

# Two conceptions of evolutionary games: reductive vs effective

Artem Kaznatcheev<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, University of Oxford, Oxford, UK

<sup>2</sup>Department of Translational Hematology & Oncology Research, Cleveland Clinic, Cleveland, OH, USA

**Evolutionary game theory (EGT) was born from economic game theory through a series of analogies. Given this heuristic genealogy, a number of central objects of the theory (like strategies, players, and games) have not been carefully defined or interpreted. A specific interpretation of these terms becomes important as EGT sees more applications to understanding experiments in microscopic systems typical of oncology and microbiology. In this essay, I provide two interpretations of the central objects of games theory: one that leads to reductive games and the other to effective games. These interpretation are based on the difference between views of fitness as a property of individuals versus fitness as a summary statistic of (sub)populations. Reductive games are typical of theoretical work like agent-based models. But effective games usually correspond more closely to experimental work. However, confusing reductive games for effective games or vice-versa can lead to divergent results, especially in spatially structured populations. As such, I propose that we treat this distinction carefully in future work at the interface of EGT and experiment.**

In my views of game theory, I largely follow Rubinstein (2012): game theory is a set of fables. A collection of heuristic models that helps us structure how we make sense of and communicate about the world. Evolutionary game theory (EGT) was born of classic game theory through a series of analogies. Given this heuristic genealogy of the field, it is usually alright to not worry too much about what exactly terms like strategy, player, or game really mean or refer to. I am usually happy to leave these terms ambiguous so that they can motivate different readers to have different interpretations and subsequently push for different models of different experiments. I think it is essential for heuristic theories to foster this diverse creativity: anything goes.

However, not everyone would agree with Rubinstein and me; some people think that EGT isn't "just" heuristics. EGT is also directly empirically useful for questions in both mathematical oncology and the evolution of microorganisms. Microscopic experimental systems in which EGT has been useful include: *Escherichia coli* (Kerr et al. 2002), yeast (Gore, Youk, and Van Oudenaarden 2009; MacLean and Gudelj 2006), bacterial symbionts of hydra (Li et al. 2015a), breast cancer (Marusyk et al. 2014), pancreatic cancer (Archetti, Ferraro, and Christofori 2015), and lung cancer (Kaznatcheev et al. 2017b). But when we actually start doing experiments like these, it is no longer acceptable to be willy-nilly with fundamental objects of the theory: strategies, players, and games. The biggest culprit is the player. In particular, I think that a lot of confusion stems from saying that "cells are players".

In classical game theory, the concepts of player, strategy, and game are intertwined but relatively straightforward. Players use a rational decision process to select strategies which are then mapped by the rules of the game to payoffs – the utility given to the players. Or, as I say more awkwardly in the first column of table 1: utility is given to a player

based on its strategy which results from a rational decision process carried out by the player. All of this is summarized as the game. But, how does this classic picture translate to evolutionary games?

The easiest place to start is utility: this is almost always interpreted as changes in fitness. We can take this much as uncontroversial. But fitness itself has a complicated ontology and is part of a more general discussion in biological theory that reaches beyond EGT. Two competing interpretations are fitness as a property of an individual organism (reductive fitness leading to reductive games; see section 1) versus fitness as a summary statistic (or emergent property) of a population of subpopulation (effective fitness leading to effective games; see section 2). These two readings are summarized in columns two and three of table 1.

It is helpful to highlight the difference between these two interpretations with an analogy to physics. The setting of statistical mechanics mirrors the fitness for individuals view and defines properties like kinetic energy for individual molecules. Thermodynamics mirrors the effective fitness view and defines properties like temperature for ensembles of molecules. It simply doesn't make sense to talk of the temperature of an individual molecule. Of course, in simple models like the ideal gas, temperature is just mean kinetic energy; this would correspond to an unstructured (inviscid) population in biology where the effective fitness of a (sub)population is just the average reductive fitness of the individuals that make up that subpopulation. But this ideal case seldom happens in nature. In general, there are many ways like recombination systems, spatial subdivision, and admixture in which structured populations depart from panmixis (mean-field). This can be very relevant to how we interpret games, as I discuss in section 3 in the context of spatial structure.

<sup>1</sup>Since we are primarily interested in fitness differences, it is usually subpopulations and not the population of all organisms.

Classic game theory		Evolutionary game theory	
Utility		Fitness	
Player	is given to a	Individual organism	Population <sup>1</sup>
Strategy	based on its	Phenotype (pure strategy)	Distribution of phenotypes (mixed) <sup>2</sup>
Rational decision process	which results from a	Fixed from birth <sup>3</sup>	Replicator dynamics
Game	all summarized as a	Reductive game	Effective game

Table 1: Defining evolutionary game theory by two different analogies to classic game theory.

## 1 Fitness for individuals & reductive games

Since many evolutionary game theorists are computational modelers or think in terms of simulations and agent-based models, we take fitness as a property of an individual organism. In that case, we can define players as the organisms that receive the payoff from local interactions that happen between pairs of organisms (or more for multi-player games). The summary of this local interaction is what I would call the reductive game.

In a the most common EGT setting, what the organisms do in the game is fixed by their genes. Under this reductive interpretation players don't alter their strategies. This makes it easy to present classic vs evolutionary game theory as two extremes on the spectrum of decision making. In classic game theory, players are unbounded rational decision-makers. In evolutionary game theory, players are the most bounded possible: they make no decisions at all; their behavior is genetically fixed. Or, as I say more awkwardly in the second column of table 1: fitness is a property of an individual organism based in its phenotype, which is fixed from birth. All of this is summarized as the reductive game.

The proportion of agents in the population is then updated according to an evolutionary process like replicator dynamics. In particular, in the section I'll discuss how two different implementations of reductive fitness can both give us replicator dynamics. In section 1.1, I will focus on fixed population sizes with fitness as probability to reproduce. In section 1.2, I will consider exponentially growing populations with fitness as number of offspring.

<sup>2</sup>If the fitness functions are linear then interpreting the distribution of phenotypes as a mixed strategy is straightforward. However, fitness functions are not necessarily linear; see Li et al. (2015b) for a clear example. In the case of a fitness function given by a  $m$ -order polynomial, we can still interpret the distribution of phenotypes as a mixed strategy in an  $(m+1)$ -player game.

<sup>3</sup>Fixed from birth is the typical case but it is also natural for the evolutionary game theorists to gradually build up the decision making capacities of the player (McNamara 2013). For example, tag-based models can be seen as taking a minimal step upwards by allowing the focal to decide to cooperate or defect conditioned on the tag of the alter. Although this characterization behavior is simple, it can allow for a rich analysis as a form of minimal cognition (Beer 2003) and we can associate a cognitive cost for this extra decision-making ability (Kaznatcheev 2010). Or push this genotype to behavior map even further (Kaznatcheev, Montrey, and Shultz 2014) by having evolving agents act rationally on their evolved perceptions of the game payoffs and (potentially-biased) estimates of other's probability to cooperate (for this direction, see also Masel (2007)).

