

Fibroblasts and alectinib switch the evolutionary games that non-small cell lung cancer plays

Artem Kaznatcheev^{1,2}, Jeffrey Peacock³, David Basanta⁴, Andriy Marusyk⁵, and Jacob G. Scott^{2,6}

¹Department of Computer Science, University of Oxford

²Department of Translational Hematology & Oncology Research, Cleveland Clinic

³Department of Radiation Oncology, Moffitt Cancer Center

⁴Department of Integrated Mathematical Oncology, Moffitt Cancer Center

⁵Department of Cancer Imaging and Metabolism, Moffitt Cancer Center

⁶Department of Radiation Oncology, Cleveland Clinic

Heterogeneity in strategies for survival and proliferation among the cells which constitute a tumour is a driving force behind the evolution of resistance to cancer therapy. The rules mapping the tumour’s strategy distribution to the fitness of individual strategies can be represented as an evolutionary game. We develop a game assay to measure effective evolutionary games in co-cultures of alectinib-sensitive and alectinib-resistant non-small cell lung cancer. The games are not only quantitatively different between different environments, but targeted therapy and cancer associated fibroblasts qualitatively switch the type of game being played by the *in-vitro* population from Leader to Deadlock. This observation provides the first direct empirical confirmation of a central theoretical postulate of evolutionary game theory in oncology: we can treat not only the player, but also the game. Although we concentrate on measuring games played by cancer cells, the measurement methodology we develop can be used to advance the study of games in other microscopic systems by providing a quantitative description of non-cell-autonomous effects.

Tumours are heterogeneous, evolving ecosystems [1, 2], comprised of sub-populations of neoplastic cells that follow distinct strategies for survival and propagation [3]. The success of a strategy employed by any single neoplastic sub-population is dependent on the distribution of other strategies, and on various components of the tumour microenvironment, like cancer associated fibroblasts (CAFs) [4]. The EML4-ALK fusion, found in approximately 5% of non-small cell lung cancer (NSCLC) patients, leads to constitutive activation of oncogenic tyrosine kinase activity of ALK, thereby “driving” the disease. Inhibitors of tyrosine kinase activity of ALK (ALK TKI) have proven to be highly clinically efficacious, inducing tumour regression and prolonging patient survival [5, 6]. Unfortunately, virtually all of the tumours that respond to ALK TKIs eventually relapse [7] – an outcome typical of inhibitors of other oncogenic tyrosine kinases [8]. Resistance to ALK TKI, like most targeted therapies, remains a major unresolved clinical challenge. Despite significant advances in deciphering the resultant molecular mechanisms of resistance [9], the evolutionary dynamics of ALK TKI resistance remains poorly understood. The inability of TKI therapies to completely eliminate tumour cells has been shown to be at least partially attributable to protection by aspects of the tumour microenvironment [10]. CAFs are one of the main non-malignant components of tumour microenvironment and the interplay between them and tumour cells is a major contributor to microenvironmental resistance, including cytokine mediated protection against ALK inhibitors [11].

To study the eco-evolutionary dynamics of these various factors, we interrogated the competition between treatment naive cells of ALK mutant NSCLC cell line H3122 – a “workhorse”

for studies of ALK+ lung cancer – and a derivative cell line in which we developed resistance to alectinib – a highly effective clinical ALK TKI [12] – by selection in progressively increasing concentrations of the drug [13]. We aimed to come to a quantitative understanding of how these dynamics were affected by clinically relevant concentrations of alectinib (0.5 μ M; see [14]) in the presence or absence of CAFs, isolated from a lung cancer. To achieve this, we developed a novel assay for quantifying effective games that is of independent interest to the general study of microscopic systems.

Mono vs mixed cultures & cost of resistance

To establish baseline characteristics, we performed assays in monotypic cultures of parental (alectinib-sensitive) and resistant cell lines with and without alectinib and CAFs. To gather temporally-resolved data for inferring growth rates, we used time lapse microscopy to follow the expansion of therapy resistant and parental cells, differentially labeled with stable expression of selectively neutral GFP and mCherry fluorescent proteins, respectively. From the time series data, we inferred the growth rate with confidence intervals for each one of 6 experimental replicates in four different experimental conditions (total of 24 data points, each with confidence intervals), as seen in Figure 1. As expected, alectinib inhibited growth rates of parental cells (DMSO vs Alectinib: $p < .005$; DMSO + CAF vs Alectinib + CAF: $p < .005$), whereas the growth rate of the resistant cells was not affected. And, as previously reported [11], CAFs partially rescued growth inhibition of parental cells by Alectinib (Alectinib vs Alectinib + CAF: $p < .005$; Alectinib + CAF vs DMSO: $p < .005$), without




Figure 1: Monotypic culture growth rates for parental (red) and resistant (green) cells in indicated experimental conditions. Confidence intervals on each experimental replicate is from confidence on the estimate of growth rate for that single replicate according to the Theil-Sen estimator. Comparisons between experimental conditions (of 6 replicates each) are made using Wilcoxon rank-sum. In addition to conditions linked by lines with reported p-values, conditions labeled by 'a' and 'b' are pairwise distinguishable with $p < .005$.

impacting growth rates of resistant cells. alk+

The classic model of resistance posits that the resistant phenotype receives a benefit in drug (in our case: Alectinib or Alectinib + CAF) but is neutral, or even carries an inherent cost, in the absence of treatment (DMSO or DMSO + CAF). For example, experimentalists frequently regard resistance granting mutations as selectively neutral in the absence of drug, and the modeling community often goes further by considering explicit costs like up-regulating drug efflux pumps, investing in other defensive strategies, or lowering growth rate by switching to sub-optimal growth pathways [3, 15]. If we limited ourselves to these monotypic assays, then our observations would be consistent with this classic model of resistance.

But we did not limit ourselves to monotypic assays. Our experience observing non-cell-autonomous biological interactions [16] and modeling eco-evolutionary interactions [17–19] in cancer led us to suspect that the heterotypic growth rates would differ from monotypic culture. Cell-autonomous fitness effects are ones where the benefits/costs to growth rate are inherent to the cell: the presence of other cells are an irrelevant feature of the micro-environment and the growth rates from monotypic cultures provide all the necessary information. Non-cell-autonomous effects allow fitness to depend on a cell's micro-environmental context, including the frequency of other cell types: growth rates need to be measured in competitive fitness assays over a range of seeding frequencies. Other microscopic experimental systems in which frequency dependent fitness effects have been considered include, but are not limited to: *Escherichia coli* [20, 21], yeast [22, 23], bacterial symbionts of hydra [24], breast cancer [16] and pancreatic cancer [25]. Hence, we continued our experiments over a range of initial proportions of resistant and parental cells in mixed cultures for each of the four experimental conditions.

Figure 2 shows the resulting growth rates of each cell type in the co-culture experiments for all experimental (color, shape) and initial conditions (opacity is parental cell proportion). In the heterotypic culture – unlike monotypic – CAFs slightly improved the growth rates of the parental cells, even in DMSO. More strikingly, even in the absence of drug, resistant cells tend to have a higher growth rate than parental cells in the

same environment (i.e. proportion of parental cells in the co-culture). This is evident from most DMSO points being above the dotted diagonal line ($y = x$) corresponding to equal growth rate of the sub-populations (this is quantified in figure 4b and is further discussed in section ‘Leader and Deadlock games in NSCLC’). The higher fitness of resistant cells in co-culture that we observe in the absence of drug is not consistent with the classic model of resistance. This is a result that cannot be seen from just the monotypic experiments, which are compatible with the classic model. Although the highest fitness of resistant cells might not surprise clinicians as much as the biologists: in clinical experience, tumours that have acquired resistance are often more aggressive than before they were treated, even in the absence of drug.

A reductionist could rationalize our observations by saying that we actually selected for two different qualities in our resistant line: (i) a general growth advantage, and (ii) resistance to Alectinib. This is a reasonable hypothesis, but it faces a few challenges. First, both parental and resistant cells were evolved for the same length of time, with escalating dosages of DMSO for the former and Alectinib for the latter (see Mediavilla-Varela *et al.* [26] and Supplementary Appendix A). Thus, (i) cannot be due to just subculturing, but is somehow linked to drug. Second, there is no growth rate advantage of resistant cells in monoculture (see Figure 1); the advantage is only revealed when parental and resistant cells are cultured with a common proportion of parental cells. Finally, to even make the distinction between (i) and (ii), one has to implicitly assume that resistance has to be neutral or costly by definition. For an oncologist, however, both (i) and (ii) would constitute clinical resistance if they led to a tumour escaping therapeutic control. By using a definition of clinical resistance that is broad enough to capture both aspects, we observe resistance that is neither neutral nor costly in DMSO co-culture.

