alpha}$$

$$= \frac{2g}{2\sqrt{2gx + 2g/\alpha}}\frac{dx}{dt}$$

$$= \frac{g}{\sqrt{2gx + 2g/\alpha}}v$$

$$= g$$
(6.88)

Together with eqn (6.1), the control b(t) for this segment will be:

$$b(t) = (\dot{v} + g)e^{\alpha x}/g$$

$$= 2e^{\alpha x}$$
(6.89)

which is identical to the singular arc in eqn (6.87) (shown as the transfer segment in figure 6.7(a)).

6.6 Discussion

We have shown, using basic classical mechanics techniques along with orbit transfer ideas [60, 80], that in almost all parameter cases, it is a bang-bang control schedule that achieves the transfer of the ball to the levitation point x^* in minimum time, with minimum air volume, and expending minimum energy. There are, however, some parameter values for which a minimum energy transfer is only achieved with the use of an extra transfer segment that connects two bang-bang arcs. The transfer formula is obtained analytically and is an example of the use of singular control methods [21]. It might be of some interest to formulate similar optimal control problems associated with other forms of levitation, such as acoustic levitation or magnetic levitation.

It is worth noting that aside from choosing an appropriate time-dependent function b(t) to levitate the ball to the elliptic point, there are other mechanisms we could exploit. The simplest would be to add a small amount of dissipation to the model, proportional to the velocity:

$$m\ddot{x} = -mg + mgb(t)\exp(-\alpha x) - \gamma \dot{x}.$$
(6.90)

Here γ is our dissipation parameter. We show in figure 6.8 a trajectory that ends up at x^* (as $t \to \infty$) after oscillating around it more and more tightly.



Figure 6.2: (a) Phase curves in the (x, v) plane for 5 different values of the energy E_b . x_{max} marks the maximum height that the ball can achieve on the energy curve $E_b = E_b^0$, while the elliptic point x^* marks the levitation point where the forces balance $E_b = E_b^*$. (b) Energy curves $E_b = E_b^0$ going through the origin for four different values of b. (c) Phase curves of free fall (b = 0) with different values of energy. (d) x_{max} marks the intersection of the two curves defined by the left and right hand sides of eqn (6.8).



Figure 6.3: Diagram showing two different paths from point O to point x^* . Orbits corresponding to energy $E_b = E_b^0$ for two different values, $b = b_1$ (blue) and $b = b_2 < b_1$ (red). Height x^* represents the elliptic fixed point associated with $b = b_1$. The value of b_2 is chosen so that x_{max} for the red orbit intersects x^* . The pink orbit corresponds to a free-fall trajectory b = 0 with an initial upward launch velocity v^* so that the trajectory with the blue orbit.



Figure 6.4: Examples of piecewise constant blowing schedules that move the ball from the origin to the target point $(x, v) = (x^*, 0)$. Parameter values are given by: $b \in [0, 6]$, $\alpha = 1m^{-1}, g = 9.8m/s^2, x^* = log(2)meter, m = 1kg$, and we choose T = 1.34s for time scale. (a) Two-step schedule: b(t) starts with $b = \frac{\alpha x^* e^{\alpha x^*}}{e^{\alpha x^*} - 1}$, then it switches to $b = e^{\alpha x^*}$. (b) The corresponding trajectories for the two-step schedule. Note that the ball will remain at equilibrium on second step, so the trajectory only has one segment. (c) Three-step schedule with blowing on the first segment, followed by free falling (b = 0) for the second segment. (d) The corresponding trajectory for (c). The trajectories will get closer and closer to vertical axis with the increase of initial b(t). (e) Examples of general three-step schedule, with b_1 for first segment and b_2 for the second. (f) The corresponding b(t) for the examples of general three step schedule. (g) A schematic of a four-step schedule, where $t_1 t_2$ and t_3 mark three switching points.(h) The corresponding trajectory for (g)



Figure 6.5: The schematic of the trajectory driven by a general blowing function b(t) as indicated in blue. For comparison, we show what we call the *bang-bang* strategy, starting with maximum blowing b_{max} from t = 0 to $t = t_s$, then free fall b = 0 from $t = t_s$ to $t = t_f$. Any general trajectory associated with b(t) is trapped between the two segments of the other orbit.



Figure 6.6: The bang-bang (on-off) strategy is optimal to attain the transfer in minimum time and using minimum air volume. (a) The actual trajectory for the transfer; (b) The time-dependent blowing schedule.



