Plasticity in evolutionary games

Slimane Dridi

Institute for Advanced Study in Toulouse, University of Toulouse 1 Capitole, France
E-mail: `slimane.dridi@gmail.com`
Website: `slimaned.github.io`

1 Abstract

The ability to respond appropriately to environmental cues is fundamental to the success of all forms of life. However, previous theoretical studies of the evolution of plasticity make such diverse assumptions that the conditions under which plasticity can emerge in evolving populations are unclear when fitness is frequency-dependent. We study the effect of adding plastic types to symmetric evolutionary games. Since frequency dependence induces an evolutionary change in the environment of players, one might expect that plastic individuals who can adapt their phenotypes to the environment would have a fitness advantage over simpler purely genetically determined phenotypes. In our model, plastic individuals can detect the type of their opponent before an interaction and condition their action on it. Even though it might appear to be an outstanding advantage, such an ability cannot guarantee invasion of pure types in all games as long as plasticity is costly. We classify games according to whether plasticity can or cannot invade a population of pure types and become the evolutionarily stable strategy. In games where the standard replicator dynamics converge to a pure state, costly plasticity cannot invade an equilibrium population. One can guarantee co-existence of plasticity with pure types in games with mixed equilibria, mirroring the result of frequency-independent models that plasticity thrives when there is variance in fitness. Costly plasticity can however be locally stable in many games, but the way to achieve stability is not to play the best response to any possible encountered type. Rather, part of the stability success of plastic types is based on establishing Pareto-efficiency as residents. Zero-sum games always allow for the global stability of plastic types. This study offers a more principled way of thinking about the evolutionary emergence of plasticity in social scenarios and helps demonstrate that such an emergence is strongly dependent on the type of game individuals are faced with.

Keywords: reaction norms; frequency dependence; natural selection; evolutionary games; variance in fitness.

1 Introduction

Most phenotypes are plastic; indeed, traits are in general only expressed in reaction to an environmental cue. The immune system is programmed to detect pathogens, quorum sensing in bacteria conditions gene expression on cell density, plant growth depends on external light via photosynthesis (Diggle et al., 2007; Bergstrom and Dugatkin, 2016). Being plastic thus seemingly provides a selective advantage to biological organisms. The conditions favoring plasticity evolution have been studied in detail in evolutionary biology, and a general conclusion that can be drawn from classical theory is that a varying environment is required for plastic traits to provide a fitness benefit (Gomulkiewicz and Kirkpatrick, 1992; Gavrilets and Scheiner, 1993a,b). However, most of these classical models were developed under the assumption that fitness depends on the environment but not on the phenotype of other organisms in the population, i.e., in the absence of frequency dependence. It is an open question under which conditions plasticity emerges in frequency-dependent scenarios. In particular, it is unclear how the idea that varying environments favor plasticity evolution can be generalized to cases with frequency-dependent selection.

Because the key to plasticity evolution is variance in fitness, one actually does not require that the rules of the game change over evolutionary time (Dridi and Lehmann, 2014; Weitz et al., 2016; Hilbe et al., 2018) for frequency dependence to generate variance in fitness (Graves and Weinreich, 2017). Indeed, in frequency-dependent selection, other types make up part of the environment, and as long as there is some level of mixing between types (e.g., there is not full assortment between individuals of similar phenotype), any individual is going to interact with different types and hence is going to experience variance in fitness. These observations suggest that even in the absence of changes in the structure of the evolutionary game determining individuals' fitness, there is potentially a selection pressure in favor of plasticity. In order to further examine this question, we must determine what it means to have a plastic phenotype in a game theoretical context.