## 1.1 Moran: fitness as probability to reproduce

In a Moran process (Moran 1958; Taylor et al. 2004), we imagine that a population is made up of a fixed number  $N$  of individuals. An agent is selected to reproduce in proportion to their game payoff, and their offspring replaces another agent in the population, chosen uniformly at random. This gives us a very clear individual account of fitness as a measure of the probability to place a replicate into the population.

Traulsen, Claussen, and Hauert (2006) wrote down the Fokker-Planck equation for the above Moran process, and then use Ito-calculus to derive a Langevin equation for the evolution of the proportions of each strategy  $x_k$ . The fluctuations in this stochastic equation scale with  $1/\sqrt{N}$  and so vanish in the limit of large  $N$ . This reduces them to a deterministic limit of the replicator equation in Maynard Smith form (Maynard Smith 1982), with the fitness functions as the payoff functions:

$$\dot{x}_k = x_k \frac{w_k - \langle w \rangle}{\langle w \rangle} \quad (1)$$

where  $\langle w \rangle$  is the average fitness and the extra condition of  $\langle w \rangle > 0$  is introduced.<sup>4</sup>

Alternatively, we might be interested in directly getting the Taylor form (Taylor and Jonker 1978) of replicator dynamics:

$$\dot{x}_k = x_k (w_k - \langle w \rangle) \quad (2)$$

Traulsen, Claussen, and Hauert (2006) show how to achieve this, too. Instead of birth-death, they consider an imitation process. Two agents are selected uniformly at random, and individual if the payoff of the first individual is  $w_1$  and the second is  $w_2$  then the first copies the second with probability  $p = \frac{1}{2} + \frac{s}{2} \frac{w_1 - w_2}{\Delta w}$  where  $\Delta w$  is the maximum possible gap in the payoff of two agents in the model. With this version, they get the Taylor form, with the fitness as payoff, again. They are not the first to derive the Taylor form replicator equation from imitation processes. In fact, Schlag (1998) went further by showing that with the proportional imitation rule (only copy those that have higher payoffs, in proportion to how much higher the payoff is), you not only get the Taylor form

<sup>4</sup>In the limit of weak-selection, this form is equivalent to the Taylor form. Even without weak-selection, the two systems of equations differ only by dynamic time rescaling and thus have the same fixed points, orbits, and paths. If we only care about this in our analysis then we can use the equations interchangeably.

replicator equation in a large population limit but also that this local update rule is optimal from the individual agent’s perspective in certain social learning settings.

## 1.2 Exponential: fitness as number of offspring

One of the biggest difference between ecological modeling in micro- vs. macro- organisms is that macro-organisms seldom have the opportunity to undergo exponential growth; they are almost always at carrying capacity. So what if we want to model this ecological difference – the fact that the total population grows or that the cell density in the Petri dish changes?

Populations don’t have to be constant to achieve replicator dynamics. Consider  $m$  types of cells with  $N_1, \dots, N_m$  individuals each, leaving  $w_1, \dots, w_m$  offspring each for the next generation. These offspring numbers  $w_k$  could be functions of various other parameters. The population dynamics are then described by the set of  $m$  differential equations:  $\dot{N}_k = w_k N_k$  for  $1 \leq k \leq m$ . Now, with  $N = N_1 + \dots + N_m$ , we can look at the dynamics of  $x_k = N_k/N$ :

$$\dot{x}_k = \frac{\dot{N}_k}{N} - \frac{N_k \dot{N}}{N^2} \quad (3)$$

$$= \frac{w_k N_k}{N} - \frac{N_k}{N} \frac{\sum_{i=1}^m w_i N_i}{N} = \frac{N_k}{N} (w_k - \sum_{i=1}^m w_i \frac{N_i}{N}) \quad (4)$$

$$= x_k (w_k - \langle w \rangle) \quad (5)$$

which is just the replicator dynamics. If the  $w_k$  are functions of proportions then replicator dynamics can perfectly describe an exponentially growing population.

## 2 Fitness for populations & effective games

But fitness doesn’t have to be defined individually. An alternative perspective is to see fitness as defined only as a summary statistic or emergent property of populations. This is the perspective that makes the most sense when operationalizing fitness in microscopic systems; especially when using typical fitness measures like growth-rates. In that case, the player is the (sub)population that receives the payoff of fitness. The game then becomes the macroscopic coupling between (sub)populations made up of microscopic agents. It is even misleading to call this coupling an “interaction” since that suggest something too active and direct; as I show in section 2.5, the coupling could be as indirect as two populations feeding on a single resource in batch culture. This population-level description is what I call an effective game. Given its roots in operationalization of microscopic systems, the effective games can be measured directly and we recently developed a game assay for this purpose (Kaznatcheev et al. 2017b).

This perspective has some curious consequences. Since the players are populations, the individuals organisms – or behaviorally identical classes of them – are the strategies. The distribution of phenotypes in the population is then interpreted as a mixed strategy. The player is not static but carries out a ‘decision process’ specified by the rules of the evolutionary dynamics. This is usually described by the replicator equation.

Make what you will of the correspondence between replicator dynamics and Bayesian inference, reinforcement learning and other forms of rational decision making (Arora, Hazan, and Kale 2012; Börgers and Sarin 1997). To me, this seems like both a closer correspondence to the aspirations of classical game theory and easier to link to experiment.

To summarize as in the third column of table 1: fitness is a summary statistic of a (sub)population based in its distribution of phenotypes, which is updated according to replicator dynamics. All of this is summarized as the effective game.

What does this mean for “cells are players”? For the reductive game, each individual cell is a player. For the effective game, the population of cells is a player.

### 2.1 Replating: fitness as fold change

Consider the following idealized protocol that is loosely inspired by Archetti, Ferraro, and Christofori (2015) and the *E. coli* Long-term evolution experiment (Lenski et al. 1991; Ribeck and Lenski 2015; Wiser, Ribeck, and Lenski 2013). We will (E1) take a new petri dish or plate; (E2) fill it with a fixed mix of nutritional medium like fetal bovine serum; (E3) put a known number  $N_I = N_1^I + \dots + N_m^I$  of  $m$  different cell on the medium (on the first plate we will also know the proportion of A and B in the mixture); (E4) let them grow for a fixed amount of time  $\Delta t$  which will be on the order of a couple of cell cycles (keeping us in the growth phase); (E5) scrape the cells off the medium and measure the final numbers  $N_F = N_1^F + \dots + N_m^F$ ; and (E6) return to step (E1) while selecting  $N$  cells at random from the ones we got in step (E5) to seed step (E3).