Figure 6.7: Minimum energy strategy.(a)phase space. Simulated with parameters $b_{max} = 20, \alpha = 1m^{-1}, x^* = ln(5)m, g = 9.8m/s^2, m = 1kg$ (b)corresponding schedule for (a) .(c)phase space simulated with parameters: $b_{max} = 5, \alpha = 1m^{-1}, x^* = ln(5)m, g = 9.8m/s^2, m = 1kg$.(d) corresponding optimal control and it's still bang-bang control.



Figure 6.8: Constant levitation with dissipation. The numbers for the simulation are $b = 5, \alpha = 1m^{-1}, g = 9.8m/s^2, x^* = log(5)m, m = 1kg, \gamma = 1kg/s$

Reference List

- CA Aktipis, VSY Kwan, KA Johnson, SL Neuberg, and CC Maley. Overlooking evolution: a systematic analysis of cancer relapse and therapeutic resistance research. *PloS One*, 6(11):e26100, 2011.
- [2] Eitan Altman, Rachid ElAzouzi, Yezekeael Hayel, and Hamidou Tembine. An evolutionary game approach for the design of congestion control protocols in wireless networks. In 2008 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, pages 547–552. IEEE, 2008.
- [3] D Andersson, NQ Balaban, F Baquero, P Courvalin, P Glaser, U Gophna, R Kishony, S Molin, and T Tonjum. Antibiotic resistance: turning evolutionary principles into clinical reality. *FEMS Microbiology Rev.*, 2020.
- [4] Dan I Andersson and Bruce R Levin. The biological cost of antibiotic resistance. *Current opinion in microbiology*, 2(5):489–493, 1999.
- [5] DI Andersson and BR Levin. The biological cost of antibiotic resistance. Current Opinion in Microbiology, 2:489–493, 1999.
- [6] Camille Stephan-Otto Attolini and Franziska Michor. Evolutionary theory of cancer. Annals of the New York Academy of Sciences, 1168(1):23–51, 2009.
- [7] Robert Axelrod, David E Axelrod, and Kenneth J Pienta. Evolution of cooperation among tumor cells. Proc. Natl. Acad. Sci., 103(36):13474–13479, 2006.
- [8] Katarina Bacevic, Robert Noble, Ahmed Soffar, Orchid Wael Ammar, Benjamin Boszonyik, Susana Prieto, Charles Vincent, Michael E Hochberg, Liliana Krasinska, and Daniel Fisher. Spatial competition constrains resistance to targeted cancer therapy. *Nature communications*, 8(1):1995, 2017.
- [9] D Basanta and ARA Anderson. Exploiting ecological principles to better understand cancer progression and treatment. *Interface Focus*, 3(4):20130020, 2013.
- [10] D Basanta, RA Gatenby, and ARA Anderson. Exploiting evolution to treat drug resistance: combination therapy and the double bind. *Molecular pharmaceutics*, 9(4):914– 921, 2012.
- [11] M Baym, LK Stone, and R Kishony. Multidrug evolutionary strategies to reverse antibiotic resistance. *Science*, 351(6268):DOI: 10.1126/science.aad3292, 2016.

- [12] RA Beckman, GS Schemmarm, and CH Yeang. Impact of genetic dynamics and singlecell heterogeneity on the development of personalized non-standard medicine strategies for cancer. *Proc. Natl. Acad. Sci.*, 109(36):14586–14591, 2012.
- [13] Sharon Bewick, Ruoting Yang, and Mingjun Zhang. Embedding evolutionary game theory into an optimal control framework for drug dosage design. In 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pages 6026–6029. IEEE, 2009.
- [14] CI Bliss. The toxicity of poisons applied jointly. Annals of Applied Biology, 26:585–615, 1939.
- [15] K Bohl, S Hummert, S Werner, D Basanta, A Deutsch, S Schuster, G Theiben, and A Schroeter. Evolutionary game theory: molecules as players. *Molecular BioSystems*, 10(12):3044–3065, 2014.
- [16] I Bozic and MA Nowak. Timing and heterogeneity of mutations associated with drug resistance in metastatic cancers. Proc. Natl. Acad. Sci., 111(45):15964–15968, 2014.
- [17] I Bozic and MA Nowak. Resisting resistance. Ann. Rev. of Can. Bio., 1:203–221, 2017.
- [18] K Bratton, U Ledzewicz, and H Schattler. Modeling and control of heterogeneous tumors under chemotherapy. *Biomath Comm.*, 1(1):1–2, 2014.
- [19] JS Brown and K Stankova. Game theory as a conceptual framework for managing insect pests. Current Opinion in Insect Sci., 21:26–32, 2017.
- [20] William L Brown and Edward O Wilson. Character displacement. Systematic Zoology, 5(2):49–64, 1956.
- [21] Arthur Earl Bryson. Applied Optimal Control: Optimization, Estimation and Control. Routledge, 2018.
- [22] Cécile Carrere. Optimization of an in vitro chemotherapy to avoid resistant tumours. Journal of Theoretical Biology, 413:24–33, 2017.
- [23] Gustavo F Carvalhal, Saima N Daudi, Donghui Kan, Dana Mondo, Kimberly A Roehl, Stacy Loeb, and William J Catalona. Correlation between serum prostate-specific antigen and cancer volume in prostate glands of different sizes. Urology, 76(5):1072– 1076, 2010.
- [24] J Chacon, J Saenz, L de la Torre, JM Diaz, and F Esquembre. Design of a low-cost air levitation system for teaching control engineering. *Sensors*, 17(2321), 2017.
- [25] TC Chou and DC Rideout. Synergism and Antagonism in Chemotherapy. Academic Press, NY, 1991.
- [26] Andrew J Coldman and JM Murray. Optimal control for a stochastic model of cancer chemotherapy. *Mathematical biosciences*, 168(2):187–200, 2000.