Plasticity in social evolution has been assumed – without investigating its evolutionary emergence – in order to address many different questions, from studies of cooperation in the iterated prisoner's dilemma (Axelrod and Hamilton, 1981; Nowak, 2006; Press and Dyson, 2012; Adami and Hintze, 2013; Stewart and Plotkin, 2013, 2016) to investigations of signalling (Zahavi, 1975; Grafen, 1990; Maynard-Smith and Harper, 2004), and learning (McElreath and Boyd, 2007; Dridi and Lehmann, 2015, 2016), so we lack a unified view on plasticity in social evolution. These studies have provided interesting insights into our understanding of how natural selection shapes complex strategies for repeated interactions. However, complex strategies that allow an individual to condition behavior on the environment, on opponents' behavior, or on memory of past events are only possible if, in the first place, individuals possess the ability to express plastic social phenotypes. Previous work seemingly does not address the question of the evolutionary emergence of social plasticity, so it remains unclear what is the main advantage of plasticity in social evolution.

In this paper, we adopt one of the simplest implementations of plasticity that we can think of, namely we assume that plastic individuals can detect the type of their opponent before an interaction takes place, and can condition their action on the detected type. At first, this might seem to be a considerable advantage to plastic types, but we will see below that even the smallest cost impedes plasticity to dominate other strategies in all circumstances. Another potential concern is that the genetic and molecular machinery necessary to perform a combination of strategy detection and appropriate response might be complex to evolve even for the most basic forms of plasticity. A perfect response to existing types is unlikely to emerge out of a background of pure genetic determination. Indeed, previous research suggests that such perfect responses might be very difficult to evolve (McNamara et al., 1999; André and Day, 2007). For this reason, we allow our plastic types to adopt any possible response to their opponents, in contrast to a previous work on the topic (Banerjee and Weibull, 1995), where these authors have assumed that plastic types always play a best response to pure types. Here, we rather study how the evolutionary success of plastic types depends on their response to pure types. Moreover, the perceptual system allowing one to infer others' strategies might at first also be defective if it evolves from a state where there was no perceptual system of this kind in the ancestral population. We capture such imperfections by imposing a fitness cost on the expression of the plastic phenotype.

In the following, we define a model that makes our assumptions more precise, and analyze the evolutionary performance of plastic types when pitted against individuals who can only express a fixed pure strategy in a normal-form game. We start by giving a special focus to 2×2 games and analyze the replicator dynamics for four standard games of cooperation: the Prisoner's dilemma, the Stag-hunt game, the Snowdrift game, and a Mutualism game. We then provide classes of games where plastic types can or cannot be globally or locally stable under the standard replicator dynamics.

2 Model

We consider the standard model of evolutionary game theory (Taylor and Jonker, 1978) of a well-mixed population in which players are matched randomly in pairs to play a 2-player n-action game, which is called throughout the pure-type game. We denote the set of actions by $\mathcal{A}$. The population consists of $n+1$ types: the first $n = |\mathcal{A}|$ types are called pure types, and correspond to each pure action while the $(n+1)$-th type can detect others' type before choosing an action. This plastic type, denoted p, has a strategy described by $\mathbf{z} = (\mathbf{z}_1, \ldots, \mathbf{z}_n, \mathbf{z}_{n+1})$, where $\mathbf{z}_i = (z_{i1}, \ldots, z_{in}) \in \Delta\mathcal{A}$ is the mixed strategy adopted by p when faced with type $i \in \mathcal{A} \cup \{p\}$. The symbol $\Delta\mathcal{A}$ denotes the n-dimensional simplex, such that z_{ik} is the probability that the plastic type plays strategy k against type i. With these definitions, the Cartesian product $\prod_{i=1}^{n+1} \Delta\mathcal{A} = (\Delta\mathcal{A})^{n+1}$ is the strategy set of the plastic type. The mixed strategy $\mathbf{z}_{n+1} = \mathbf{z}_p \in \Delta\mathcal{A}$ is the strategy adopted by a plastic individual when faced with another plastic individual. The payoff of type $i \in \mathcal{A}$ against type $j \in \mathcal{A}$ is denoted $\pi(i, j)$, with the convention that the payoff goes to the individual whose strategy appears in the first position in parentheses. For interactions involving the plastic type p, we generally write the payoff $\pi(p, i) = \pi(\mathbf{z}_i, i)$ to emphasize the dependence of the payoff on the mixed strategy $\mathbf{z}_i$ of type p against i. We identify the payoffs $\pi(\mathbf{z}_i, i)$ and $\pi(i, \mathbf{z}_i)$ with the expected payoff generated by the mixed strategy of the plastic type, that is