Frequency dependence in fitness functions

Although not common in cancer biology, competitive fitness assays are a gold standard for studying bacteria. But they are typically conducted with a single initial ratio of the two competing cell types. However, in Figure 2, if we view the

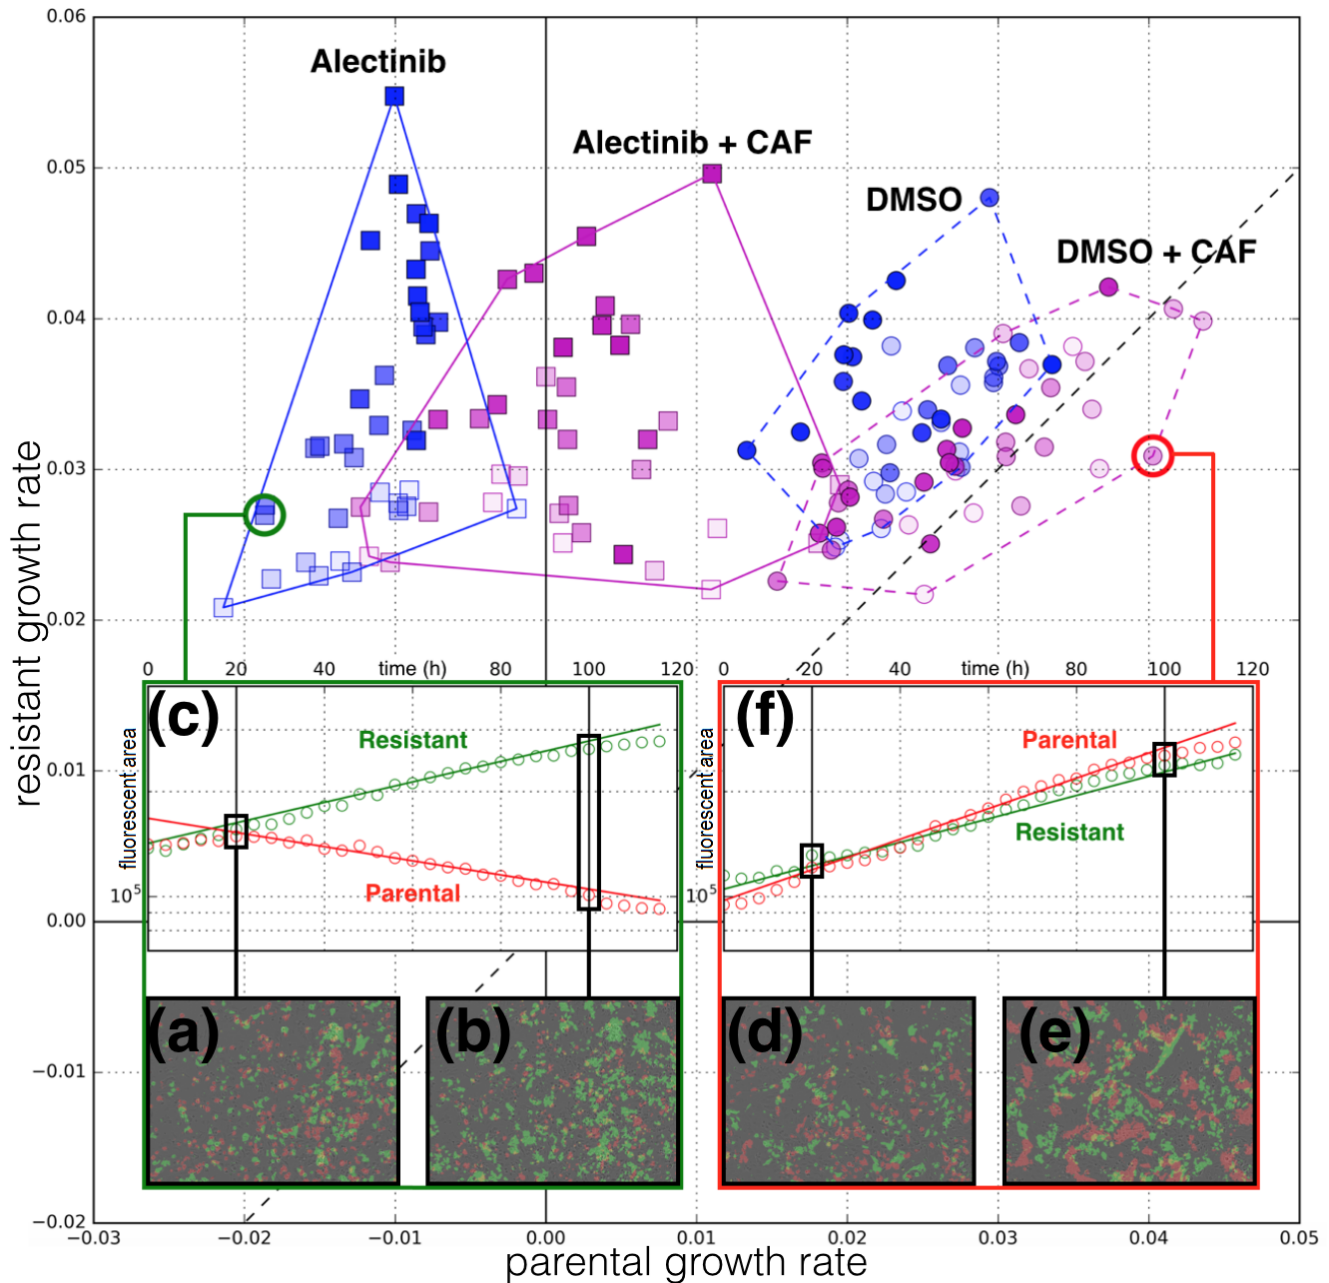


Figure 2: Coculture growth rates across four experimental conditions. Each point in the main figure is a separate replicate of a competitive fitness assay with initial proportion of parental cells represented by opacity and experimental condition represented by shape (DMSO: circle; Alectinib: square) and colour (no CAF: blue; + CAF: magenta). Each replicate's x -position corresponds to the measured parental growth rate and y -position for resistant growth rate; the dotted block line corresponds to the line of equal fitness between the two at $x = y$. (a,b,d,e) In each replicate, we quantify population size by fluorescent area of each cell type from 30 time-lapse microscopy images. (c,f): time-series of parental and resistant population size for two example wells. Inset x -axis is time, y -axis is log of population size. Median growth rate and confidence intervals (omitted) were estimated for each well using the Theil-Sen estimator, and serve as the coordinates in the main figure. See Figure 3 for growth rate confidence intervals and supplemental section B for detailed discussion.




Figure 3: Fitness functions for competition of parental vs. resistant NSCLC. For each plot: growth rate with confidence intervals versus initial proportion of parental cells. This is the same data, measured in the same way, as figure 2. Red data points are growth rates of parental cells, and green for resistant cells. Dotted lines represent the linear fitness function of the least-squares best fit; fit error visualised in Figure 4b. The black dotted line is the gain function for parental (see Figure 4a), it is well below the $y = 0$ line in the Alectinib conditions (indicating the strong advantage of resistance) and thus cut out of the figure. See supplemental section C for more discussion and equations for lines of best fit.

initial proportion of parental to resistant cells as a variable parameter represented by opacity then we can see a hint of frequency dependence in both parental and resistant growth rates. There we see an increase in fitness of both cell types as the initial proportion of parental cells (represented by the opacity of each point) increases. In other words, the variance in growth rates between points within each of the four environmental conditions is not due to just noise but can be explained by the micro-environmental variable of initial proportion of parental cells.

This is shown more clearly in Figure 3. In all four conditions, we see that the growth rate of the resistant and parental cell lines depends on the initial proportion of parental cells. To capture the principle first-order part of this dependence, we consider a line of best fit between initial proportion of parental cells and the growth rates. See equations 1-8 in Supplementary Appendix C (or the matrix entries in Figure 4b) for these lines of best fit. Interpretable versions of these lines of best fit (see Supplementary Appendix D) can be expressed as a regularized fitness function w_S^C where $S \in \{P, R\}$ indexes the parental or resistant strategy and $C \in \{\text{DMSO}, \text{DMSO} + \text{CAF}, \text{Alectinib}, \text{Alectinib} + \text{CAF}\}$ indexes the experimental condition. For a description of regularization see Supplementary Appendix D. Finally, for a discussion of higher-order fitness functions, see Supplemental Appendix F.

In three of the conditions, the resistant cell growth rates increase with increased seeding proportion of parental cells, while the parental growth rates remain relatively constant (in the case of no CAFs) or slightly increasing (in the case of alectinib + CAFs). For example, in DMSO, this suggests that parental cells' fitness is independent of resistant cells: $w_P^{\text{DMSO}} = 0.025$. Parental fitness in DMSO could be well characterized as cell-autonomous. However, resistant cells in monotypic culture have approximately the same fitness as parental cells (Figure 2a), but they benefit from the parental cells in co-culture: $w_R^{\text{DMSO}} = 0.025 + 0.015p$ (where p is the proportion of parental cells). Their fitness is non-cell autonomous.