- [27] Joseph H Connell. The influence of interspecific competition and other factors on the distribution of the barnacle chthamalus stellatus. *Ecology*, 42(4):710–723, 1961.
- [28] Charles Darwin and William F Bynum. The origin of species by means of natural selection: or, the preservation of favored races in the struggle for life. Penguin Harmondsworth, 2009.
- [29] Luis A Diaz Jr, Richard T Williams, Jian Wu, Isaac Kinde, J Randolph Hecht, Jordan Berlin, Benjamin Allen, Ivana Bozic, Johannes G Reiter, Martin A Nowak, et al. The molecular evolution of acquired resistance to targeted egfr blockade in colorectal cancers. *Nature*, 486(7404):537–540, 2012.
- [30] Li Ding, Timothy J Ley, David E Larson, Christopher A Miller, Daniel C Koboldt, John S Welch, Julie K Ritchey, Margaret A Young, Tamara Lamprecht, Michael D McLellan, et al. Clonal evolution in relapsed acute myeloid leukaemia revealed by whole-genome sequencing. *Nature*, 481(7382):506–510, 2012.
- [31] Pedro M Enriquez-Navas, Yoonseok Kam, Tuhin Das, Sabrina Hassan, Ariosto Silva, Parastou Foroutan, Epifanio Ruiz, Gary Martinez, Susan Minton, Robert J Gillies, et al. Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Science Translational Medicine*, 8(327):327ra24–327ra24, 2016.
- [32] PM Enriquez-Navas, JW Wojtkowiak, and RA Gatenby. Application of evolutionary principles to cancer therapy. *Cancer research*, 75(22):4675–4680, 2015.
- [33] A Fisher, I Vazquez-Garcia, and V Mustonen. The value of monitoring to control evolving populations. *Proc. Natl. Acad. Sci.*, 112(4):1007–1012, 2015.
- [34] J Foo and Michor F. Evolution of resistance to anti-cancer therapy during general dosing schedules. J. Theor. Bio., 263(2):179–188, 2010.
- [35] J Foo and Michor F. Evolution of acquired resistance to anti-cancer therapy. J. Theor. Bio., 355:10–20, 2014.
- [36] Jill Gallaher, Pedro M Enriquez-Navas, Kimberly A Luddy, Robert A Gatenby, and Alexander RA Anderson. Spatial heterogeneity and evolutionary dynamics modulate time to recurrence in continuous and adaptive cancer therapies. *Cancer Research*, 2018.
- [37] Jill A Gallaher, Pedro M Enriquez-Navas, Kimberly A Luddy, Robert A Gatenby, and Alexander RA Anderson. Adaptive vs continuous cancer therapy: Exploiting space and trade-offs in drug scheduling. 2017.
- [38] RA Gatenby. A change of strategy in the war on cancer. *Nature*, 459(7246):508–509, 2009.
- [39] RA Gatenby and J Brown. The evolution and ecology of resistance in cancer therapy. Cold Spring Harbor Perspectives in Medicine, page doi: 10.1101/cshperspect.a033415, 2017.