$$\pi(\mathbf{z}_i, i) = \sum_{j \in \mathcal{A}} z_{ij} \pi(j, i) \qquad \text{and} \qquad \pi(i, \mathbf{z}_i) = \sum_{j \in \mathcal{A}} z_{ij} \pi(i, j). \tag{1}$$

When two plastic individuals meet, they both use their strategy $\mathbf{z}_p$ against a plastic type, so their payoff reads

$$\pi(\mathbf{z}_p, \mathbf{z}_p) = \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{A}} z_{pi} z_{pj} \pi(i, j), \tag{2}$$

where we assumed that both plastic types adopt the same responsive strategy. We do not consider selection on the plastic response itself, $\mathbf{z}$, in this paper, so all plastic types will always have the same responsive strategy. We are interested in tracking the vector of frequencies of the types $\mathbf{x} = (x_1, \ldots, x_n, x_{n+1}) \in \Delta^{n+1}$ such that $\sum_{i=1}^{n+1} x_i = 1$. We write $w_i(\mathbf{x}), i = 1, \ldots, n$, for the fitness of type i when the population is in state $\mathbf{x}$ which is calculated as the average payoff at state $\mathbf{x}$, or

$$w_i(\mathbf{x}) = \sum_{j=1}^{n+1} x_j \pi(i, j). \tag{3}$$

We further assume that type p pays a cost $k > 0$ for expressing a plastic response so that its fitness reads

$$w_p(\mathbf{x}) = \sum_{j=1}^{n+1} x_j \pi(p, j) - k. \tag{4}$$

The frequency of any type i evolves according to the replicator dynamics, which are given by the differential equations

$$\dot{x}_i = x_i \left(w_i(\mathbf{x}) - \bar{w}(\mathbf{x}) \right), \qquad i \in \mathcal{A} \cup \{p\}, \tag{5}$$

where

$$\bar{w}(\mathbf{x}) = \sum_{i=1}^{n+1} x_i w_i(\mathbf{x}) \tag{6}$$

is the average fitness in the population at state $\mathbf{x} \in \Delta^{n+1}$. A further element of notation is that we denote by $\phi(\cdot, \mathbf{x}_0) : \mathbb{R}^+ \to \Delta^{n+1}$ the global solution trajectory (the *flow*) to eq. 5 that passes through $\mathbf{x}_0$, which means that $\phi(t, \mathbf{x}_0)$ is the vector of frequencies at time t given that the vector frequencies passes through $\mathbf{x}_0$. We denote by $\phi_i(t, \mathbf{x}_0)$ the i^{th} element of $\phi(t, \mathbf{x}_0)$, with $i \in \mathcal{A} \cup \{p\}$.

3   2×2 games of cooperation

We begin by studying simple examples of games in order to form intuition about the evolutionary success of plasticity in social interactions. More specifically, we consider four 2×2 games of cooperation traditionally studied in evolutionary biology: the Prisoner's Dilemma game (PD), the Stag-hunt game (SH), the Snowdrift game (SD), and the Mutualism game (MG). We label the two possible actions as C (for Cooperation) and D (for Defection), and this generates a plastic type with a three-dimensional trait $z = (z_C, z_D, z_p) \in [0,1]^3$, where z_i is the probability to cooperate when facing type $i \in \{C, D, p\}$.