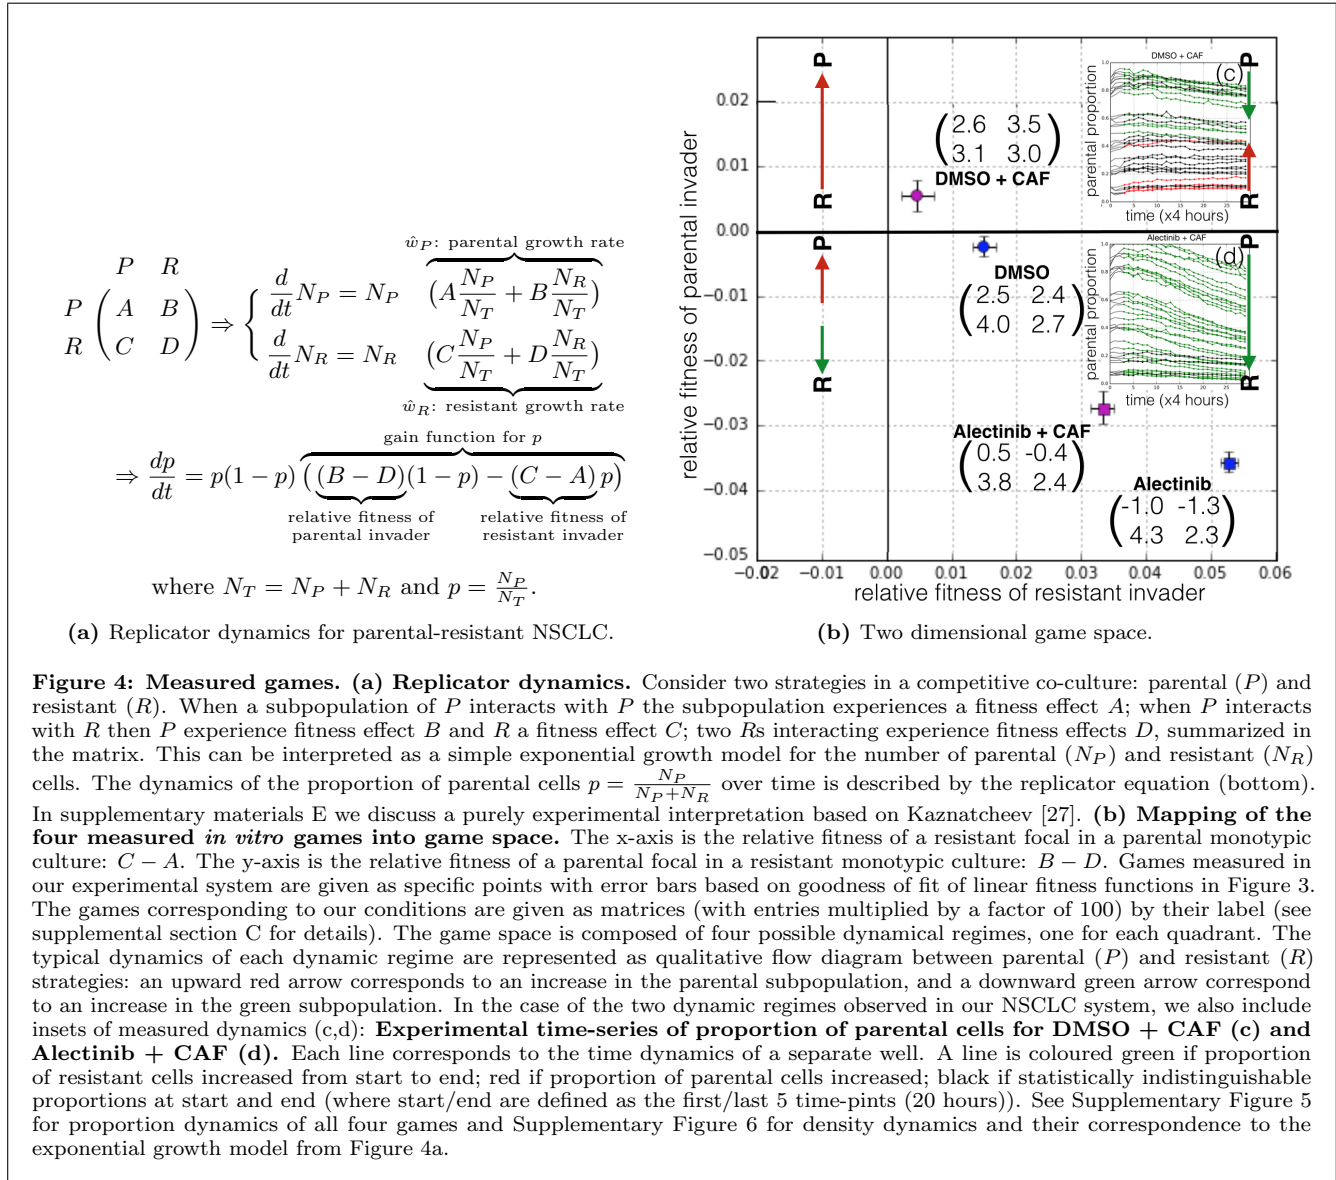
The positive coefficient in front of p suggests commensalism between resistant and parental cells, i.e. resistant cells benefit from the interaction with the parental cells, without exerting positive or negative impact on them.

The DMSO + CAF case differs from the other three in that we see a constant – although elevated $w_R^{\text{DMSO} + \text{CAF}} = 0.03$ – growth rate in resistant cells; but a linearly decreasing (in p) growth rate of parental cells: $w_P^{\text{DMSO} + \text{CAF}} = 0.025 + 0.01(1 - p)$ (or, equivalently: $w_P^{\text{DMSO} + \text{CAF}} = 0.03 - 0.01(\frac{1}{2} - p)$). This could be interpreted as CAFs switching the direction of commensalism between parental and resistant cells. Further, the stable co-existence enabled by the cross in fitness functions offers an alternative to the widely held view of pre-existent resistance as a rare cell-autonomous binary switch. We discuss this alternative in section ‘Treating the game & heterogeneity’.

Leader and Deadlock games in NSCLC

The tools of evolutionary game theory (EGT) are well suited for making sense of frequency-dependent fitness [17–19, 25, 28–31]. In EGT, a game is the rule mapping the population's strategy distribution to the fitness of individual strategies. Previous work has considered games like snowdrift [23], stag hunt [24], rock-paper-scissors [20], and public goods [22, 25] although none have been designed to measure evolutionary games directly. Instead, prior work followed a two track approach. They compared – usually qualitatively – the results of experiments (track one) to numeric simulations or analytic dynamics of these specific games (track two). We unite these parallel tracks by defining the effective game as an assayable hidden variable of a population and its environment. As such, we are not aiming to test EGT as an explanation of these particular data but rather using these data as an illustration of how we should operationally define games. With this game assay, we are aiming to quantitatively describe our system in the language of EGT.

To measure the game that describes the non-cell-autonomous interactions in NSCLC, we focus on the gain



function (see [19, 32] for a theoretical perspective): the increase in growth rate that a hypothetical player would get in ‘switching’ from being parental to resistant with all other variables held constant. In other words, we quantify how the difference between resistant and parental growth rates varies with initial proportion of parental cells. The relatively good fit of a linear dependence of growth rates on parental seeding proportion allows us to describe the interaction as a matrix game – a well-studied class of evolutionary games (see a description in Figure 4a). Note that this linearity is not guaranteed to be a good description for arbitrary experimental systems. For example, the game between the two Betaproteobacteria *Curvibacter* sp. AEP1.3 and *Duganella* sp. C1.2 was described by a quadratic gain function [24]. For extensions of the game assay to non-matrix games, see Supplemental Appendix F

Here, as a first step toward quantifying an evolutionary game, we are satisfied with the first-order approximation of interactions provided by matrix games. Two strategy matrix games have a convenient representation in a two dimensional game-space and can produce all possible linear gain functions. More importantly, from a linear gain function, it is possible to infer the corresponding matrix game. Since the game type and resultant dynamics are invariant under constant offsets to the columns, we can infer the game played by the cancer cells (see the model in Figure 4a for details). This is the output of our game assay. We plot the inferred games in a game-space spanned by the theoretical fitness advantage a single resistant invader would have if introduced into a parental monotypic culture versus the fitness advantage of a parental invader in a resistant monotypic culture; as shown in Figure 4b. In this representation, there are four qualitatively different types of games corresponding to the four quadrants, each of which we illustrate with a dynamic flow. We can see that the game corresponding to DMSO + CAF – although quantitatively similar to DMSO – is of a qualitatively different type compared to all three of the other combinations.

We can also convert our inferred fitness functions from Figure 3 into a payoff matrix. We do this by having each row correspond to a strategy’s fitness function with the column entries as the $p = 1$ and $p = 0$ intersects of this line of best fit. These are abstract phenomenological quantities that could be implemented by various biological or physical processes [27]. If we look at our empirical measurements for DMSO + CAF (upper-right quadrant Figure 4b) we see the Leader game, and Deadlock in the other three cases (we will use DMSO to illustrate the Deadlock game).

The Deadlock game observed in DMSO is in some ways the opposite of the popular Prisoner’s Dilemma (PD) game (in fact, Robinson & Goforth [33] call it the anti-PD). If we interpret parental as cooperate and resistant as defect then, similar to PD, each player wants to defect regardless of what the other player does (because $4.0 > 2.5$ and $2.7 > 2.4$; payoff numbers used in these examples are from the matrix entries we measured in Figure 4) but hopes that the other player will cooperate (because $4.0 > 2.7$). However, unlike PD, mutual cooperation does not Pareto dominate mutual defection (because $2.5 < 2.7$) but is instead strictly dominated by it. Thus, the players are locked into defection. In our

system, this corresponds to resistant cells having an advantage over parental in DMSO.

The Leader game observed in DMSO + CAF is one of Rapoport [34]’s four archetypal 2×2 games and a social dilemma related to the popular game known as Hawk-Dove, Chicken, or Snowdrift (in fact, Robinson & Goforth [33] call it Benevolent Chicken). If we interpret parental as ‘lead’ (for Snowdrift: wait) and resistant as ‘work’ (for Snowdrift: shovel) then similar to Snowdrift, mutual work is better than both leading (because $3.0 > 2.6$) and thus no work being done (for Snowdrift: both waiting and thus not getting out of the snowdrift) but each player would want to lead while the other works (because $3.5 > 3.0$). However, unlike Snowdrift, mutual work is not better than the “sucker’s payoff” of working while the other player leads (because $3.1 > 3.0$). Rapoport [34] sees this as a tension with a player switching from a “natural” point of mutual work to lead and thus benefit both players ($3.5 > 3.0, 3.1 > 3.0$), but if the second player also does the same and becomes a leader then all benefit disappears (because 2.6 is the smallest payoff). In our system, this corresponds to cells in the tumour experiencing selective pressure to lose some but not all of its resistance in DMSO + CAF.

Note that the above intuitive stories are meant as heuristics, and the effective games that we measure are summaries of population level properties [27, 35]: the population is the player and the two types of cancer cells are the strategies. This means that the matrix entries should not be interpreted as direct interactions between cells, but as general couplings between subpopulations corresponding to different strategies. The coupling term includes not only direct interactions, but also indirect effects due to spatial structure, diffusible goods, contact inhibition, etc.. But this does not mean that an effective game is not interpretable. For example, the Deadlock game captures the phenomenon of the resistant population always being fitter than parental (for example, in DMSO). We noted this effect intuitively in Figure 2 (also see section Mono vs mixed cultures & cost of resistance) from replicates being above the $y = x$ diagonal. Measuring a Deadlock game for DMSO with confidence intervals that do not extend outside the bottom right quadrant of the game space in figure 4b allows us to show the statistical significance of our prior intuitive understanding. Effective games allow us to quantify frequency-dependent differences in growth rates.

Treating the game & heterogeneity

To our knowledge, neither of the Leader and Deadlock games are considered in the prior theoretical EGT literature in oncology. Given that the Deadlock of drug-resistant over drug-sensitive cells is a challenge for classic models of resistance – as discussed in section Leader and Deadlock games in NSCLC – we would be particularly interested in theoretical models of resistance that produce the Deadlock game. In addition to challenging theorists by adding two new entries to the catalogue of games that cancers play, our results also support existing theoretical work in mathematical oncology that considers treatment (or other environmental differences) as changes between qualitatively different game regimes [17–19, 31]. In this framework, treatment has the goal not to directly

target cells in the tumour, but instead to perturb the game they are playing and allow evolution to drive down or control unwanted cancer subclones through competition. Before this study, this possibility has been largely taken as a theoretical postulate. In our system, we can view an untreated tumour as similar to DMSO + CAF and thus following the Leader game. Treating with alectinib (move to Alectinib + CAF) or eliminating CAFs through a stromal directed therapy (move to DMSO), moves the game into the lower-right quadrant of Figure 4b, and the game becomes a Deadlock game. This switch allows us to show that this theoretical construct of EGT – that treatment can qualitatively change the type of game – has an experimental implementation. Unfortunately, our system is *in vitro* and even if it generalized to *in vivo*: neither of the games leads to a therapeutically desirable outcome for the patient.