- [40] RA Gatenby, AS Silva, RJ Gillies, and BR Frieden. Adaptive therapy. Cancer Research, 69(11):4894–4903, 2009.
- [41] Robert A Gatenby, Robert J Gillies, and Joel S Brown. Of cancer and cave fish. Nature Reviews Cancer, 11(4):237, 2011.
- [42] RJ Gillies, D Verduzco, and RA Gatenby. Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. *Nature Reviews Cancer*, 12(7):487–493, 2012.
- [43] G Gomez, WS Koon, MW Lo, JE Marsden, and SD Ross. Connecting orbits and invariant manifolds in the spatial restricted three-body problem. *Nonlinearity*, 17(5):1571, 2004.
- [44] F Gould. The evolutionary potential of crop pests. *American Scientist*, 79:496–507, 1991.
- [45] Peter R Grant. Convergent and divergent character displacement. *Biological journal* of the Linnean Society, 4(1):39–68, 1972.
- [46] Mel Greaves and Carlo C Maley. Clonal evolution in cancer. Nature, 481(7381):306– 313, 2012.
- [47] P Grover and SD Ross. Designing trajectories in a planet-moon environment using the controlled keplerian map. J. of Guidance, Control, Dyn., 32(2):436–443, 2009.
- [48] Erik Gullberg, Sha Cao, Otto G Berg, Carolina Ilbäck, Linus Sandegren, Diarmaid Hughes, and Dan I Andersson. Selection of resistant bacteria at very low antibiotic concentrations. *PLoS pathogens*, 7(7):e1002158, 2011.
- [49] Elsa Hansen, Robert J Woods, and Andrew F Read. How to use a chemotherapeutic agent when resistance to it threatens the patient. *PLoS biology*, 15(2):e2001110, 2017.
- [50] M Hegreness, N Shoresh, D Damiann, D Hartl, and R Kishony. Accelerated evolution of resistance in multidrug environments. *Proc. Nat'l Acad. Sci.*, 105(37):13977–13981, 2008.
- [51] Keith S Hoek and Colin R Goding. Cancer stem cells versus phenotype-switching in melanoma. Pigment cell & melanoma research, 23(6):746–759, 2010.
- [52] Josef Hofbauer and Karl Sigmund. Evolutionary Games and Population Dynamics. Cambridge University Press, 1998.
- [53] S Hummert, K Bohl, D Basanta, A Deutsch, S Werner, G Theiben, A Schroeter, and S Schuster. Evolutionary game theory: cells as players. *Molecular BioSystems*, 10(12):3044–3065, 2014.
- [54] G Jamsen, R Gatenby, and CA Aktipis. Opinion: Control vs. eradication: Applying infectious disease strategies to cancer. Proc. Natl. Acad. Sci., 112(4):937–938, 2015.

- [55] Hem Raj Joshi. Optimal control of an hiv immunology model. Optimal control applications and methods, 23(4):199–213, 2002.
- [56] Kristel Kemper, Oscar Krijgsman, Paulien Cornelissen-Steijger, Aida Shahrabi, Fleur Weeber, Ji-Ying Song, Thomas Kuilman, Daniel J Vis, Lodewyk F Wessels, Emile E Voest, et al. Intra-and inter-tumor heterogeneity in a vemurafenib-resistant melanoma patient and derived xenografts. *EMBO Molecular Medicine*, 7(9):1104–1118, 2015.
- [57] S Kim, TD Lieberman, and R Kishony. Alternating antibiotic treatments constrain evolutionary paths to multidrug resistance. *Proc. Nat'l Acad. Sci.*, 40:14494–14499, 2014.
- [58] NL Komarova and D Wodarz. Drug resistance in cancer: Principles of emergence and prevention. Proc. Natl. Acad. Sci., 102(27):9714–9719, 2005.
- [59] S Konishi, M Harada, Y Ogami, Y Daiho, Y Mita, and H Fujita. Measurements and analysis on characteristics of a conveyance system using air levitation. *IEEE 6th International Conference on Emerging Technologies and Factory Automation*, page DOI: 10.1109/ETFA.1997.616274, 1997.
- [60] WS Koon, MW Lo, JE Marsden, and SD Ross. Heteroclinic connections between periodic orbits and resonance transitions in celestial mechanics. *Chaos*, 10(427):doi:10.1063/1.166509, 2000.
- [61] K-A Kreuzer, P Le Coutre, O Landt, I-K Na, M Schwarz, K Schultheis, A Hochhaus, and B Dörken. Preexistence and evolution of imatinib mesylate-resistant clones in chronic myelogenous leukemia detected by a pna-based pcr clamping technique. Annals of Hematology, 82(5):284–289, 2003.
- [62] Pierre Laurent-Puig, Deniz Pekin, Corinne Normand, Steve K Kotsopoulos, Philippe Nizard, Karla Perez-Toralla, Rachel Rowell, Jeff Olson, Preethi Srinivasan, Delphine Le Corre, et al. Clinical relevance of kras-mutated subclones detected with picodroplet digital pcr in advanced colorectal cancer treated with anti-egfr therapy. *Clinical Cancer Research*, 2014.
- [63] O Lavi, MM Gottesman, and D Levy. The dynamics of drug resistance: A mathematical perspective. Drug Res. Updates, 15(1-2):90–97, 2012.
- [64] U Ledzewicz and HM Schattler. Optimal bang-bang controls for a two-compartmental model in cancer therapy. J. Optimization Theory and Appl., 114(3):609–637, 2002.
- [65] U Ledzewicz, S Wang, H Schattler, N Andre, MA Heng, and E Pasquier. On drug resistance and metronomic chemotherapy: A mathematical modeling and optimal control approach. *Math. Biosciences and Eng.*, 14(1):217–235, 2017.
- [66] Ernest Bruce Lee and Lawrence Markus. Foundations of optimal control theory. Technical report, Minnesota Univ Minneapolis Center For Control Sciences, 1967.
- [67] S Lev. Pontryagin. mathematical theory of optimal processes, 1987.