From comparing steps E3 and E5, we can get the experimental population growth rates as:

$$w_k := \frac{N_k^F - N_k^I}{N_k^I \Delta t} \quad (6)$$

this can be rotated into a mapping  $N_I \mapsto N_F$  given by  $N_k^F = N_k^I (1 + w_k \Delta t)$ .

From defining the initial and final population sizes  $N^{\{I,F\}} = \sum_{k=1}^m N_k^{\{I,F\}}$ , we can compare the initial and final proportions of each cell type  $x_k^I = \frac{N_k^I}{N^I}$  and

$$x_k^F = \frac{N_k^F}{N^F} = x_k^I \frac{1 + w_k \Delta t}{1 + \langle w \rangle \Delta t} \quad (7)$$

where  $\langle w \rangle = \sum_{k=1}^m x_k^I w_k$ .

So far we were looking at a discrete process. But we can approximate it with a continuous one. In that case, we can define  $x_k(t) = x_k^I$ ,  $x_k(t + \Delta t) = x_k^F$  and look at the limit as  $\Delta t$  gets very small:

$$\dot{x} = \lim_{\Delta t \rightarrow 0} \frac{x_k(t + \Delta t) - x_k(t)}{\Delta t} \quad (8)$$

$$= \lim_{\Delta t \rightarrow 0} \frac{x_k^I}{\Delta t} \left( \frac{1 + w_k \Delta t}{1 + \langle w \rangle \Delta t} - 1 \right) = \lim_{\Delta t \rightarrow 0} x_k \frac{w_k - \langle w \rangle}{1 + \langle w \rangle \Delta t} \quad (9)$$

$$= x_k (w_k - \langle w \rangle) \quad (10)$$

We recover replicator dynamics as an explicit experimental interpretation for all of our theoretical terms.

We didn’t 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 that we spent so much time debating about are simply buried in the definition of experimental fitness. More importantly, we provided a precise description of how we will measure this quantity. This allows us to hide the pedantics of microdynamics inside of how we measure. Somehow, this might seem unsatisfying in that we just named something but aren’t testing anything. This is where we can begin to twist our experimental knobs a bit to start understanding our system, as I do in section 2.3.

If we are able to peek inside the system more, as – for example, Kaznatcheev et al. (2017b) do with time-lapsed microscopy – then we can also replace fold-change by more specific measurements like inferred growth rates. An advantage is that the goodness-of-fit of exponential models can provide a good estimate of the error associated with these measurements. But the cost is a slightly more specific set of assumptions on the microdynamics of our system. However, since many models can lead to transient exponential growth curves for various microdynamic implementations, these assumptions still don’t have to be as stringent as the definitions in section 1.2.

Does this mean that just by playing with definitions and writing down the replicator equation I have shown that any replating experiment follows replicator dynamics? That would be ridiculous. But we aren’t far. To actually establish this, we need only one more piece: we would need to show that the replatings have the Markov property. In other words, that if the protocol is kept fixed then from one plate to another then the proportion of cells (and maybe a little noise) is all we need to specify the proportion of cells that we will have at the end of the grow phase; there is no dependence of plates that are further back in time. Of course, there are many ways this Markov property might break down.

The easiest way to break the Markov property is if the state-space we consider is insufficient. For example, we assume that the ‘pure’ cultures we start with all have the same strategy. What if the culture is heterogeneous in terms of their game behavior? In that case, for a two-strategy game, we need to treat the game as not between type A and type B but between types A, A’ and B where we don’t have an experimental procedure to differentiate between type A and A’. Alternatively, the extra variable we need to add to the state space might not be an extra cell (sub)type but something like a toxin that survives between replatings.<sup>5</sup>

Of course, expanding the state space might be futile in some cases. The classic examples of this would be cellular or epigenetic memory, certain kinds of phenotype switching, or rapid mutations. In that case, our state space will expand with each replating and it is better to abandon replicator dynamics and adopt some technique that was built for handling

<sup>5</sup>This can be done by binding to the surface of the cells – although experimentalists often ‘wash’ cells to avoid this specific effect – or by being absorbed without harm by one cell type to only be actively or passively pumped out of them later and affect the other cell type – a typical example in oncology might be a chemo-resistant cell pumping out the drug after replating and poisoning it’s non-chemo resistant partners.

historicity. To find this out, though, we would need to first design careful experimental protocols and run our cells and see the long-term behavior breaking the Markov property. Alternatively, if we are lucky and nature cooperates by granting our protocol the Markov property then we will find ourselves in the fortunate position of having a relatively simple operationalist theory of evolutionary games.

## 2.2 Measuring the gain function directly

We don’t necessarily need to measure separate fitness functions for each cell type. It is more important to know the fitness differences, which we can measure directly instead.

Suppose the proportion of cell line A in a mixture is  $p$ , with cell line B making up  $1 - p$ , and the fitnesses of the cells are  $w_A(p)$  and  $w_B(p)$ , respectively. Then the replicator dynamics are given by:

$$\frac{dp}{dt} = p(1-p) \underbrace{(w_A - w_B)}_{\text{gain function}} \quad (11)$$

Our goal in identifying the game is to measure this gain function (see Peña, Lehmann, and Nöldeke (2014) and Kaznatcheev et al. (2017a) for their importance): the increase in growth rate from ‘switching’ from strategy A to B with  $p$  held constant. For this we will use a simple calculus trick: consider the log-odds  $s = \ln \frac{p}{1-p}$ . Then

$$\frac{ds}{dt} = \frac{dp}{dt} \left( \frac{1}{p} + \frac{1}{1-p} \right) = \frac{dp/dt}{p(1-p)} = w_A - w_B \quad (12)$$

By looking at the log-odds of  $p$  instead of just  $p$ , we have ‘factored out’ the logistic growth part of the equation. Now, to measure the gain function, we just have to measure the derivative of  $s$ . Unfortunately, experimentalists do not have a derivative detector in the lab, so we have to approximate the derivative by looking at the change in  $s$  over a short period of time.

In the case of the basic replating experiments you have a natural discretization of time:  $p_{\text{in}}$  can be the proportion of type-A cells at the start of your experiment, and  $p_{\text{out}}$  can be the proportion of type-A cells when you replate. You can then run the experiments for several initial values  $p_{\text{in}}$  and plot your results as  $\Delta s = \ln \frac{p_{\text{out}}}{1-p_{\text{out}}} - \ln \frac{p_{\text{in}}}{1-p_{\text{in}}}$  versus  $p_{\text{in}}$ . This graph is your gain function.