A particularly important difference between Leader and Deadlock dynamics is the existence of an internal fixed point in Leader but not in Deadlock. Fixed points are a property of equilibrium dynamics: in the most general case, even on very long timescales these fixed points might not be realized due to the evolutionary constraints of population size [36] or computation [37, 38]. Thus, it is important to check to what extent this qualitative difference can translate to a quantitative difference in finite time horizons. In our system, we can see a quantitative difference in the convergence towards the fixed point in the DMSO + CAF condition of Figure 4c, and no such convergence in the other three cases (Figure 4d for Alectinib + CAF; Supplementary Figure 5). Since the strength of selection (magnitude of the gain function) is small near a fixed point, the change in p also slows in the DMSO + CAF condition. We provide a more robust analysis of this in Supplemental Appendices C and F by focusing on a fixed *region* of finite width instead of just a fixed *point*.

Since the DMSO + CAF condition is our closest to an untreated patient, it might have important consequences for latent resistance. Many classical models of resistance assume a rare preexistent mutant taking over the population after the introduction of drug. In our experimental system, however, if the resistant strategy is preexistent then negative frequency dependent selection will push the population towards a stable polyclonal tumour of resistant and sensitive cells before the introduction of drug. This allows for much higher levels of preexisting heterogeneity in resistance than predicted by the classical picture. If similar games occur *in vivo* then such preexisting heterogeneity could be a possible *evolutionary mechanism* behind the speed and robustness of treatment resistance to targeted therapies in patients.

Conclusion

Drug-sensitive (parental) and resistant cells interact not only with alectinib, but also with each other and micro-environmental factors like CAFs. The relative fitness advantage of resistant over parental cells is a function of the initial proportion of sensitive cells – the gain function characterizing replicator dynamics. In co-culture, resistant cells have an advantage over parental cells even in DMSO. Measuring a linear gain function has enabled us to develop an assay

that represents the inter-dependence between parental and resistant cells as a matrix game. Not only are these games quantitatively different among the four environmental conditions – see Figure 4b – but they are also of two qualitatively different types: a Leader game in the case of DMSO + CAF and Deadlock in the other three cases.

This ability of treatment to qualitatively change the type of game being played provides the first direct experimental demonstration of the principal “don’t treat the player, treat the game”. This principle is often used in theoretic work when various treatments (or environmental differences) are modeled as changes in a game’s parameters that allow evolution to steer the tumour towards a more desirable result (for examples, see [17–19, 31, 39, 40]). Empirically, this principle has inspired or built support for interventions like buffer therapy [41], vascular renormalization therapy [42], or adaptive therapy [43] that target the micro-environment and interactions instead of just attacking the cancer cell population. But so far, the evolutionary games have not been measured directly but followed the two track approach: the games were prescribed based on biological or clinical intuition and then affirmed by qualitative agreement with clinical or experimental observations. The success of the Zhang *et al.* [43] trial suggests that therapeutic strategies based on modulating competition dynamics are feasible. This highlights the need for a formal experimental method like our game assay that directly measures the games that cancer plays.

Unfortunately, the Leader and Deadlock games that we measured are understudied in mathematical oncology, and we hope that our observation of them will motivate theorists to explore them in more detail. One difference between these game types is already clear: in the case of Leader there is negative frequency dependence selection toward a coexistence of parental and resistant cells – which we confirm for DMSO + CAF in Figure 4c – while for Deadlock there is selection towards a completely resistant tumour. Since DMSO + CAF is the closest analogue to a pre-treatment patient in our *in vitro* system, this suggests that there might be much higher levels of initial heterogeneity in drug resistance than prior theory would suggest. As such, we urge theorists to reconsider the assumption of the rare pre-existing resistant clone. Of course, our results are for a single *in vitro* system. But if similar results hold *in vivo* and/or for other cancers, it could help explain the ubiquity and speed of resistance that undermines our abilities to cure patients or control their disease in the long term. We will not know this unless we set out to quantify the non-cell autonomous processes in cancer. Building a catalogue of the games cancers play – by adopting our game assay in other cancers, and other experimental contexts – can help resolve this and other questions.

Acknowledgements

JGS would like to acknowledge the NIH Loan Repayment program for their generous support of his research in general as well as Miles for Moffitt for support of this work. We would also like to thank Mohamed Abazeed, Peter Jeavons, and Konstantine Kaznatcheev for helpful discussions.

Appendices

In these supplementary appendices, we develop and discuss the tools used to define and measure our game assay. The structure is as follows:

- A Description of materials used, the experimental method, and time-lapse microscopy.
- B Basic quantification of experimental images and how growth rates and associated error are measured within each well. This is the definition of fitness used throughout our text. Explains Figures 1 and 2 from the main text.
- C Defines parental proportions (p) and contrasts the evolutionary dynamics of proportions (Supplementary Figure 5) with the ecological dynamics of densities (Supplementary Figure 6). Defines the fitness functions based on lines of best fit between fitness and proportion from Figure 3 and presents the actual lines of best fit as equations 1-8. Explains how the linear fitness functions are converted into gain functions and games; and how those games are plotted in Figure 4b. Justifies the use of linear functions in terms of explanatory value and presents the model residuals in Supplementary Figure 7.
- D Presentation of interpretable fitness functions w_S^C from section ‘Frequency dependence in fitness functions’ of the main text in the context of regularization. As a visual check of the regularization, Supplementary Figure 8 shows what the games would look like if based on the regularized fitness functions w_S^C instead of the unregularized fitness functions \hat{w}_S^C shown in Figure 4b.
- E Experimental interpretation of replicator dynamics as an alternative to exponential model of Figure 4a.
- F Generalization of game assay to non-linear fitness functions. Provides the 3rd order fitness functions (Supplemental Figure 9) or independent mixed order fitness functions (Supplemental Figure 11) that would be selected by information criteria (Supplemental Figure 10) if one treats the game assay not as a definition of games but as a model selection problem for parameter fitting. As a qualitative check, Supplementary Figure 12 shows the agreement in game space between the higher order games and our measured matrix games.

A Materials and experimental method

Cell lines H3122 cell line was obtained from Dr. Haura (Moffitt Cancer Center). Cell line identity was validated by the Moffitt Cancer Center Molecular Genetics core facility using short tandem repeats (STR) analysis. Primary lung cancer associated fibroblasts were obtained from S. Antonia lab (Moffitt Cancer Center), following the protocols approved by the USF Institutional Review Board. CAFs were isolated as previously described in Mediavilla-Varela *et al.* [26] and expanded for 3-10 passages prior to the experiments. The alectinib resistant derivative cell line was obtained through escalating inhibitor concentration protocol, as described in Dhawan *et al.* [13]. Alectinib sensitive parental H3122 cells were cultured in DMSO for the same length of time, as the alectinib resistant derivate.

Stable GFP and mCherry expressing derivative cell H3122 cell lines were obtained through lentiviral transduction with pLVX-AcGFP (Clontech) and mCherry (obtained from K. Mitsiades, DFCI) vectors, respectively. We cultured both H3122 cells and CAFs in RPMI media (Gibco brand from Thermo Scientific), supplemented with 10% FBS (purchased from Serum Source, Charlotte, NC). Regular tests for mycoplasma contamination were performed with MycoScope PCR based kit from GenLantis, San Diego, CA.

Experimental set-up. The cells were harvested upon reaching 70% confluence and counted using Countess II automatic cell counter (Invitrogen). CAFs were counted manually to avoid segmentation artifacts. Mixtures of parental and resistant H3122 cells were prepared at 8 different ratios: all-resistant, 9:1 resistant to parental, 4:1, 3:2, 2:3, 1:4, 1:9, and all-parental. For the determination of competitive growth rates, 2,000 H3122 cells from the 8 mixtures were seeded with or without 500 CAF cells in 50 μ L RPMI media per well into 384 well plates (Corning, catalogue #7200655), with different ratios of differentially labelled parental and alectinib resistant variants: with 6 wells used for each resistant:parental ratio in each of the 4 conditions. 20 hours after seeding, alectinib – purchased from ChemieTek (Indianapolis, IN) – or DMSO vehicle control, diluted in 20 μ L RPMI was added to each well, to achieve final alectinib concentration of 500 nM/L [14]. Time lapse microscopy measurements were performed every 4 hours in white light, as well as green and red fluorescent channels using Incucyte Zoom system from Essen Bioscience.

B Measuring population sizes and fitnesses

Fluorescent area as units of population size We measured fluorescent area from time-lapse images via python code using the OpenCV package and used this as our units of size for populations. See Kaznatcheev [27] for a discussion of fitness and replicator dynamics under various definition of population size. We cleaned images by normalizing them (GFP and mCherry intensities vary over different orders of magnitude), removed vignetting with CLAHE, and finally thresholded to identify fluorescent regions. We eliminated salt-and-pepper noise from the thresholded images with the opening morphological transform. See Figures 2a,b,d,e for examples of the image analysis. The resultant area is then taken as a measure of population size for the purposes of computing fitnesses.


Growth rate as fitness. We use growth rate as our measure of fitness. In order to minimise the impact of growth inhibition by confluency, we analyzed the competitive dynamics during the first 5 days of culture, when the cell population was expanding exponentially. See section E for a discussion of the impact of measurement length. We learned growth rate along with a confidence interval from the time-series of population size in each well using the Theil-Sen estimator [44, 45]. Since the Theil-Sen estimator is a rank-based median method (unlike least-squares, which is a numeric-based mean method), it is more robust to noise and does not need to choose between a linear or log representation for computing the error-term (since log transforms do not change rank orders). The robustness to rare but large magnitude noise is useful for our purposes because such errors do not reflect biological function or noise but are more likely to be due to errors in image processing, for example in response to sudden condensation on the well plate. See Figures 2c,f for examples of fitting.