- [68] RC Lewontin. The units of selection. Annual Reviews of Ecol. Syst., 1:1–18, 1970.
- [69] Erez Lieberman, Christoph Hauert, and Martin A Nowak. Evolutionary dynamics on graphs. Nature, 433(7023):312–316, 2005.
- [70] M Lipsitch and BR Levin. The population dynamics of antimicrobial chemotherapy. Antimicrobial Agents and Chemotherapy, 41(2):363–373, 1997.
- [71] S Loewe. The problem of synergism and antagonism of combined drugs. Arzneim. Forsch., 3:285–290, 1953.
- [72] Alfred J Lotka. Contribution to the theory of periodic reactions. The Journal of Physical Chemistry, 14(3):271–274, 2002.
- [73] Rory B Martin. Optimal control drug scheduling of cancer chemotherapy. Automatica, 28(6):1113–1123, 1992.
- [74] Lauren MF Merlo, John W Pepper, Brian J Reid, and Carlo C Maley. Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*, 6(12):924–935, 2006.
- [75] M Dror Michaelson, Stephane Oudard, Yen-Chuan Ou, Lisa Sengeløv, Fred Saad, Nadine Houede, Peter Ostler, Arnulf Stenzl, Gedske Daugaard, Robert Jones, et al. Randomized, placebo-controlled, phase iii trial of sunitinib plus prednisone versus prednisone alone in progressive, metastatic, castration-resistant prostate cancer. *Journal* of Clinical Oncology, 32(2):76–82, 2013.
- [76] JB Michel, P Yeh, R Chait, RC Moellering, and R Kishony. Drug interactions modulate the potential for evolution of resistance. *Proc. Nat'l Acad. Sci.*, 105(39):14918–14923, 2008.
- [77] A Sorana Morrissy, Livia Garzia, David JH Shih, Scott Zuyderduyn, Xi Huang, Patryk Skowron, Marc Remke, Florence MG Cavalli, Vijay Ramaswamy, Patricia E Lindsay, et al. Divergent clonal selection dominates medulloblastoma at recurrence. *Nature*, 529(7586):351–357, 2016.
- [78] PK Newton and Y Ma. Nonlinear adaptive control of competitive release and chemotherapeutic resistance. *Phys. Rev. E*, 99(022404), 2019.
- [79] PK Newton and Y Ma. Nonlinear adaptive control of competitive release and chemotherapeutic resistance. *Physical Review E*, 99(2):022404, 2019.
- [80] PK Newton and Y Ma. Nonlinear adaptive control of competitive release and chemotherapeutic resistance. *Phys. Rev. E*, 99(022404), 2019.
- [81] D Nichol, P Jeavons, A G Fletcher, R A Bonomo, P K Maini, J L Paul, R A Gatenby, A R A Anderson, and J G Scott. Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comput Biol*, 11(9):e1004493, 2015.
- [82] D Nichol, P Jeavons, AG Fletcher, RA Bonomo, PK Maini, JL Paul, RA Gatenby, ARA Anderson, and JG Scott. Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comp. Bio.*, 11(9), 2015.
- [83] L Norton. A gompertzian model of human breast cancer growth. Can. Res., 48(24 Part 1):7067–7071, 1988.
- [84] L Norton and R Simon. Tumor size, sensitivity to therapy, and design of treatment schedules. *Cancer Treatment Reports*, 61(7):1307, 1977.
- [85] Martin A Nowak. Evolutionary Dynamics. Harvard University Press, 2006.
- [86] Peter C Nowell. The clonal evolution of tumor cell populations. Science, 194(4260):23–28, 1976.
- [87] Martin J Osborne and Ariel Rubinstein. A course in game theory. MIT press, 1994.
- [88] MC Perry. The Chemotherapy Source Book. Lippincott Williams & Wilkins, 2008.
- [89] Daniel P Petrylak, Nicholas J Vogelzang, Nikolay Budnik, Pawel Jan Wiechno, Cora N Sternberg, Kevin Doner, Joaquim Bellmunt, John M Burke, Maria Ochoa de Olza, Ananya Choudhury, et al. Docetaxel and prednisone with or without lenalidomide in chemotherapy-naive patients with metastatic castration-resistant prostate cancer (mainsail): a randomised, double-blind, placebo-controlled phase 3 trial. The Lancet Oncology, 16(4):417–425, 2015.
- [90] Angela Oliveira Pisco, Amy Brock, Joseph Zhou, Andreas Moor, Mitra Mojtahedi, Dean Jackson, and Sui Huang. Non-darwinian dynamics in therapy-induced cancer drug resistance. *Nature communications*, 4:2467, 2013.
- [91] Lev Semenovich Pontryagin. *Mathematical Theory of Optimal Processes*. Routledge, 2018.
- [92] JE Prussing and BA Conway. Orbital Mechanics, 2nd Ed. Oxford University Press, 2013.
- [93] S Radu, WS Koon, MW Lo, JE Marsden, SD Ross, and RS Wilson. Halo orbit mission correction maneuvers using optimal control. Automatica, 38:571–583, 2002.
- [94] AF Read, T Day, and S Huijben. The evolution of drug resistance and the curious orthodoxy of aggressive chemotherapy. Proc. Natl. Acad. Sci., 108(2):10871–10877, 2011.
- [95] Nicholas P Restifo, Mark E Dudley, and Steven A Rosenberg. Adoptive immunotherapy for cancer: harnessing the t cell response. *Nature Reviews Immunology*, 12:269–281, 2012.
- [96] Catherine Roche-Lestienne and Claude Preudhomme. Mutations in the abl kinase domain pre-exist the onset of imatinib treatment. In *Seminars in Hematology*, volume 40, pages 80–82. Elsevier, 2003.