Here we run into the important question of “how short is short enough?” If we run the experiment for too short of a time then the change in  $p$  will be overwhelmed by the measurement noise, but if we run for too long before measuring  $p_{\text{out}}$  then it doesn’t make sense to say we are measuring the derivative. To overcome this, we have to focus on the tangent line of our function as a local linearization. If we have access to high-resolution in time data of the growth during our preliminary experiments then we can plot the resulting growth curve as we would normally: proportion versus time. We then want to take the biggest time-window on which this growth curve is well-approximated by a constant-fitness logistic growth. This logistic growth is the linearization of our function, and as long as it is in approximate agreement, we can say that we are measuring the derivative.



Since we are considering experimental data, it is important to look at the errors associated with our measurements. I don't mean the variance between different runs in different Petri dishes, although that is also important, but the accuracy of the proportions of our initial seeds and the precision of our measurements. There are several ways we could propagate the errors from  $p$  to  $s$ , but as an estimates:

$$\sigma_{\Delta s} = \sigma_p \sqrt{\frac{1}{p_{\text{in}}^2(1-p_{\text{in}})^2} + \frac{1}{p_{\text{out}}^2(1-p_{\text{out}})^2}} \quad (13)$$

So the error is amplified by  $4\sqrt{2}$  near  $p_{\text{in}} = p_{\text{out}} = 0.5$  and the amplification increases as the proportions approach 0 or 1. For example, for  $p_{\text{in}} = p_{\text{out}} = 1/m$  and large  $m$ , it becomes approximate a factor of  $m\sqrt{2}$ . Of course, each experimental set up will serve as a different meter stick, and we will need to do the metrology more carefully for each.

## 2.3 Partitioning fitness with mixed media

For a conditioned medium experiment, a petri-dish is filled with a liquid containing some food (the medium) and then cells of type X from some established culture are introduced onto this medium. They spend some fixed amount of time in that dish, consuming a portion of whatever is in the medium and potentially releasing their own chemicals into the liquid. After a relatively short amount of time – something on the order of the cell cycle – the cells are separated from the medium, usually by running the mixture through a filter that is fine enough to catch cells but let macro-molecules through. This gives us the post-X medium in a flask, which we pour out into a fresh petri-dish and now introduce a known number of cells of type Y onto it. Again, we grow this mixture for a fixed amount of time on the order of a cell cycle and then the cells. This allows us to work with monotypic cultures, which can be helpful if the tools for marking or sorting cells by type are unavailable.

We can extend the conditioned media experiments to mixed media experiments: take a beaker of conditioned medium from cells of type A and another from cells of type B and mix them in some proportion  $p$ . Now we can grow cells of type X on this mixed medium and the ratio of cells that we introduced to the cells we count at the end will give us the fitness of X as a function of the proportion  $p$  of type A medium in the A-B mix. By considering different values of mixed media  $p$  added to monoculture, we can do a mixed-media version of the game assays described in section 2.1 and Kaznatcheev et al. (2017b).

This gives another operationalization of fitness: experimental medium-mediated fitness. When we grow cells X on media conditioned by a proportion  $p$  of type A conditioned media in the A-B mix then we can call the corresponding ratio between initial and final cells as  $w_X^{MM}(p)$ .

Now, if we compare the experimental medium-mediated fitness to the experimental fitness  $w_X(p)$  that we defined above to  $w_X^{MM}(p)$  then for a short enough replating time, we would expect the two numbers to be nearly identical if the interaction between cells is primarily mediated by the medium:  $w_X^{MM}(p) \approx w_X(p)$ .<sup>6</sup>

<sup>6</sup> Although even here, we might want to allow for some low complex-

In general these two functions might be very different. In that case we will have made more than a measurement. We will have discovered that the experimental fitness function partitions into a direct interaction part – that could come from direct predation, very short-term or short-range media-based signals, spatial competition, or many other microdynamical effects – and a medium-mediated part:  $w_X(p) = w_X^{DI}(p) + w_X^{MM}(p)$ .<sup>7</sup> In other words, we started by trying to eliminate all reductionist basis from our work, and yet learned something about the microdynamical properties of our system. Although coming at it from the operationalist perspective forces us to be much more humble and reserved about our conclusions. If we come up with other experimental knobs to fiddle with then we can find further partitions and thus more nuance for the operationalist meaning of the fitness function.

## 2.4 Choosing units of size for populations

So far, we mostly considered populations as numbers of individuals, organisms, or cells. But what is so special about the number of cells? In this subsection, I want to question the reasons to focus on individual cells (at the expense of other choices) as basic atoms.

So, let's look at what we could mean by 'size of population'. The obvious definition is number of cells, and if all we did was *in silico* simulations then it is the definition we could stuck to. Especially for agent-based models, it is very tempting to have cells as your agents and building everything up around them. But consider two populations that have the same number of cells and everything else is equal, *except* ...

1. ... the cells in the first population are metabolically twice as active as cells in the second population. In this case, the more active cells can easily strain their environment more, as they use more resources to fuel themselves. If your limiting resource in the petri dish is growth medium then the more metabolically active cells will consume more of it.<sup>8</sup> With slower metabolic activity, the cell becomes less of an effect not only on its own future, but also on other cells it interacts with – for example, by moving around less or releasing fewer cytokines and thus interacting with fewer other cells. In this case, the more natural set of physical units might be the power consumption in terms of watts-used or ATP-used. This might be more compatible with the metabolic theory of ecology. (Brown et al. 2004)

ity interpolation between the two functions, like proportionality, log-proportionality, or logit-proportionality and accept the approximate equality if the two functions are equal under this comparison instead of strict numerical equality which would correspond to the simplest map: identity.

<sup>7</sup> Again, the best way to partition might be something other than '+' (that corresponds to identity or proportionality) and we might use another low complexity function to partition like  $\log(w_X(p) = \frac{w_X^{DI}(p)w_X^{NN}(p)}{w_X^{DI}(p)w_X^{NN}(p)+1})$ . Whatever choice is made, should be consistent with the choice for approximation in footnote 6 so that the domain to which the different  $w$ s map have a consistent mathematical structure.

<sup>8</sup> The extreme case of this is cells that are completely inactive or maybe even dead. This doesn't come up as much in simulation, since I can just cleanly remove dead agents.

2. ... the cells in the first population are twice as big as cells in the second population.<sup>9</sup> From the point of view of games mediated by things like diffusive factors or cell-cell contact, the bigger cells will have more area to absorb/release factors or to contact other cells. In we are working *in vitro*, larger cells also exhaust the limiting factor of free space quicker than small cells.<sup>10</sup> On top of this, size feeds back into the first point, with larger cells usually doing more things metabolically and in terms of other activity. In this case, the more natural set of physical units might be area-covered.<sup>11</sup>

Without individual cells to ground us, reductionist story telling becomes more difficult; something that can be both a plus or a minus:

On the one hand, it is hard to imagine how 10 watts-used by cancer interacts with 10 watts-used by fibroblasts, instead we are forced to make these measurements experimentally. Since these measurements are almost always at the level of populations, we don't feel a need to make sense of them in reductionist terms of how a single watt-used interacts with another watt-used. You might have noticed that even the word 'interaction' felt awkward in the last two sentences. Watt or area use invite us to recognize the importance of both the size of the other population and the environment more generally. This makes it easier to notice evolutionary games with only indirect interactions.