We perform a sensitivity analysis with respect to the parameters z_C, z_D, z_p. For each value of these parameters, we determine the fate of plasticity when competing in populations consisting of pure types. We will use the standard notation for the payoffs of 2×2 social dilemma games, often used in the biological literature, i.e., $\pi(C, C) = R$, $\pi(C, D) = S$, $\pi(D, C) = T$, and $\pi(D, D) = P$.

3.1   Prisoner's dilemma game

In this section we define the various conditions for plasticity to invade pure types or to resist invasion, using the generic payoffs R, S, T, and P, so the inequalities below are valid for all games of cooperation studied in this article. However, in this particular section we discuss the validity of these inequalities for the Prisoner's Dilemma game. The payoffs of the Prisoner's dilemma satisfy $T > R > P > S$ and $R \geq (T + S)/2$. In this game, there is one dominant action, D. The outcome (D, D) is the only stable equilibrium of the pure-type game. It follows that the plastic type cannot invade a population of defectors, which means that the inequality

$$\pi(i, i) < \pi(z_i, i) - k \tag{7}$$

can never be satisfied for $i = D$. But can plasticity invade cooperators, and is it immune against invasion by pure types? As to the first question, setting $i = C$ in eq. 7 we find that any strategy such that $z_C < (R - T + k)/(R - T)$ would guarantee that plastic mutants invade a monomorphic population of cooperators (i.e., the plastic type should defect with positive probability against cooperators). As to the second question regarding the ability of plastic types to resist invasion by pure types, we have to solve the inequalities

$$\pi(i, z_i) < \pi(z_p, z_p) - k \tag{8}$$

for z_D and z_p, replacing $i = C$ and $i = D$. In Fig. 1A–B, we show the regions of the space defined by combinations of (z_D, z_p) and (z_C, z_p) that satisfy eq. 8, which constitute the set of plastic types that are immune against the invasion by defectors and cooperators respectively in the Prisoner's dilemma game. Note that eq. 8 implies that the plastic type can be immune against the invasion by defectors only if $z_p > 0$ (this can be shown by setting $i = D$ and $z_p = 0$ in eq. 8, which then cannot be satisfied because $T > P$).

What are the conditions on the payoffs of the game, (R, S, T, P), that allow these inequalities to be satisfied? The feasible payoffs of a plastic type as a resident (the right-hand side of eq. 8) are on the line $x = y$ within the convex hull of feasible payoffs. This means that the payoff $\pi(z_p, z_p)$ satisfies

$$\min\{R, P, \frac{T + S}{2}\} \leq \pi(z_p, z_p) \leq \max\{R, P, \frac{T + S}{2}\}. \tag{9}$$

Combining the above inequality with eq. 8, we then see that plasticity cannot be locally stable against mutant cooperators if