- [97] Alessandro Romanel, Delila Gasi Tandefelt, Vincenza Conteduca, Anuradha Jayaram, Nicola Casiraghi, Daniel Wetterskog, Samanta Salvi, Dino Amadori, Zafeiris Zafeiriou, Pasquale Rescigno, et al. Plasma ar and abiraterone-resistant prostate cancer. *Science Translational Medicine*, 7(312):312re10–312re10, 2015.
- [98] IM Ross. A Primer on Pontryagin's Principle in Optimal Control, 2nd Ed. Collegiate Press, 2015.
- [99] SD Ross and DJ Scheeres. Multiple gravity assists, capture, and escape in the restricted three-body problem. *SIAM J. Appl. Dyn. Sys.*, 6(3):576–596, 2007.
- [100] D Russ and R Kishony. Additivity of inhibitory effects in multidrug combinations. *Nature Microbiology*, 3:1339–1345, 2018.
- [101] Bean San Goh, George Leitmann, and Thomas L Vincent. Optimal control of a preypredator system. *Mathematical Biosciences*, 19(3-4):263–286, 1974.
- [102] WH Sandholm, E Dokumacı, and F Franchetti. Dynamo: diagrams for evolutionary game dynamics. software, 2012.
- [103] Roland F Schwarz, Charlotte KY Ng, Susanna L Cooke, Scott Newman, Jillian Temple, Anna M Piskorz, Davina Gale, Karen Sayal, Muhammed Murtaza, Peter J Baldwin, et al. Spatial and temporal heterogeneity in high-grade serous ovarian cancer: A phylogenetic analysis. *PLoS Med*, 12(2):e1001789, 2015.
- [104] H Segre, N DeMalach, Z Henkin, and R Kadmon. Quantifying competitive exclusion and competitive release in ecological communities: A conceptual framework and a case study. *PloS One*, 11(8):e0160798, 2016.
- [105] S Seton-Rogers. Chemotherapy: Preventing competitive release. Nature Rev. Can., 16(4):199–199, 2016.
- [106] Sreenath V Sharma, Diana Y Lee, Bihua Li, Margaret P Quinlan, Fumiyuki Takahashi, Shyamala Maheswaran, Ultan McDermott, Nancy Azizian, Lee Zou, Michael A Fischbach, et al. A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1):69–80, 2010.
- [107] AS Silva and RA Gatenby. A theoretical quantitative model for evolution of cancer chemotherapy resistance. *Biology Direct*, 5(25):1–25, 2010.
- [108] J Maynard Smith. Game theory and the evolution of fighting. On evolution, pages 8–28, 1972.
- [109] SH Strogatz. Nonlinear Dynamics and Chaos. Westview Press, 2015, 2nd Ed.
- [110] P Sudalagunta, MC Silva, RR Canevarolo, R Gatenby, Gillies R, R Baz, MB Meads, KH Shain, and AS Silva. A pharmacodynamic model of clinical synergy in multiple myeloma. *EBioMedicine*, 54:102716, 2020.