On the other hand, these more abstract units and operationalist perspective can hinder the imagination, and it can often become more difficult to explain the work or to design new experiments. The alternative units I considered here also obscure the discrete nature of cells, a discreteness that is essential to life and can have significant side-effects on model conclusions (for discussion, see Durrett and Levin (1994) and Shnerb et al. (2000)). Although in the case of replicator dynamics, it is not clear how much this would matter since they already ignore the discrete nature of life through the use of ODEs.

## 2.5 Effective games without interactions

Let  $p$  be the proportion of A-cells in the population, and  $y$  be the nutritional content of the medium – normalized so that the most nutrient rich mix possible has  $y = 1$  and distilled

<sup>9</sup>In the case of cancer, the tumor corresponding to the more voluminous population would be much more burdensome to the patient. In fact, tumor burden is often measured and reported as volume in x-ray or other imaging. The number of cells in the tumor is then inferred from these volumetric measures by assuming (or measuring outside the body) the size of a typical cancer cell.

<sup>10</sup>The importance of area has come up in thinking about prostate cancer metastases to the bone. (Araujo et al. 2014) Osteoclasts and osteoblasts, take up drastically different amounts of area on the bone, and they are only of significant consequence to the model if they are in contact with the bone (else they are not remodeling it). Area-On-Bone becomes the important variable here.

<sup>11</sup>Of course, we could try to express the above in terms of individual cells by converting back and forth between numbers of cells and watts-used or area-covered. Practically, this would mean finding a conversion factor which amounts to a measure of how much power or area a typical cell uses. But in doing so, we have swept some amount of heterogeneity under the rug – after all, each cell takes up a different amount of space or uses a different amount of energy, especially when facing new circumstances like chemotherapy – and it is not clear what useful thing we got in return.

water has  $y = 0$ . For each cell we will have some (analytic) feeding function  $f_A(y)$  and  $f_B(y)$  which translates between the nutritional content of the medium and the organism's fitness, such that:

$$\dot{p} = p(1 - p)(f_A(y) - f_B(y)) \quad (14)$$

It is important to note that  $f_A$  and  $f_B$  are functions of  $y$  and completely independent from  $p$ . At this point, we might be tempted to stop by saying that since experimental step (E2) uses a fixed mix of nutritional medium, we can just treat  $f_A(y) - f_B(y)$  as a constant and thus (excluding the neutral case) we will always have the population converge to all-A or all-B depending on the sign. We will see no (non-trivial) evolutionary game dynamics. This is the standard intuition behind Gaus' exclusion principle: two species cannot co-exist on a single abiotic resource (Hardin 1960).

But stopping here would be a bit disingenuous. The reason that we have to renew the medium on each cycle of the experiment is because it gets consumed between the replating. Further, the rate of consumption might differ between the two cell types. Suppose that each cell type consume the nutrients at rate  $2k_A$  and  $2k_B$ , such that if  $y_{in}$  was our initial level of nutrients in step (E2) then our final level it step (E5) is  $y_{out} = y_{in}(1 - 2k_A x - 2k_B(1 - p))$ .

We can assume that the cell cycle is significantly slower than the metabolic cycle, and start working with our average consumption (calling  $y_{in}$  just  $y$ ):  $\langle y \rangle_p = y(1 - k_B + p(k_B - k_A))$ .

Now, our dynamics become:

$$\dot{p} = p(1 - p)(f_A(\langle y \rangle_p) - f_B(\langle y \rangle_p)) \quad (15)$$

and suddenly our gain function is no longer independent of  $p$ .<sup>12</sup> That is the main slight-of-hand, but let's take the trick to its conclusion.

We will expand the gain function, noting that we picked analytic feeding functions, so  $f_A(y) = \sum_{n=0}^{\infty} A_n y^n$  for some sequence  $\{A_n\}$  and similar for  $f_B(y)$  but with  $\{B_n\}$ :

$$\Gamma(p) = f_A(\langle y \rangle_p) - f_B(\langle y \rangle_p) \quad (16)$$

$$= \sum_{n=0}^{\infty} (A_n - B_n) \langle y \rangle_p^n \quad (17)$$

$$= \sum_{n=0}^{\infty} (A_n - B_n) y^n (1 - k_B + p(k_B - k_A))^n \quad (18)$$

$$= \sum_{n=0}^{\infty} \frac{(A_n - B_n) y^n}{(k_A - k_B)^n} \left( p - \frac{1 - k_B}{k_A - k_B} \right)^n \quad (19)$$

The last line is the power series of some analytic function  $\Gamma(p)$  with coefficients  $\{\Gamma_n = \frac{(A_n - B_n) y^n}{(k_A - k_B)^n}\}$  around the point  $x_0 = \frac{1 - k_B}{k_A - k_B}$ . In particular, given any desired (analytic) gain function  $\Gamma(p)$ , there is some choice of feeding functions  $f_A$  and  $f_B$  (thus, their corresponding coefficients  $\{A_n\}$  and  $\{B_n\}$ ) and  $k_A$  and  $k_B$  such that the population follows identical dynamics. In other words, in this experimental set up, we can recreate any evolutionary game dynamics without having

<sup>12</sup>It might be helpful to switch to our prior notation by noticing that in this case  $w_{\{A,B\}}(p) = f_{\{A,B\}}(\langle y \rangle_p)$

the cells interacting but just based on how they turn nutrition into reproduction. In particular, we can implement games like Hawk-Dove to have co-existence of A and B on a single abiotic resource, thus violating the competitive exclusion principle.

### 3 Spatially structured populations

Spatial structure in game theory is so important that a particular approach to it even has its own sub-field name: evolutionary graph theory (Lieberman, Hauert, and Nowak 2005; Maciejewski and Puleo 2014; Shakarian, Roos, and Johnson 2012; Szabó and Fath 2007). Durrett and Levin (1994) and Shnerb et al. (2000) provide a particularly good demonstration of how much spatial structure and stochasticity can matter as they build from mean-field approaches (of which the inviscid replicator dynamics with which we started is an example) to patch models of discrete individuals to reaction-diffusion equations to full-fledged interacting particle systems. We know that space can promote cooperation (Nowak and May 1992; Ohtsuki et al. 2006), or inhibit it (Hauert and Doebeli 2004), or complicate the whole discussion around it (Killingback and Doebeli 1996). Spatial structure can so drastically change the nature of the reductive game such that the mean-field analysis is completely inapplicable. In section 3.1, I will consider how we might work from the bottom up, by transforming a reductive game through an incorporation of spatial structure. But in section 3.2, I'll consider the reverse direction: extracting the contribution of space from an effective game that would normally swallow-up or abstract the spatial structure into measurement.