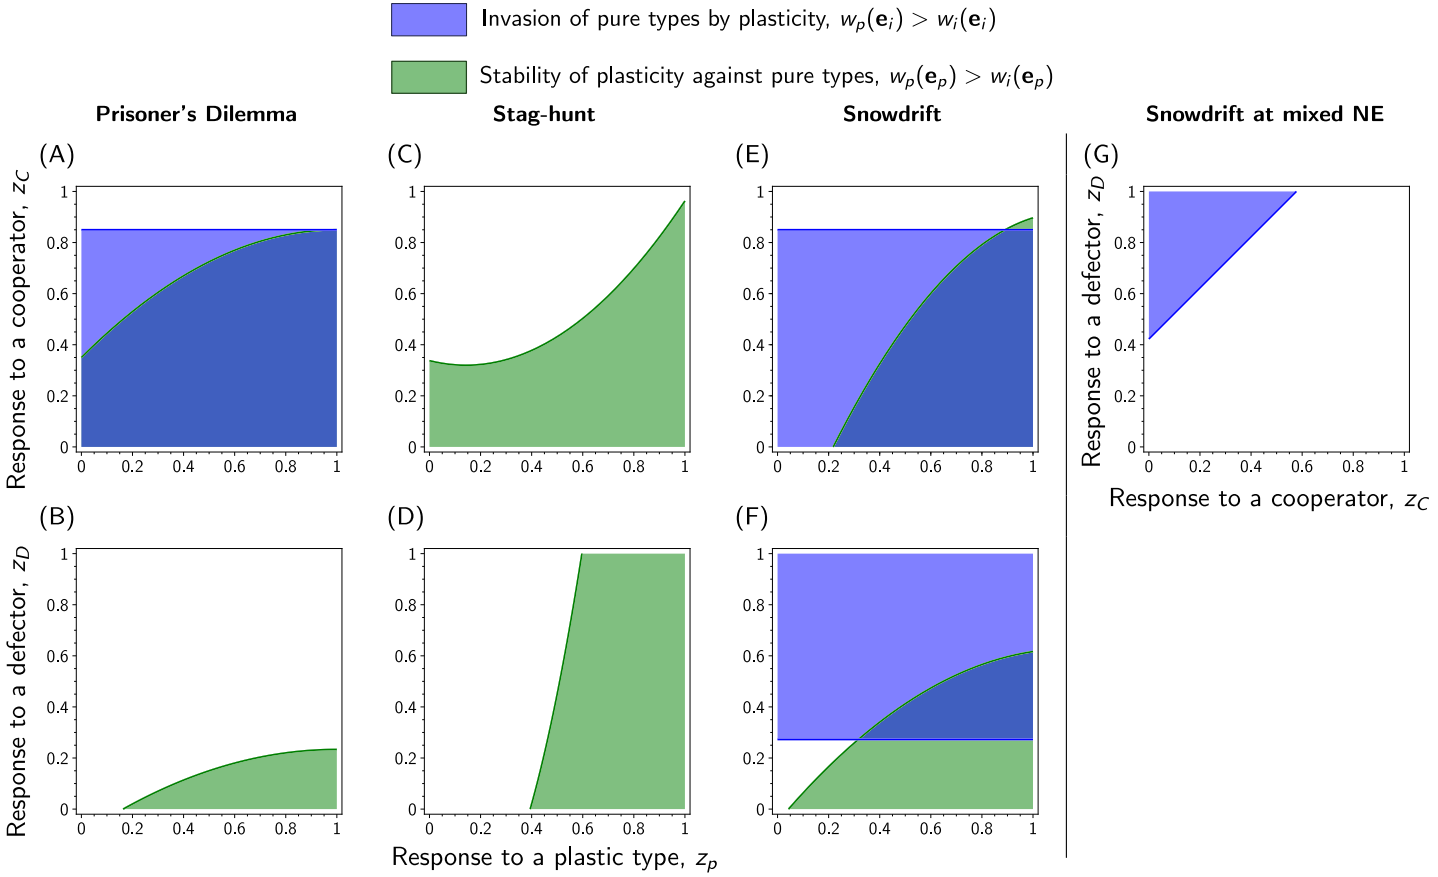$$\max\{R, P, \frac{T + S}{2}\} \leq \min\{R, S\}. \tag{10}$$

Fig. 1: Regions of stability and invasion of plastic types. (A)–(F) For each social dilemma game (columns) we show the response of the plastic type, that is the combinations (z_p, z_C) (top row) and (z_p, z_D) (bottom row) that allow invasion of monomorphic populations of pure types by plasticity (blue), and the ones that allow stability of plasticity against the invasion by pure types (green). We recall here that z_i is the probability to cooperate of the plastic type against each type $i \in \{C, D, p\}$. (G) Combinations of (z_C, z_D) that allow plasticity to invade a population of pure types at the mixed ESS of the Snowdrift game. If (z_C, z_D) are chosen in the blue region of (G) and (z_C, z_D, z_p) are further chosen at intersection of the green and blue regions of panels (E)–(F), plasticity is globally stable in the Snowdrift game. Parameter values: Prisoner's dilemma game: $(R, S, T, P) = (3, 1, 5, 2)$; Stag-hunt game: $(R, S, T, P) = (8, 0, 4, 3)$; Snowdrift game: $(R, S, T, P) = (3, 0.1, 5, -1)$; $k = 0.3$.