The learned parental growth rate and resistant growth rate of each well are used as the y coordinates in the monoculture experiments of Figure 1 (along with errors on the growth rate) and as the x and y coordinates of the main part of Figure 2. Due to too much information content, the errors on the growth rates are omitted in Figure 2, but they are shown explicitly as error-bars in Figure 3. Note that this means that each point in Figure 1 and the main part of Figure 2 and each pair of points in Figure 3 (one green and one red at the same x-position) correspond to one biological replicate, with the error term coming from the confidence interval on the growth-rate estimate from the 30 time-series points that we recorded for each biological replicate (see section E on how this relates to the accuracy-precision trade-off). Thus, each of the 6 wells corresponding to a given resistant:parental ratio (in each of the 4 conditions) has its own independent growth rate with associated error. The wells are not averaged together: each acts as its own data point (with noise) for later analysis (that propagates the noise).


Figure 2 as map of analysis flow. Along with showing all the data, Figure 2 serves as a map to the above analysis pipeline. The inset subfigures and figure can be understood in the following order:

- [a,b,d,e] Within each image from the series generated by time-lapse microscopy: identify the fluorescent regions for GFP and mCherry and calculate their areas to serve as units of population size (GFA and CFA).
- [c,f] For parental (mCherry) and resistant (GFP) plot the population sizes from each image in the series on a semilog grid as population vs. time. Find the slope of the two lines to serve as parental and resistant fitness.
- [main] Use the parental fitness as x value and resistant as y value to plot each well as a data-point according to the above process, and color the point according to its experimental condition (with opacity for initial parental proportion; see Section C). For ease of viewing: put a convex hull binding polygon around each well data-point dependent on their experimental condition.

Given the complexity of Figure 2, it is tempting to ask for a simple summary statistic of the data in the main figure. But it is not reasonable to ask of the “average” growth rate in Figure 2 because each point differs not only along the four experimental conditions of the environment, but also along the micro-environmental conditions of the initial parental proportion (represented by the opacity that is explained in SA C). Averaging over this information would be akin to assuming that the growth rates are cell-autonomous. It would be attributing the variance in growth rates to noise instead of the independent variable or initial parental proportion. As such, the game assay developed in rest of the paper can be viewed as a method for summarizing Figure 2 when the underlying process is non-cell-autonomous. And the games derived through Figure 3 and presented in Figure 4b are the summary of the data in the main part of Figure 2.



Supplementary Figure 5: Evolutionary dynamics of proportion of parental cells versus time for competition of parental vs. resistant NSCLC. Each line corresponds to the time dynamics of a separate well. A line is coloured green if proportion of resistant cells increased from start (time step 3 to 8) to end (time step 24 to 29); red if proportion of parental cells increased; black if statistically indistinguishable proportions at start and end.

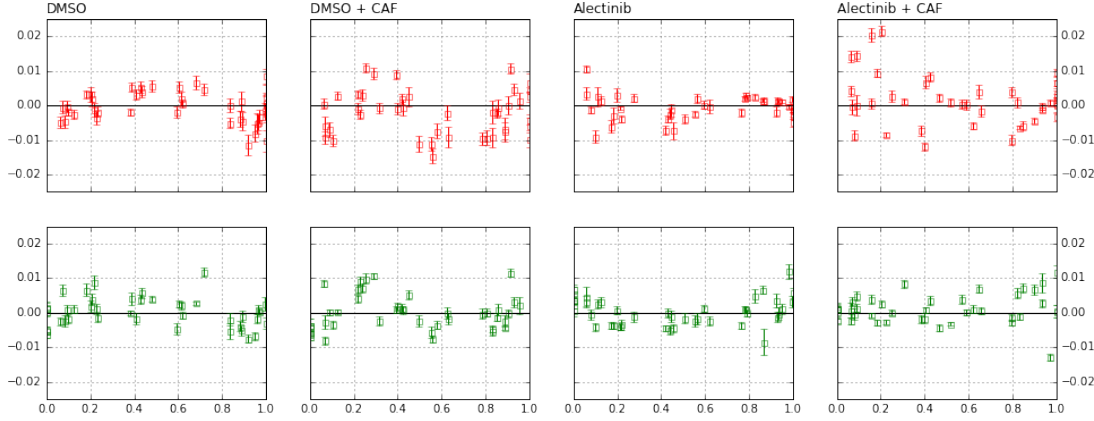


Supplementary Figure 6: Dynamics of population sizes of resistant cells versus parental cells. Axis are fold change in CNU normalized from the seeding proportion of each well, with x-axis for parental and y-axis for resistant. Foreground: raw data. Each line corresponds to the time dynamics of a separate well. A line is coloured green if proportion of resistant cells increased from start to end; red if proportion of parental cells increased; black if statistically indistinguishable proportions at start and end (using the same conventions as Figure 5). Background: flow diagram for model from Figure 4a. Each coloured point shows the proportion of parental-resistant (red-green) at that point. Arrow going from red to green indicates parental proportion increased, if from green to red then resistant proportion increased.

C Measuring fitness functions and games

Proportions. Since raw population sizes have different units (GFP Fluorescent Area (GFA) vs mCherry Fluorescent Area (RFA)), we converted them to common cell-number-units (CNU) by learning the linear transform that scales GFA and RFA into CNU. We defined proportions based on this common CNU as $p = N_P / (N_P + N_R)$ where $N_{\{P,R\}}$ is the CNU size of parental and resistant populations. The transform of GFA to RFA into CNU is associated with an error that is propagated to measures of p as σ_p . Thus, although we used 8 different ratios of resistant to parental cells with 6 wells per condition seeded at each of the ratios, we do not average over these 6 wells but associated each with its own proportion $p \pm \sigma_p$ from the initial image. This helps us control for systemic noise from field of view and our image processing algorithm. The time dynamics of p can be seen in the insets of Figure 4b for DMSO and DMSO+CAF or in Supplementary Figure 5 for all conditions.

Neglecting ecological dynamics. Throughout this report, we focus on evolutionary dynamics: changes in proportion of strategies. However, one could also consider the ecological dynamics: changes in densities of strategies. It is not only proportions that are changing in our experimental system but also the densities. These ecological dynamics are not the focus of our report, but we present them in Supplementary Figure 6 for completeness. Here, we also compare the prediction of the model based on our measured games and the exponential growth interpretation in Figure 4a to the observed data. There is overall agreement between data and model. But this is based on the traditional two track approach on qualitative agreement. Instead, we prefer to focus on the single track measurement of evolutionary dynamics described in the rest of these supplementary materials. Future work can aim to extend our



Supplementary Figure 7: Residuals for the fitness functions. The x-axis is proportion and y-axis is residuals of the lines of best fit from Figure 3 for parental (Red,Top) and resistant (Green,Bottom).

approach to also include ecological dynamics.

Lines of best fit as fitness functions. To measure the fitness functions we plotted fitness of each cell-type in each well vs seeding proportion (p) of parental cells in Figure 3. The x-axis proportion of parental cells (p) was computed from the first time-point: see section E for an interpretation of this as a measurement of dp/dt or as a series of competitive fitness assays. We estimated the line of best-fit and error on parameters for this data using least-squares weighted by the inverse of the error on each data point (i.e. $\text{weight}_{p,w} = 1/\sqrt{\sigma_p^2 + \sigma_w^2}$). This provides the error estimates on the line's parameters that we use later. The lines of best fit (with coefficients rounded to the thousandths for presentation) from weighted least-squares are:

$$\hat{w}_P^{\text{DMSO}} = 0.025 - 0.001(1 - p) = 0.025p + 0.024(1 - p) \quad (1)$$

$$\hat{w}_R^{\text{DMSO}} = 0.027 + 0.013p = 0.04p + 0.027(1 - p) \quad (2)$$

$$\hat{w}_P^{\text{DMSO + CAF}} = 0.026 + 0.009(1 - p) = 0.026p + 0.035(1 - p) \quad (3)$$

$$\hat{w}_R^{\text{DMSO + CAF}} = 0.03 + 0.001p = 0.031p + 0.03(1 - p) \quad (4)$$

$$\hat{w}_P^{\text{Alectinib}} = -0.01 - 0.002(1 - p) = -0.01p - 0.013(1 - p) \quad (5)$$

$$\hat{w}_R^{\text{Alectinib}} = 0.023 + 0.02p = 0.043p + 0.023(1 - p) \quad (6)$$

$$\hat{w}_P^{\text{Alectinib + CAF}} = 0.005 - 0.009(1 - p) = 0.005p - 0.004(1 - p) \quad (7)$$

$$\hat{w}_R^{\text{Alectinib + CAF}} = 0.024 + 0.014p = 0.038p + 0.024(1 - p) \quad (8)$$

Summarizing fitness functions as games. For the final column of our presentation of \hat{w}_S^C in equations 1-8, we rewrote the fitness functions in a suggestive form of $\hat{w}_P^C = Ap + B(1 - p)$ and $\hat{w}_R^C = Cp + D(1 - p)$. This is done to show at a glance where the matrix entries in Figure 4b come from. This is because the $p = 0$ and $p = 1$ intercepts of the fitness functions serve as the entries of the game matrices. Note that in Figure 4b, we multiplied the entries by 100 for easier presentation. The game point are calculated from the matrices as $x := C - A$ and $y := B - D$, and the error is propagated from the error estimates on fitness function's parameters.