#### 3.1 Approximating spatial structure

A typical study of space in evolutionary game theory will start with the reductive game and then simulate that interaction over a model of space to show a surprising difference in dynamics between the spatial model and the mean-field. Exceptional works like that of Ohtsuki and Nowak (2006) (more recently, Nanda and Durrett (2017)), provide a general method for combining a reductive game with spatial structure. Although in its original presentation, Ohtsuki and Nowak (2006) focused on dynamics on  $k$ -regular random graphs, I think it can be useful to frame their work as a general first-order approximation of an arbitrary spatial structure.

First, let me present replicator equation as a zeroth-order approximation of spatial structure. Without knowing anything about spatial structure, the roughest guess we can make of the probability of interacting with another agent is just to say that we sample agents from a distribution given by their proportion in the population – a mean-field approximation. This gives us a utility function for agents of type  $i$  as  $[Gx]_i$  where  $G_{ij}$  is the payoff of an agent of type  $i$  interacting with an agent of type  $j$  and  $x$  is a vector of proportions of agents of each type. From here on in, our hands are tied and the math forces us to replicator dynamics:  $\dot{x}_i = x_i([Gx]_i - x^T Gx)$ .

But the perfect sampling used in the calculation of fitness effectively makes interactions global. To overcome this, we can say that instead of the fitness being the mean-field  $[Gx]_i$ , we instead sample  $M$  interaction partners from the distribu-

tion given by  $x$  and use these local interaction groups for our fitness calculation. This would be our 0.5th-order approximation. In this case, Hilbe (2011) showed that the result is still replicator dynamics (although with a different time-scale that is irrelevant to the functional form) but with a modified payoff matrix  $G' = G + \frac{1}{M-2}(G - G^T)$ . Where  $G^T$  is the matrix transpose of  $G$ . This is a great way to reintroduce some local effects, but the groups of  $M$  agents are constantly re-sampled and fitness-competition still happens at a global level; in other words, there is no spatial structure. Hence the 0.5 and not 1.

To get things completely localized, we will assume a fixed population size, and make our replication procedure more explicit to make a first-order approximation of spatial structure. Since the population size is fixed, we can only get a new agent if an old one dies, this gives us a great way to localize. Once a focal agent dies, there is some neighborhood of the focal agent with  $k$  agents that compete for the focal spot. Now, we can have some extra information, instead of just keeping track of the proportion of agents of type  $i$  given by  $x_i$ , we can keep track of neighbors. More explicitly, we will use pair-approximation (Matsuda et al. 1987; Van Baalen 2000) to keep track of proportion  $x_{ij}$ : the probability of seeing an agent of type  $i$  in the neighborhood of an agent of type  $j$ .

This tells us who is competing for for the vacated spot, but it doesn't let us calculate their fitness because we would need to know the probability  $x_{ijk}$  of seeing an agent of type  $i$  near an agent of type  $j$  and  $k$ . To update that probability, we would need to know more long range effects like  $x_{ijkl}$ , etc. Hence, for a first-order approximation, we truncate the series here and approximate the further effects by saying that  $x_{ijk} = x_{ij}$ .<sup>13</sup> Working under these assumptions, Ohtsuki and Nowak (2006) showed that we still get replicator dynamics but with a modified payoff matrix  $G''$  given by:

$$\text{ON}_k(G) = G + \underbrace{\frac{1}{k-2}(\bar{\Delta}\bar{1}^T - \bar{1}\bar{\Delta}^T)}_{\text{assortativeness from local dispersal}} + \underbrace{\frac{1}{(k+1)(k-2)}(G - G^T)}_{\text{finite sampling from death-birth updating}} \quad (20)$$

where vector  $\Delta$  which is the diagonal of the game matrix  $G$ , i.e.  $\Delta_i = G_{ii}$  and  $\bar{1}$  is the all ones vector; thus  $\bar{1}\bar{\Delta}^T$  ( $\bar{\Delta}\bar{1}^T$ ) is a matrix with diagonal elements repeated to fill each row (column). Note that we have  $k$  agents competing for a spot, and each one samples  $k-1$  other agents (since one spot they are neighboring was just vacated) so  $M = k(k-1)$ . Thus, the last term is Hilbe (2011) finite sampling effect and the first term is the spatial structure.

<sup>13</sup>Since we assumed that the neighbors of the perished agent  $i$  are drawn from the same sort of distribution as the neighbors' neighbors, we have ignored extra correlations that might arise from looking out to distance two or more, hence the first order nature of the approximation. Of course, like any first-order approximation, if higher-orders are important then it will not agree with experimental data. However, if we just look at data without a first-order theory then we wouldn't even know that higher-order terms are important. Thus, the first-order approximation is always a good first step; if empirical results contradict it then at least we know where to look, second-order and higher correlations in the distributions of neighbors.

To restate: the effective game  $G'' = \text{ON}_k(G)$  then has the same mean-field replicator dynamics as the reductive game  $G$  carried out on the spatial structure. This allows us to use our tools of evolutionary game theory to analyze the transformed game and thus learn something about the system implemented by the reductive game and spatial structure. We can think of this approach as bottom up: start with the reductive game, find the corresponding effective game (or other description of macroscopic population-level dynamics) and then use these as a prediction to compare against observed phenomena.

As an example of the typical study of space in EGT, consider Kaznatcheev, Scott, and Basanta (2015). Here, we started with the reductive Go-vs-Grow game between invasive (INV; Go) and autonomously growing (AG; Grow) cells (Basanta, Hatzikirou, and Deutsch 2008). We wanted to know how the game was effected by different spatial structure in the bulk versus a static boundry of a tumour, so we transformed it according to the Ohtsuki and Nowak (2006) transform:

$$\begin{array}{cc} \text{INV} & \text{AG} \\ \text{INV} & \left( \begin{array}{cc} \frac{1}{2}b + \frac{1}{2}(b-c) & b-c \\ b & \frac{1}{2}b \end{array} \right) \\ \text{AG} & \end{array} \quad (21)$$

$$\begin{array}{c} \downarrow \\ \text{ON}_k(\circ) \\ \downarrow \\ \left( \begin{array}{cc} \frac{1}{2}b + \frac{1}{2}(b-c) & b \frac{2k-3}{2(k-2)} - c \frac{2k^2-k-1}{2(M-2)} \\ b \frac{2k-5}{2(k-2)} + c \frac{k+3}{2(M-2)} & \frac{1}{2}b \end{array} \right) \end{array} \quad (22)$$

where  $M = k(k-1)$ .