Similarly, using eq. 8, plasticity cannot be locally stable against mutant defectors if

$$\max\{R, P, \frac{T+S}{2}\} \leq \min\{T, P\}. \tag{11}$$

The last two inequalities are not satisfied in the Prisoner's dilemma, which means that a plastic type can always find a strategy that makes it stable against the invasion by pure types.

In sum, in the Prisoner's Dilemma game, the best outcome for a plastic type is local asymptotic stability. It is impossible for plasticity to invade a population at the stable equilibrium of the pure-type game because the plastic type cannot do better than pure defectors.

3.2 Stag-hunt game

We define this game as one with payoffs satisfying $R > T > P > S$ and $R + S > T + P$, which entails that (C, C) and (D, D) are two pure Nash equilibria (NE) of the game, with (C, C) being the Pareto-dominant equilibrium, and (D, D) being the risk-dominant equilibrium. Consequently (C, C) and (D, D) are both asymptotically stable equilibria of the replicator dynamics in the pure-type game. Solving inequality 7 for the payoffs of the Stag-hunt game, we find that plastic types cannot invade any monomorphic population

of pure types. However, solving inequality 8, we find that plasticity can be locally stable (Fig. 1C–D). This game is a good illustration of the idea that invasion ability and local stability of our plastic types might be impossible to reconcile. Indeed, the plastic type needs to anti-coordinate with C and D in order to ensure local stability, but the best it can do in monomorphic populations of either C or D is to coordinate; and this best effort is not even enough to be able to invade monomorphic populations.

We finally mention a result in the Stag-hunt game that may be of interest beyond evolutionary biology. Namely, in our model, costless plasticity in the Stag-hunt game is a solution to the problem of converging to the payoff-dominant equilibrium. Even though our main focus in this paper is on the case where the cost of plasticity is positive ($k > 0$), we mention this result because most evolutionary or learning processes known to us converge to the risk-dominant equilibrium, (D, D) (or stochastic evolutionary processes admit a stationary distribution that puts more mass on the risk-dominant equilibrium, e.g., Young, 1993). Here, if one sets $k = 0$ and $(z_C, z_D, z_p) = (1, 0, 1)$, then the face of the simplex such that defectors are at frequency 0 is globally stable. All points on this face are neutrally stable, i.e., any mix of cooperators and plasticity is globally stable. In such populations, everyone cooperates and thus achieves the maximum possible payoff in the Stag-hunt game.

In sum, in the Stag-hunt game, the best outcome for costly plasticity is local asymptotic stability, which is achieved through anti-coordination with the pure types. A plastic type cannot invade monomorphic populations because the best it can do is to coordinate with them but the plastic type must pay a cost. We finally noted that if we allow plasticity to be costless, then it can allow cooperation to be achieved from any initial condition, even when the population is close to the equilibrium where everyone defects.

3.3 Snowdrift game

In the Snowdrift game, the payoffs are such that $T > R > S > P$. This game thus calls for closer attention because it has a symmetric mixed NE. In order to achieve global stability, a plastic type must be able to invade any mixture of the pure types, including the NE mixture. This game is in fact an excellent illustration of the idea that playing the best response to every pure type does not always yield the best outcome for a plastic type. Indeed, in this game, playing the best response to cooperators and defectors guarantees that plasticity will invade monomorphic populations and the mixed NE, but this strategy does not guarantee local stability of plasticity. However there exists a strategy, which is not the best response, that guarantees global stability of plasticity, provided the payoffs allow it. Such a strategy is found by solving simultaneously the inequality in eq. 7, together with the following inequality

$$\pi(\mathbf{x}_{\text{NE}}) < x_{C,\text{NE}}\pi(z_C, C) + x_{D,\text{NE}}\pi(z_D, D) - k, \qquad (12)$$

which is the condition for plasticity to invade a population at the mixed NE of the Snowdrift game. It turns out to be quite difficult to reduce the inequalities in eq. 7 and eq. 12 (a system of inequalities involving linear and quadratic ones), but we provide in Fig. 1E–G an example showing that there exists a strategy and payoffs of the Snowdrift game such that they all hold, i.e., such that plasticity is globally stable.