Gain functions, game space, and fixed points A particularly important equation for studying two strategy games is the gain function. This represents the relative fitness difference between two strategies. Thus, it is a measure of selection strength and a proxy for the rate of evolution. The parental gain function (i.e. gain function for p in Figure 4a and equation 23) is given by $\hat{g}_P^C(p) = \hat{w}_P^C(p) - \hat{w}_R^C(p)$; and the resistant gain function (i.e. gain function for $q = 1 - p$) is $\hat{g}_R^C(q) = \hat{w}_R^C(1 - p) - \hat{w}_P^C(1 - p) = -\hat{g}_P^C(1 - p)$. The end-points of this gain function determine the game coordinates in the game space of Figure 4b with $(x, y) := (\hat{g}_R^C(0), \hat{g}_P^C(0)) = (-\hat{g}_P^C(1), \hat{g}_P^C(0))$. These points can be interpreted as the idealized quantities of relative fitness of a resistant invader in parental monoculture ($\hat{g}_P^C(0)$) and relative fitness of a parental

invader in resistant monoculture ($\hat{g}_R^C(0) = -\hat{g}_P^C(1)$). If these two coordinates have the same sign then the gain function has to cross 0 in getting from $p = 0$ to $p = 1$ and thus the dynamics have a fixed point. If the two coordinates are both positive (top right quadrant of Figure 4b) then the fixed point is stable, if both are negative (bottom left quadrant of Figure 4b) then the fixed point is unstable. In our experimental system, only the DMSO + CAF condition has a fixed point at 0.53 ± 0.14 (rounded to the nearest percent).

Width and height of fixed regions. Since we propagate the errors on our measurement from the image all the way to the game, we find it more helpful to think of an experimental fixed point not as a point but as a fixed region $p \in (0.39, 0.67)$ of finite width. This can provide an alternative explanation for the apparent slowness of convergence to the fixed point in Figure 5. Some of the fixed region’s width is noise in measurement, but some could be due to true variance between wells: in particular, even if the reductive game is the same, the spatial structure will be slightly different in each well and thus there will be a slightly different effective game. As such, apparent slowness in Figure 5 might be from different lines being very close to slightly different fixed points that are all within the fixed region’s width. An alternative view to width is in terms of height: a fixed region corresponds not just to the point where $\hat{g}_P^{\text{DMSO} + \text{CAF}}$ crosses 0 but to the region where the gain function crosses 0 ± 0.0014 (rounded to the nearest thousandth). We call this the fixed region’s height (and use it in section F). This height is due to propagation of error and can be interpreted as our measurement not being able to distinguish relative growth rates in $(-0.0014, 0.0014)$ from zero. In the case of the other three conditions (DMSO, Alectinib, and Alectinib + CAF), in going from $p = 0$ to $p = 1$, the gain function do not pass within their fixed region height of zero, and thus no fixed regions exist.

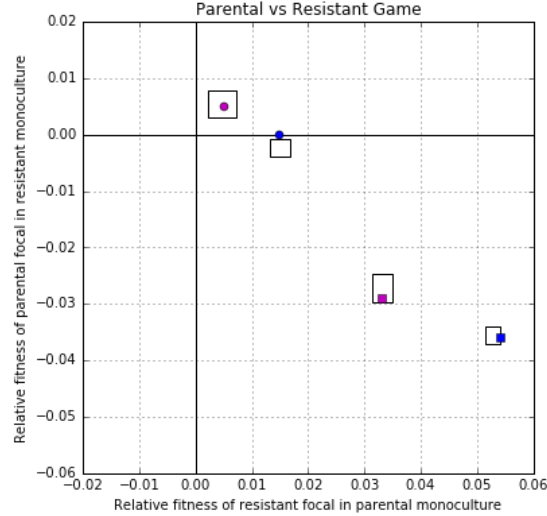
Lines and matrix games. Although slight deviations from a linear fit – that might not be attributable to noise alone – might be present in the data (see Supplementary Figure 7), we do not think that they justify considering higher-order fitness functions (although we discuss higher-order functions in section F for completeness). This is due to the higher explanatory value of linear models and our hope to influence the well-established study of matrix games in microscopic systems. Some good EGT work has recently been done on non-linear games [24, 25, 31], but this is very little compared to the immense literature on matrix games. More importantly, we think that our focus on matrix games is better viewed not from the perspective of model selection but rather as an operational definition of effective games. We are not aiming to provide the best or most predictive account of non-small-cell lung cancer in the petri dish, but rather a method for measuring (matrix) games. If the error of the measured (matrix) games ends up very high – which is not the case from the error bars in Figure 4b – then we know that this first order approximation of interactions is not sufficient and higher orders should be pursued. However, we will not know this unless we first have a robust method for measuring the lower order terms.

D Regularization and interpretable fitness functions

Regularization is a machine learning technique for reducing over-fitting by biasing towards more succinct models. It is the use of *a priori* knowledge on what constitutes a simpler or more likely model to anchor our inference. A classic example of this is preferring lower-order over higher-order polynomials for describing data unless there is overwhelming evidence otherwise. Of course, what constitutes overwhelming evidence depends on the goals of the scientists. If the only goal is prediction then cross-validation is a good way to test how heavily inference should be regularized. But if the goal is explanation then accordance with existing theory is another important factor to consider.

As such, our choice of focusing on linear fitness function in section C and Figure 3 can be seen as a form of regularization. In particular, we can see our inference procedure as either restricted to the hypothesis class of linear functions, or as considering the hypothesis class of all polynomials but with prohibitively high costs for non-zero components (l_0 regularization) on orders beyond linear. But we prefer to think of it in terms of operationalization. By introducing a game assay, we are defining the hidden variable of (matrix) games in terms of the measurement procedure that we described in sections B and C.

Interpretable fitness functions. An uncontroversial case of regularization in our report is the presentation of w_S^C in section ‘Frequency dependence in fitness functions’. There, we restrict beyond linear fitness functions to focus on conceptually simple ones. In particular, we favor cell-autonomous functions over frequency dependent ones (i.e. l_0 regularization on the fitness function coefficients) and we favor coefficients that are shared between different S and C . This results in the following regularized fitness functions:



Supplementary Figure 8: Mapping of the regularized fitness functions for the four conditions into game space. The x-axis is the relative fitness of a resistant focal in a parental monotypic culture: $C - A$. The y-axis is the relative fitness of a parental focal in a resistant monotypic culture: $B - D$. Games measured in our experimental system are specified by the bounding boxes corresponding to the range of their errors. The games corresponding to the regularized fitness functions in equations 9-16 are given as points.

$$w_P^{\text{DMSO}} = 0.025 \quad (9)$$

$$w_R^{\text{DMSO}} = 0.025 + 0.015p \quad (10)$$

$$w_P^{\text{DMSO} + \text{CAF}} = 0.025 + 0.01(1 - p) \text{ (or } 0.03 - 0.01(\frac{1}{2} - p)) \quad (11)$$

$$w_R^{\text{DMSO} + \text{CAF}} = 0.03 \quad (12)$$

$$w_P^{\text{Alectinib}} = -0.010 \quad (13)$$

$$w_R^{\text{Alectinib}} = 0.025 + 0.018p \quad (14)$$

$$w_P^{\text{Alectinib} + \text{CAF}} = 0.005 - 0.009(1 - p) \text{ (or } 0.009(\frac{1}{2} - p)) \quad (15)$$

$$w_R^{\text{Alectinib} + \text{CAF}} = 0.025 + 0.013p \quad (16)$$

Note that for both P and R strategies, we used the proportion of the other strategy ($1 - p$, p) as the parameter that captures the non-cell-autonomous contribution. In equations 11,15, we also consider the parameter $\frac{1}{2} - p$ because of the elegant form it provides.

We can compare these regularized fitness functions w_S^C to the non-regularized \hat{w}_S^C in equations 1-8. As can be seen, all w_S^C are close to their respective \hat{w}_S^C and are actually within the error estimates on \hat{w}_S^C . We can see the regularization in action with a push towards a constant base fitness of 0.025 shared by $w_{\{P,R\}}^{\text{DMSO}}$, $w_P^{\text{DMSO} + \text{CAF}}$, and $w_R^{\{\text{Alectinib}, \text{Alectinib} + \text{CAF}\}}$. The absence of frequency dependent perturbation terms for $w_P^{\{\text{DMSO}, \text{Alectinib}\}}$ and $w_R^{\text{DMSO} + \text{CAF}}$ suggests that these strategies can be explained in terms of cell-autonomous processes. However, the other strategies in the other contexts ask for a non-cell-autonomous explanation.

Games from interpretable fitness functions. For a visual confirmation that the regularization of w_S^C in equations 9-16 are reasonable, we can transform them into regularized games. We do this in the same way as we did for transforming the non-regularized \hat{w}_S^C in equations 1-8 into the game-points of figure 4b. The results are in Supplementary Figure 8. The regularized games (points) are within the confidence rectangles of the measured games (boxes), with the exception of DMSO which is just outside its box. This is reasonable given that the boxes correspond to error: i.e. around 2/3rds confidence.

E Experimental definition of replicator dynamics

Consider a well that is seeded with an initial number N_P^I of parental and N_R^I of resistance cells; total number $N^I = N_P^I + N_R^I$. Let $N_{\{P,R\}}^F$ be the number of {parental,resistant} cells after being grown for an amount of time Δt . From this, the experimental growth rate can be defined based on fold change as:

$$w_{\{P,R\}} := \frac{N_{\{P,R\}}^F - N_{\{P,R\}}^I}{N_{\{P,R\}}^I \Delta t} \quad (17)$$

this can be rotated into a mapping $N^I \mapsto N^F$ given by $N_{\{P,R\}}^F = N_{\{P,R\}}^I (1 + w_{\{P,R\}} \Delta t)$.

By defining the initial and final proportion of parental cells as $p^{\{I,F\}} = N_P^{\{I,F\}} / N^{\{I,F\}}$, we can find the mapping:

$$p^F = \frac{N_P^F}{N^F} = p^I \frac{1 + w_P \Delta t}{1 + \langle w \rangle \Delta t} \quad (18)$$

where $\langle w \rangle = p^I w_P + (1 - p^I) w_R$. This is the discrete-time replicator equation.

We can approximate this discrete process with a continuous one by defining $p(t) = p^I$, $p(t + \Delta t) = p^F$ and looking at the limit as Δt gets very small:

$$\dot{p} = \lim_{\Delta t \rightarrow 0} \frac{p(t + \Delta t) - p(t)}{\Delta t} \quad (19)$$

$$= \lim_{\Delta t \rightarrow 0} \frac{p^I}{\Delta t} \left(\frac{1 + w_P \Delta t}{1 + \langle w \rangle \Delta t} - 1 \right) \quad (20)$$

$$= \lim_{\Delta t \rightarrow 0} p \frac{w_P - \langle w \rangle}{1 + \langle w \rangle \Delta t} \quad (21)$$


$$= p(w_P - \langle w \rangle) \quad (22)$$

$$= p(1 - p) \underbrace{(w_P - w_R)}_{\text{gain fn for } p} \quad (23)$$

Thus, we recover replicator dynamics as an explicit experimental interpretation for all of our theoretical terms. Note that we did not make any assumptions about if things are inviscid or spatial; if we are talking about individual or inclusive fitness; or, if we have growing populations in log phase or static populations with replacement. All of these microdynamical details are buried in the definition of experimental fitness. This allows us to focus on effective games [27] and avoid potential confusions over aspects like spatial structure [35].