- [111] George W Swan. Role of optimal control theory in cancer chemotherapy. Mathematical biosciences, 101(2):237–284, 1990.
- [112] Ian F Tannock, Ronald de Wit, William R Berry, Jozsef Horti, Anna Pluzanska, Kim N Chi, Stephane Oudard, Christine Théodore, Nicholas D James, Ingela Turesson, et al. Docetaxel plus prednisone or mitoxantrone plus prednisone for advanced prostate cancer. New England Journal of Medicine, 351(15):1502–1512, 2004.
- [113] Kok Lay Teo, C Goh, and K Wong. A unified computational approach to optimal control problems. 1991.
- [114] Arne Traulsen, Jens Christian Claussen, and Christoph Hauert. Coevolutionary dynamics: From finite to infinite populations. *Physical Review Letters*, 95(23):238701, 2005.
- [115] Arne Traulsen, Jens Christian Claussen, and Christoph Hauert. Coevolutionary dynamics in large, but finite populations. *Physical Review E*, 74(1):011901, 2006.
- [116] S Venkatesan and C Swanton. Tumor evolutionary principles: How intratumor heterogeneity influences cancer treatment and outcome. In American Society of Clinical Oncology educational book. American Society of Clinical Oncology. Meeting, volume 35, page e141, 2016.
- [117] B Walsh and MW Blows. Abundant genetic variation + strong selection = multivariate genetic constraints: A geometric view of adaptation. Annu. Rev. Ecol. Evol. Syst., 40:41–59, 2009.
- [118] J West, Z Hasnain, P Macklin, and PK Newton. An evolutionary model of tumor cell kinetics and the emergence of molecular heterogeneity and gompertzian growth. SIAM Review, 58(4):716–736, 2016.
- [119] J West, Z Hasnain, J Mason, and PK Newton. The prisoner's dilemma as a cancer model. Converg. Sci. Phys. Oncol., 2(3):035002, 2016.
- [120] J West, Y Ma, and PK Newton. Capitalizing on competition: An evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment. J. Theor. Bio., 455:249–260, 2018.
- [121] J West, Y Ma, and PK Newton. Towards multi-drug adaptive therapy. Cancer Research, 80(7):1578–1589, 2020.
- [122] J West and PK Newton. Chemotherapeutic dose scheduling based on tumor growth rates provides a case for low-dose metronomic high-entropy therapies. *Cancer Research*, pages doi: 10.1158/0008-5472.CAN-17-1120, 2017.
- [123] J West and PK Newton. Cellular interactions constrain tumor growth. Proc. Nat'l Acad. Sci., 116(6):1918–1923, 2019.