Unlike the reductive game in equation 21, the effective game in equation 22 allows for a fully invasive tumour when  $\frac{k+1}{k^2+1} \geq \frac{c}{b}$ . From this, we drew the surprising conclusion that a tumour can have much more invasive phenotypes at the boundary than the bulk.

### 3.2 Operationalizing spatial structure

But when we apply the typical pipeline: how do we know that the local interactions of the reductive game are the right ones to start with? In the example above, how do we know that the reductive Go-vs-Grow game is given by equation 21? For macroscopic systems like human or other large animals, we might be able to directly observe or maybe even design the reductive game. In microscopic systems like cancer, however, we tend to guess these games from intuitions acquired by looking at population level experiments. Unfortunately, these experiments seldom explicitly account for the effect of their spatial structure. Hence, they are actually intuitions about the effective game that we then feed into our models as the reductive game. This is the common confusion about spatial structure in microscopic systems. We are taking a game from a top-level view, feeding it into the bottom level, getting a different result at the top-level and then publishing that surprising conclusion.

This is backwards. At best, it is just telling us that our intuitions about the game were wrong – since correct intuitions about the reductive game should yield the observed effective games. At worst, this is a type-errors and thus incoherent: we are feeding in an effective game where we should be putting a reductive game.

Instead, I propose starting with our measurement of effective games and pushing it down with an operationalization of spatial structure. We can achieve this, by giving an experimental definition of the local environment of cells. This local perspective might be very different from the perspective that an experimenter has of the system as a whole.

Let's suppose we are considering a system with two possible strategies A and B then we can define the following functions:

- Let  $\rho_A : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$  be the distribution over number of type A and type B partners that a player of strategy A encounters; i.e.  $\rho(k_A, k_B)$  is the probability of encountering  $k_A$  many players of type A and  $k_B$  many players of type B during the timescale relevant to the calculation of local fitness. This is similar to  $p_{ij}$  in our discussion of the Ohtsuki-Nowak transform.
- Let  $\hat{\omega}_A : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}$  be the local fitness function for an agent of type A; i.e.  $\hat{\omega}_A(k_A, k_B)$  is the local fitness of an agent of type A that encountered  $k_A$  many players of type A and  $k_B$  many players of type B during the timescale relevant to the calculation of local fitness.

Define  $\rho_B, \hat{\omega}_B$  analogously for strategy B. This allows us to write down the replicator dynamics at the level of the whole population as  $\dot{p} = p(1-p)\Gamma(p)$  where the gain function  $\Gamma(p)$  is given by:

$$\Gamma(p) = \mathbb{E}_{(k_A, k_B) \sim \rho_A} [\hat{\omega}_A(k_A, k_B)] - \mathbb{E}_{(k_A, k_B) \sim \rho_B} [\hat{\omega}_B(k_A, k_B)] \quad (23)$$

Since  $\rho_{\{A, B\}}$  depends on  $p$ , let us name this mapping  $S(p; \vec{s}) \mapsto (\rho_A, \rho_B)$  and introduce an extra state vector  $\vec{s}$  which might also change with time according to some general relationship:  $\dot{\vec{s}} = T(\vec{s}; p)$ . This relationship between  $p, \vec{s}$  and  $\rho$  is meant to capture the functional role of space or any other discrepancy between the local and global perspectives. The hope is that in practical settings,  $\vec{s}$  is simple or non-existent or the dynamics of  $\vec{s}$  can be decoupled from the dynamics of  $p$  by something like a separation of timescales, like what happens from using weak selection in the derivation of Ohtsuki and Nowak (2006).

Now, suppose that we were able to find a system where the equations for  $\vec{s}$  can be decoupled from  $p$ . In that case, this formalism is that we can start to operationalize space. The first step to measure the gain function  $\Gamma(p)$  as described in section 2.2. But measuring  $S$  is more difficult because it is encoding much more information about the microdynamical structure. A good first guess might be to take something like a structured core biopsy and define an interaction radius  $r$ . Then go through taking each cell as a focal agent and count how many cells of type A and B are within distance  $r$  of the focal agent. The result is an empirical estimate for  $\rho_A, \rho_B$ . Repeat for different initial  $p$  to get as many points of function  $S$  as desired. Clearly, if  $r$  is taken as the diameter



of the slide then  $S$  will be basically an identity map since  $p = \frac{\rho_A}{\rho_A + \rho_B}$  in that case. At the other extreme, taking  $r$  as less than a cell radius will make  $S$  into a constant map with  $\rho_A = \rho_B = 1$ . For intermittent values of  $r$ , we will potentially have a variance in different local densities for each focal agent, and picking a good  $r$  will depend on trade-offs between the level of error introduced by this variance versus the level of error that's introduced in the propagation from raw gain function to game.

Unfortunately, just like with other effective concepts, this operationalized of spatial structure might not always have clear microdynamic interpretations. However, it does allow us to go one step closer to experimentally understanding the effects of spatial structure on populations without confusing effective and reductive games.