Better estimates of w . The problem with the definition of w in equation 17 is that it depends on just two time points, and thus not good for quantifying error. In our experimental system, we are able to peek inside the system with time-lapse microscopy. This allows us to get more than just the initial and final population sizes and replace fold-change by the more specific measurements of inferred growth rates for $w_{\{P,R\}}$ that we describe in section B. An advantage of this approach is that the goodness-of-fit of the exponential growth model provides a good estimate of the error associated with each measurement of w . Thus, we are able to quantify error within each well and not just between experimental replicates in different wells with similar initial conditions.

Accounting for finite Δt . The small time definition of the derivative can be thought of as a way to approximate a function by local linearizations. It is why for simulations, modelers often use the discrete time replicator dynamics to represent continuous time replicator dynamics: effectively using the discretization as a simple ODE solver/plotter. In the limit of Δt going to 0, this linearization recapitulates the function. Unfortunately in practice, our experimental system cannot take the limit as Δt goes to 0 because of a precision-accuracy trade-off. Accuracy increases as Δt decreases because the continuous dynamics is approximated by more and more, shorter and shorter straight lines. But – from an experimentalist’s perspective – the precision decreases because any measurement is noisy: if we measure growth rate over a shorter period of time then we are less certain whether our measurement reflects reality or noise. For very short measurements, we might get higher accuracy (assuming biological factors like time from seeding to adherence could be ignored) but would have incredibly low precision (due to only one, two or three time points from which to calculate growth rate). As we increase the time of the experiment, the accuracy might decrease but the precision will tend to increase. This is a classic



Supplementary Figure 9: Cubic Fitness functions for competition of parental vs. resistant NSCLC. For each plot: growth rate with confidence intervals versus initial proportion of parental cells. This is the same data as Figure 3. Red data points are growth rates of parental cells, and green for resistant cells. Dotted lines represent the 3rd-order (cubic) fitness function of the least-squares best fit. The black dotted line is the gain function for parental (see Figure 4a), it is well below the $y = 0$ line in the Alectinib conditions (indicating the strong advantage of resistance) and thus cut out of the figure.

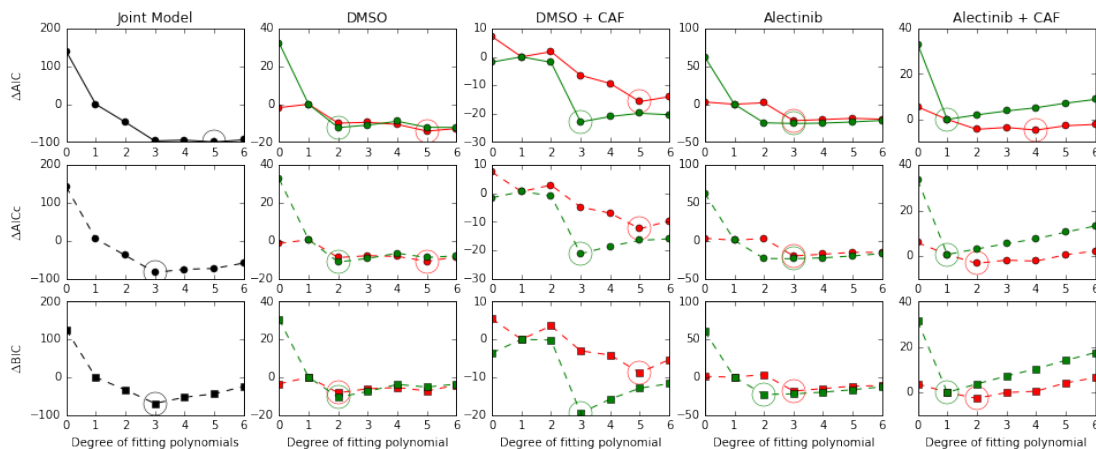
trade off between random noise (low precision) and systematic noise (linearization being progressively less accurate over larger Δt). Since each of our growth rate measurements has an associated error term (see section B), we quantify the random and systematic noise together and propagate it throughout our analysis. Given the biological constraints of our system, we judged that 5 days was a good trade-off point. This will most likely be different for other experimental systems.

Given that w are defined over a finite range of time, we need to pick a particular time-point to associate each measurement with. As is common for discrete time process, we attribute the value of the growth rate to the initial point. In particular, this means that when we make $w_{\{P,R\}}$ a function of p in the main text, then the values of growth rate are attributed to the initial proportion of parental cells and not the final one. This customary choice is further reinforced by the fact that we have a less noisy estimate of initial proportions of cells than of the final, and so other definitions would lead to less precise measurements. Finally, our procedure can be viewed as standard competitive fitness assays but with initial ratio of the two types as a varied experimental parameter. Thus, for consistency with both theoretical and experimental literature, we associated the growth rates with the initial – more controlled – seeding proportion.


F Generalizing the game assay to non-linear fitness functions

The game assay that we presented above is more interpretable and its output more easily plotted for matrix games with linear fitness functions. And in the case of our experimental system, the linear fitness functions provide an adequate fit for our purposes. However, that does not mean that the game assay has to be used only for linear games. Whereas section C used linear functions as the hypothesis class for fitting the growth-rate vs. proportion, one could use any other class of functions. An obvious candidate is polynomial fitness functions of orders higher than 1. We provide an example in Supplementary Figure 9 of a 3rd-order (cubic) fit. Visually, the cubic provides a better fit than the linear one in Figure 3, which is to be expected from the extra degrees of freedom. But qualitatively it provides the same interpretation as the linear fitness functions, including the same number of fixed points. In particular, DMSO + CAF has a single fixed point at $p = 0.52$ and (using the fixed region height from section C) a single fixed region for $p \in (0.04, 0.668)$. This is much like the linear fit, but the fixed point region is expanded. The other three conditions (DMSO, Alectinib, and Alectinib + CAF) have no fixed points and no fixed regions.

Information criteria for non-linear fits. If we treat the game assay not as a measurement and definition of games but as a model selection problem for parameter fitting then it becomes important to quantify the trade-off between the goodness of fit and model simplicity. For this, we can use techniques like the Akaike information criterion (AIC), its small-sample size correction (AICc), or the Bayesian information criterion (BIC) – or any other statistical model selection procedure. Given that (i) a polynomial of degree d has $k = d + 2$ degrees of freedom as a statistical model (+1 for zeroth order term, +1 for noise term); the eight models (4 conditions, 2 fitness functions per condition) are trained on $n = 42$ data points each; and AIC/BIC only works reasonably when $n \gg d$. For example, given




Supplementary Figure 10: AICs and BICs for best polynomial fits of a given degree (up to additive offset). In red are AIC and BIC values for models parental cell fitness functions, while in green are AIC and BIC values for models of resistant cell fitness functions. Circles surround the minimum of AIC and BIC.



Supplementary Figure 11: Fitness functions for competition of parental vs. resistant NSCLC as selected by BIC. For each plot: growth rate with confidence intervals versus initial proportion of parental cells. This is the same data as Figure 3. Red data points are growth rates of parental cells, and green for resistant cells. Dotted lines represent the fitness function of the least-squares best fit for models selected by BIC. These are a linear model for resistant fitness function in Alectinib + CAF; quadratic models for parental fitness function in Alectinib + CAF, resistant fitness function in Alectinib, and both fitness functions in DMSO; cubic The black dotted line is the gain function for parental (see Figure 4a), it is identical with the $y = 0$ line in the DMSO + CAF condition (indicating equal fitness for the two strategies) and well below it in the Alectinib conditions (indicating the strong advantage of resistance) and thus not visible in the figure.

our relatively small dataset for each model, Burnham & Anderson [46] would advocate to always prefer AICc over AIC (they suggest $n/k < 40$ as the cut off). Hence, we show the results of all three of AIC, AICc and BIC for polynomial fitness functions for degree $d \leq 6$ in Supplemental Figure 10. In this figure, a better model corresponds to a lower AIC, AICc or BIC value (lower on the y-axis). Since constant offsets in the information criteria do not matter for model selection, the axes are set so that the linear model has $\Delta\{AIC, AICc, BIC\} = 0$. The leftmost column of Supplemental Figure 10 considers the joint product model where each fitness function has the same degree – for the $d = 1$ model, this would correspond to the linear game assay as presented in section C. Both AICc and BIC select the 3rd-degree polynomial model that we discussed above. AIC doesn't differentiate strongly between the 3rd, 4th, 5th and 6th degree, but prefers slightly the 5th degree. Too much emphasis should not be placed on AIC however, given the number of parameters compared to sample size [46]. The four right columns of Supplemental Figure 10 consider independent models for each of the fitness functions across the 4 different conditions – so a total of 8 models. At the cost of extra researcher degrees of freedom, it is possible to look at the fits where the model for each of the 8 fitness functions is selected independently. Such a fit, as selected by BIC, is shown in Supplemental Figure 11. Note the two extra crossings of zero by the gain function in the DMSO + CAF case.



Supplementary Figure 12: Mapping of the AIC and BIC selected fitness functions for the four conditions into game space. The x-axis is the relative fitness of a resistant focal in a parental monotypic culture. The y-axis is the relative fitness of a parental focal in a resistant monotypic culture. Games measured in our experimental system are specified by the bounding boxes corresponding to the range of their errors. The games corresponding to the AIC and BIC selected models are given as points.