- [124] Jeffrey West, Zaki Hasnain, Paul Macklin, and Paul K Newton. An evolutionary model of tumor cell kinetics and the emergence of molecular heterogeneity driving gompertzian growth. SIAM Review, 58(4):716–736, 2016.
- [125] Jeffrey West, Yongqian Ma, and Paul K Newton. Capitalizing on competition: An evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment. *Journal of theoretical biology*, 455:249–260, 2018.
- [126] Jeffrey West and Paul K Newton. Chemotherapeutic dose scheduling based on tumor growth rates provides a case for low-dose metronomic high-entropy therapies. *Cancer research*, 77(23):6717–6728, 2017.
- [127] Julia Wilkerson, Kald Abdallah, Charles Hugh-Jones, Greg Curt, Mace Rothenberg, Ronit Simantov, Martin Murphy, Joseph Morrell, Joel Beetsch, Daniel J Sargent, et al. Estimation of tumour regression and growth rates during treatment in patients with advanced prostate cancer: a retrospective analysis. *The Lancet Oncology*, 18(1):143– 154, 2017.
- [128] KB Wood, KC Wood, S Nishida, and P Cluzel. Uncovering scaling laws to infer multidrug response of resistant microbes and cancer cells. *Cell Reports*, 6:1073–1084, 2014.
- [129] B Wu, PM Altrock, L Wang, and A Traulsen. Universality of weak selection. Phys. Rev. E, 82:046106, 2010.
- [130] N Yoon, R Vander Velde, A Marusyk, and J G Scott. Optimal therapy scheduling based on a pair of collaterally sensitive drugs. *Bull. Math. Bio.*, pages 1–34, 2018.
- [131] N Yoon, R Vander Velde, A Marusyk, and JG Scott. Optimal therapy scheduling based on a pair of collaterally sensitive drugs. *Bull. Math. Bio.*, 80(7):1776–1809, 2018.
- [132] L You, JS Brown, F Thuijsman, JJ Cunningham, RA Gatenby, J Zhang, and K Stankova. Spatial vs. non-spatial eco-evolutionary dynamics in a tumor growth model. J. Theor. Bio., 435:78–97, 2017.
- [133] MI Zelikin and VF Borisov. Singular optimal regimes in problems of mathematical economics. *Journal of Mathematical Sciences*, 130(1):4409–4570, 2005.
- [134] J Zhang, JJ Cunningham, JS Brown, and RA Gatenby. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature Comm.*, 8(1):1816, 2017.

Appendix A

replicator dynamics of tumor growth

To understand why the prisoner's dilemma is a useful paradigm for tumor growth resulting from competing subpopulations, we focus on the 2×2 case only. The standard version of the prisoner's dilemma payoff matrix [85] in a 2×2 setting in which healthy cells compete with cancer cells is:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 3 & 0 \\ 5 & 1 \end{bmatrix},$$
 (A.1)

where the first row and column correspond to the payoffs associated with the *cooperator* (C) in the PD evolutionary game, and the second row and column correspond to the payoffs associated with the *defector* (D). In the simplest tumor growth paradigm in which a population of healthy cells competes with a population of cancer cells, the healthy cells are the cooperators, while the cancer cells are the defectors. In any mixed population $\vec{x} = (x_C, x_D)^T$, $0 \le x_C \le 1$; $0 \le x_D \le 1$; $x_C + x_D = 1$, the fitness functions, $\vec{f} = (f_C, f_D)^T$, associated with the two subpopulations are:

$$\vec{f} = A\vec{x}, \tag{A.2}$$

which in component form yields:

$$f_C = (A\vec{x})_1 = 3 \cdot x_C + 0 \cdot x_D, \tag{A.3}$$

$$f_D = (A\vec{x})_2 = 5 \cdot x_C + 1 \cdot x_D,$$
 (A.4)

while the average fitness of the total population is given by the quadratic form:

$$\langle f \rangle = \vec{x}^T A \vec{x} = 3x_C^2 + 5x_C x_D + x_D^2 \ge 1.$$
 (A.5)

The average fitness of the healthy state $(x_C, x_D) = (1, 0)$ is given by $\langle f \rangle|_{(x_C=1)} = 3$, while that of the cancerous state $(x_C, x_D) = (0, 1)$ is given by $\langle f \rangle|_{(x_D=1)} = 1$, which minimizes the average fitness. Tumor growth is then modeled as a 2 × 2 evolutionary game governed by the replicator dynamical system:

$$\dot{x}_C = (f_C - \langle f \rangle) x_C, \tag{A.6}$$

$$\dot{x}_D = (f_D - \langle f \rangle) x_D. \tag{A.7}$$

It is straightforward to show:

$$\dot{x}_D = [(c-a) - (d-b)]x_D(1-x_D) \left(\frac{1}{1-\left(\frac{d-b}{c-a}\right)} - x_D\right),\tag{A.8}$$

with fixed points at $x_D = 0, 1, \frac{(c-a)}{(c-a)-(d-b)}$. From this, we can conclude (using values from eqn (A.1)) that for any initial condition containing at least one cancer cell: $0 < x_D(0) \le 1$, we have:

(i) $x_D \to 1, x_C \to 0$ as $t \to \infty$

(ii)
$$\langle f \rangle \to 1 \text{ as } t \to \infty$$
.

Condition (i) guarantees that the cancer cell population will saturate, while condition (ii) guarantees that the saturated state is sub-optimal, since $\langle f \rangle|_{(x_D=1)} < \langle f \rangle|_{(x_C=1)}$. For these

two reasons, the prisoner's dilemma evolutionary game serves as a simple paradigm for tumor growth both in finite population models, as well as replicator system (infinite population) models [118, 119, 122, 114]. Note that the 2×2 system alone it not able to account for the evolution of resistance.