## References

- Araujo, A., L. M. Cook, C. C. Lynch, and D. Basanta (2014). "An integrated computational model of the bone microenvironment in bone-metastatic prostate cancer". In: *Cancer research* 74.9, pp. 2391–2401.
- Archetti, M., D. A. Ferraro, and G. Christofori (2015). "Heterogeneity for IGF-II production maintained by public goods dynamics in neuroendocrine pancreatic cancer". In: *Proceedings of the National Academy of Sciences* 112.6, pp. 1833–1838.
- Arora, S., E. Hazan, and S. Kale (2012). "The Multiplicative Weights Update Method: a Meta-Algorithm and Applications." In: *Theory of Computing* 8.1, pp. 121–164.
- Basanta, D., H. Hatzikirou, and A. Deutsch (2008). "Studying the emergence of invasiveness in tumours using game theory". In: *The European Physical Journal B* 63.3, pp. 393–397.
- Beer, R. D. (2003). "The dynamics of active categorical perception in an evolved model agent". In: *Adaptive Behavior* 11.4, pp. 209–243.
- Börgers, T. and R. Sarin (1997). "Learning through reinforcement and replicator dynamics". In: *Journal of Economic Theory* 77.1, pp. 1–14.
- Brown, J. H., J. F. Gillooly, A. P. Allen, V. M. Savage, and G. B. West (2004). "Toward a metabolic theory of ecology". In: *Ecology* 85.7, pp. 1771–1789.
- Durrett, R. and S. Levin (1994). "The importance of being discrete (and spatial)". In: *Theoretical population biology* 46.3, pp. 363–394.
- Gore, J., H. Youk, and A. Van Oudenaarden (2009). "Snowdrift game dynamics and facultative cheating in yeast". In: *Nature* 459.7244, pp. 253–256.
- Hardin, G. (1960). "The competitive exclusion principle". In: *science* 131.3409, pp. 1292–1297.
- Hauert, C. and M. Doebeli (2004). "Spatial structure often inhibits the evolution of cooperation in the snowdrift game". In: *Nature* 428.6983, pp. 643–646.
- Hilbe, C. (2011). "Local replicator dynamics: a simple link between deterministic and stochastic models of evolutionary game theory". In: *Bulletin of mathematical biology* 73.9, pp. 2068–2087.
- Kaznatcheev, A. (2010). "The cognitive cost of ethnocentrism". In: *Proceedings of the Cognitive Science Society*. Vol. 32. 32.
- Kaznatcheev, A., M. Montrey, and T. R. Shultz (2014). "Evolving useful delusions: Subjectively rational selfishness leads to objectively irrational cooperation". In: *Proceedings of the 36th annual conference of the cognitive science society*.
- Kaznatcheev, A., J. G. Scott, and D. Basanta (2015). "Edge effects in game-theoretic dynamics of spatially structured tumours". In: *Journal of The Royal Society Interface* 12.108, p. 20150154.
- Kaznatcheev, A., R. Vander Velde, J. G. Scott, and D. Basanta (2017a). "Cancer treatment scheduling and dynamic heterogeneity in social dilemmas of tumour acidity and vasculature". In: *British Journal of Cancer*.
- Kaznatcheev, A., J. Peacock, D. Basanta, A. Marusyk, and J. G. Scott (2017b). "Fibroblasts and alectinib switch the evolutionary games that non-small cell lung cancer plays". In: *bioRxiv*, p. 179259.
- Kerr, B., M. A. Riley, M. W. Feldman, and B. J. Bohannan (2002). "Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors". In: *Nature* 418.6894, pp. 171–174.
- Killingback, T. and M. Doebeli (1996). "Spatial evolutionary game theory: Hawks and Doves revisited". In: *Proceedings of the Royal Society of London B: Biological Sciences* 263.1374, pp. 1135–1144.
- Lenski, R. E., M. R. Rose, S. C. Simpson, and S. C. Tadler (1991). "Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations". In: *The American Naturalist* 138.6, pp. 1315–1341.
- Li, X.-Y., C. Pietschke, S. Fraune, P. M. Altrock, T. C. Bosch, and A. Traulsen (2015a). "Which games are growing bacterial populations playing?" In: *Journal of The Royal Society Interface* 12.108, p. 20150121.
- (2015b). "Which games are growing bacterial populations playing?" In: *Journal of The Royal Society Interface* 12.108, p. 20150121.
- Lieberman, E., C. Hauert, and M. A. Nowak (2005). "Evolutionary dynamics on graphs". In: *Nature* 433.7023, pp. 312–316.
- Maciejewski, W. and G. J. Puleo (2014). "Environmental evolutionary graph theory". In: *Journal of theoretical biology* 360, pp. 117–128.
- MacLean, R. C. and I. Gudelj (2006). "Resource competition and social conflict in experimental populations of yeast". In: *Nature* 441.7092, p. 498.
- Marusyk, A., D. P. Tabassum, P. M. Altrock, V. Almendro, F. Michor, and K. Polyak (2014). "Non-cell autonomous tumor-growth driving supports sub-clonal heterogeneity". In: *Nature* 514.7520, p. 54.
- Masel, J. (2007). "A Bayesian model of quasi-magical thinking can explain observed cooperation in the public good game". In: *Journal of Economic Behavior & Organization* 64.2, pp. 216–231.
- Matsuda, H., N. Tamachi, A. Sasaki, and N. Ogita (1987). "A lattice model for population biology". In: *Mathematical top-*

- ics in population biology, morphogenesis and neurosciences*. Springer, pp. 154–161.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- McNamara, J. (2013). “Towards a richer evolutionary game theory”. In: *J. R. Soc., Interface* 10 (88).
- Moran, P. A. P. (1958). “Random processes in genetics”. In: *Proceedings of the Cambridge Philosophical Society*. Vol. 54, p. 60.
- Nanda, M. and R. Durrett (2017). “Spatial evolutionary games with weak selection”. In: *Proceedings of the National Academy of Sciences*, p. 201620852.
- Nowak, M. A. and R. M. May (1992). “Evolutionary games and spatial chaos”. In: *Nature* 359.6398, pp. 826–829.
- Ohtsuki, H. and M. A. Nowak (2006). “The replicator equation on graphs”. In: *Journal of theoretical biology* 243.1, pp. 86–97.
- Ohtsuki, H., C. Hauert, E. Lieberman, and M. A. Nowak (2006). “A simple rule for the evolution of cooperation on graphs and social networks”. In: *Nature* 441.7092, pp. 502–505.
- Peña, J., L. Lehmann, and G. Nöldeke (2014). “Gains from switching and evolutionary stability in multi-player matrix games”. In: *Journal of Theoretical Biology* 346, pp. 23–33.
- Ribeck, N. and R. E. Lenski (2015). “Modeling and quantifying frequency-dependent fitness in microbial populations with cross-feeding interactions”. In: *Evolution* 69.5, pp. 1313–1320.
- Rubinstein, A. (2012). *Economic fables*. Open book publishers.
- Schlag, K. H. (1998). “Why imitate, and if so, how?: A boundedly rational approach to multi-armed bandits”. In: *Journal of economic theory* 78.1, pp. 130–156.
- Shakarian, P., P. Roos, and A. Johnson (2012). “A review of evolutionary graph theory with applications to game theory”. In: *Biosystems* 107.2, pp. 66–80.
- Shnerb, N. M., Y. Louzoun, E. Bettelheim, and S. Solomon (2000). “The importance of being discrete: Life always wins on the surface”. In: *Proceedings of the National Academy of Sciences* 97.19, pp. 10322–10324.
- Szabó, G. and G. Fath (2007). “Evolutionary games on graphs”. In: *Physics reports* 446.4, pp. 97–216.
- Taylor, C., D. Fudenberg, A. Sasaki, and M. A. Nowak (2004). “Evolutionary game dynamics in finite populations”. In: *Bulletin of Mathematical Biology* 66.6, pp. 1621–1644.
- Taylor, P. D. and L. B. Jonker (1978). “Evolutionary stable strategies and game dynamics”. In: *Mathematical Biosciences* 40.1, pp. 145–156.
- Traulsen, A., J. C. Claussen, and C. Hauert (2006). “Coevolutionary dynamics in large, but finite populations”. In: *Physical Review E* 74.1, p. 011901.
- Van Baalen, M. (2000). “Pair approximations for different spatial geometries”. In: *The geometry of ecological interactions: simplifying spatial complexity* 742, pp. 359–387.
- Wiser, M. J., N. Ribeck, and R. E. Lenski (2013). “Long-term dynamics of adaptation in asexual populations”. In: *Science* 342.6164, pp. 1364–1367.