Plotting nonlinear games. Just like with the linear games, it is possible to plot nonlinear games in our 2D game space based on the $p = 0$ and $p = 1$ endpoints of the gain function. We do this in Supplementary Figure 12 with each point labeled by the degree of the corresponding polynomial fitness functions. Unsurprisingly, at a brief glance there is broad qualitative agreement – all (but one) points are in the same quadrant as the linear model – although little quantitative agreement with the linear game assay – most points are outside of the error-box corresponding to the linear game. However for a general nonlinear game, unlike with linear games, two points in the same quadrant might not correspond to the same qualitative kind of dynamic. In particular, for a general nonlinear game, a quadrant only tells us the parity of the number of roots in $(0, 1)$ – where roots are counted by their multiplicity – and the order of alternations on the flow. For more on discrete flow alternation representation of gain functions, see Peña *et al.* [32].

Fortunately, for our particular experimental system the above generality is not realized. In particular, for all degrees of the DMSO, Alectinib, and Alectinib + CAF games the gain functions have no fixed points – just like the linear case. For DMSO + CAF, degrees $d = \{1, 2, 3\}$ have one fixed point and $d = \{5, 6\}$ have 3 fixed points (although only two fixed regions: $p \in \{(0.07, 0.25), (0.38, 0.46)\}$ for $d = 5$ and $p \in \{(0.09, 0.31), (0.44, 0.50)\}$ for $d = 6$). For $d = 0$ it is impossible for any model to be in the top right or bottom left quadrant – since no constant line can be both negative and positive – and there is no fixed point, but the fitness difference for DMSO is so tiny that there is a single fixed region for the whole space $p \in (0, 1)$. The real outlier for DMSO is $d = 4$ since it has two fixed points (and is thus in the bottom right quadrant) and two fixed regions at $p \in \{(0.07, 0.11), (0.29, 0.40)\}$. Thus, the existence of fixed point(s) in DMSO + CAF and absence of fixed points in the other conditions is robust across the nonlinear models. The exact position of the fixed point(s) in DMSO + CAF, however, is not as robust to model choice.

Main Text & Appendix References

1. Merlo, L. M., Pepper, J. W., Reid, B. J. & Maley, C. C. Cancer as an evolutionary and ecological process. *Nature Reviews Cancer* **6**, 924–935 (2006).
2. Heppner, G. H. Tumor heterogeneity. *Cancer Research* **44**, 2259–2265 (1984).
3. Ibrahim-Hashim, A. *et al.* Defining cancer subpopulations by adaptive strategies rather than molecular properties provides novel insights into intratumoral evolution. *Cancer Research*, 2844 (2017).

4. Scott, J. & Marusyk, A. Somatic clonal evolution: A selection-centric perspective. *Biochimica et Biophysica Acta: Reviews on Cancer* (2017).
5. Shaw, A. T. *et al.* Crizotinib versus chemotherapy in advanced ALK-positive lung cancer. *New England Journal of Medicine* **368**, 2385–2394 (2013).
6. Peters, S. *et al.* Alectinib versus crizotinib in untreated ALK-positive non-small-cell lung cancer. *New England Journal of Medicine* **377**, 829–838 (2017).
7. Shaw, A. T. & Engelman, J. A. ALK in lung cancer: past, present, and future. *Journal of Clinical Oncology* **31**, 1105–1111 (2013).
8. Gillies, R. J., Verduzco, D. & Gatenby, R. A. Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. *Nature Reviews Cancer* **12**, 487–493 (2012).
9. Katayama, R., Lovly, C. M. & Shaw, A. T. Therapeutic targeting of anaplastic lymphoma kinase in lung cancer: a paradigm for precision cancer medicine. *Clinical Cancer Research* **21**, 2227–2235 (10 2015).
10. Marusyk, A. *et al.* Spatial proximity to fibroblasts impacts molecular features and therapeutic sensitivity of breast cancer cells influencing clinical outcomes. *Cancer Research* **76**, 6495–6506 (2016).
11. Yamada, T. *et al.* Paracrine receptor activation by microenvironment triggers bypass survival signals and ALK inhibitor resistance in EML4-ALK lung cancer cells. *Clinical Cancer Research* **18**, 3592–3602 (2012).
12. Ou, S.-H. I. *et al.* Alectinib in crizotinib-refractory ALK-rearranged non-small-cell lung cancer: a phase II global study. *Journal of Clinical Oncology* **34**, 661–668 (2015).
13. Dhawan, A. *et al.* Collateral sensitivity networks reveal evolutionary instability and novel treatment strategies in ALK mutated non-small cell lung cancer. *Scientific Reports* **7** (2017).
14. Seto, T. *et al.* CH5424802 (RO5424802) for patients with ALK-rearranged advanced non-small-cell lung cancer (AF-001JP study): a single-arm, open-label, phase 1–2 study. *The Lancet Oncology* **14**, 590–598 (2013).
15. Anderson, A. R., Weaver, A. M., Cummings, P. T. & Quaranta, V. Tumor morphology and phenotypic evolution driven by selective pressure from the microenvironment. *Cell* **127**, 905–915 (2006).
16. Marusyk, A. *et al.* Non-cell autonomous tumor-growth driving supports sub-clonal heterogeneity. *Nature* **514**, 54 (2014).
17. Basanta, D. *et al.* Investigating prostate cancer tumour–stroma interactions: clinical and biological insights from an evolutionary game. *British Journal of Cancer* **106**, 174–181 (2012).
18. Kaznatcheev, A., Scott, J. G. & Basanta, D. Edge effects in game-theoretic dynamics of spatially structured tumours. *Journal of The Royal Society Interface* **12**, 20150154 (2015).
19. Kaznatcheev, A., Vander Velde, R., Scott, J. G. & Basanta, D. Cancer treatment scheduling and dynamic heterogeneity in social dilemmas of tumour acidity and vasculature. *British Journal of Cancer* (2017).
20. Kerr, B., Riley, M. A., Feldman, M. W. & Bohannan, B. J. Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature* **418**, 171–174 (2002).
21. Maddamsetti, R., Lenski, R. E. & Barrick, J. E. Adaptation, clonal interference, and frequency-dependent interactions in a long-term evolution experiment with *Escherichia coli*. *Genetics* **200**, 619–631 (2015).
22. MacLean, R. C. & Gudelj, I. Resource competition and social conflict in experimental populations of yeast. *Nature* **441**, 498 (2006).
23. Gore, J., Youk, H. & Van Oudenaarden, A. Snowdrift game dynamics and facultative cheating in yeast. *Nature* **459**, 253–256 (2009).
24. Li, X.-Y. *et al.* Which games are growing bacterial populations playing? *Journal of The Royal Society Interface* **12**, 20150121 (2015).
25. Archetti, M., Ferraro, D. A. & Christofori, G. Heterogeneity for IGF-II production maintained by public goods dynamics in neuroendocrine pancreatic cancer. *Proceedings of the National Academy of Sciences* **112**, 1833–1838 (2015).

26. Mediavilla-Varela, M., Boateng, K., Noyes, D. & Antonia, S. J. The anti-fibrotic agent pirfenidone synergizes with cisplatin in killing tumor cells and cancer-associated fibroblasts. *BMC Cancer* **16**, 176 (2016).
27. Kaznatcheev, A. Two conceptions of evolutionary games: reductive vs effective. *bioRxiv*, 231993 (2017).
28. Maynard Smith, J & Price, G. R. The Logic of Animal Conflict. *Nature* **246**, 15 (1973).
29. Tomlinson, I. P. & Bodmer, W. F. Modelling the consequences of interactions between tumour cells. *British Journal of Cancer* **75**, 157–160 (1997).
30. Tomlinson, I. P. Game-theory models of interactions between tumour cells. *European Journal of Cancer* **33**, 1495–1500 (1997).
31. Archetti, M. Evolutionary game theory of growth factor production: implications for tumour heterogeneity and resistance to therapies. *British Journal of Cancer* **109**, 1056–1062 (2013).
32. Peña, J., Lehmann, L. & Nöldeke, G. Gains from switching and evolutionary stability in multi-player matrix games. *Journal of Theoretical Biology* **346**, 23–33 (2014).
33. Robinson, D. & Goforth, D. *The topology of the 2x2 games: a new periodic table* (Psychology Press, 2005).
34. Rapoport, A. Exploiter, Leader, Hero, and Martyr: the four archetypes of the 2 × 2 game. *Systems Research and Behavioral Science* **12**, 81–84 (1967).
35. Kaznatcheev, A. Effective games and the confusion over spatial structure. *Proceedings of the National Academy of Sciences*, 201719031 (2018).
36. Gerlee, P. & Altrock, P. M. Extinction rates in tumour public goods games. *Journal of The Royal Society Interface* **14**, 20170342 (2017).
37. Conitzer, V. *The exact computational complexity of evolutionarily stable strategies in International Conference on Web and Internet Economics* (2013), 96–108.
38. Kaznatcheev, A. Computational complexity is an ultimate constraint on evolution, 187682 (2018).
39. Nichol, D. *et al.* Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS computational biology* **11**, e1004493 (2015).
40. Basanta, D., Scott, J. G., Rockne, R., Swanson, K. R. & Anderson, A. R. The role of IDH1 mutated tumour cells in secondary glioblastomas: an evolutionary game theoretical view. *Physical biology* **8**, 015016 (2011).
41. Gatenby, R. A., Gawlinski, E. T., Gmitro, A. F., Kaylor, B. & Gillies, R. J. Acid-mediated tumor invasion: a multidisciplinary study. *Cancer research* **66**, 5216–5223 (2006).
42. Jain, R. K. Normalizing tumor microenvironment to treat cancer: bench to bedside to biomarkers. *Journal of Clinical Oncology* **31**, 2205 (2013).
43. Zhang, J., Cunningham, J. J., Brown, J. S. & Gatenby, R. A. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature communications* **8**, 1816 (2017).
44. Thiel, H. *A rank-invariant method of linear and polynomial regression analysis, I, II, III in Proceedings of Koninklijke Nederlandse Akademie van Wetenschappen A* **53** (1950), 386–392, 521–525, 1397–1412.
45. Sen, P. K. Estimates of the regression coefficient based on Kendall’s tau. *Journal of the American statistical association* **63**, 1379–1389 (1968).
46. Burnham, K. P. & Anderson, D. R. *Model selection and multimodel inference: a practical information-theoretic approach* (Springer Science & Business Media, 2003).