The Snowdrift game also provides interesting counter-examples to conjectures that one might be tempted to make: if a plastic type can invade every monomorphic population, it cannot necessarily invade a population at the mixed NE; in Fig. 1E–G, take for example $(z_C, z_D, z_p) = (0.2, 0.4, z_p)$. The converse is also false: if a plastic type can invade the mixed NE, this does not mean that it can invade every monomorphic population in the support of the NE. However, in any game without pure NE, there always exists a strategy that guarantees invasion of both mixed NE and pure monomorphic populations; in particular, playing the best response against every pure type always guarantees such an outcome.

Another statement that can be proved wrong using the Snowdrift game is that games without pure symmetric NE always make possible global stability of plasticity. In the Snowdrift game, setting T very

large creates a game that has still no pure symmetric NE, but where there is no longer a strategy that makes plasticity globally stable. The intuitive reason is that setting T very large reduces substantially the upper bound on z_D such that plasticity is locally stable, i.e., the (green) region of local stability of plasticity in Fig. 1F shrinks and then no longer intersects the (blue) region where plasticity can invade a monomorphic population of defectors.

In sum in the Snowdrift game, we found that playing the best response allows the plastic type to guarantee co-existence with the pure types. We also found that a subset of Snowdrift games, where T is not too large, also allows the plastic type to become globally asymptotically stable.

3.4   Mutualism game

In this game we assume simply that $R > 0$ and $S > 0$, while $T = P = 0$, such that cooperation is a dominant action and the only NE of the pure-type game. Despite not being a social dilemma in the sense that there is no conflict between cooperation and rational behavior, this game is still interesting when considering the possible evolution of plasticity. Indeed in this game there are values of the payoffs for which plasticity is neither locally stable nor able to invade a monomorphic population of cooperators, for any value of (z_C, z_D, z_p), which implies that any share of plastic mutants – irrespective of their strategy – will be unable to survive natural selection. From our previous analysis, we know that plasticity cannot invade a monomorphic population of cooperators, since cooperation is a dominant action. Plasticity can however invade a population of defectors by playing a sufficiently high z_D. However, if $S \geq R$, then eq. 10 is satisfied, which means that plasticity cannot be locally stable. If, on the other hand, $S < R$ then a plastic type can resist invasion by cooperators through defection against them. In Fig. 2, we show the phase portrait of trajectories in a particular Mutualism game where $S \geq R$, but also for the three other games of cooperation studied in this section.

In sum, in the Mutualism game, if $S \geq R$, the plastic type is unable to invade or be locally stable against pure types, while if $S < R$ the best outcome for plasticity is local stability.

4   Other games

Can we generalize some of the results we derived for 2×2 games of cooperation? Also, do they hold for other classes of games? In this section we first state general results regarding what is possible or impossible for plastic types and we then focus on the special class of zero-sum games. In all of our results, we assume that the pure-type game is generic (i.e., no two outcomes yield the same payoffs for any player).

We begin by recalling here a classical result that we will be using throughout this section.

Theorem (Static equilibria and the replicator dynamics; Hofbauer and Sigmund, 1998; Webb, 2007). *Let $\mathbb{A}$ be the set of asymptotically stable equilibria of the replicator dynamics (eq. 5), $\mathbb{N}$ the set of symmetric Nash equilibria of the game, $\mathbb{N}^*$ the set of strict symmetric Nash equilibria of the game, and $\mathbb{F}$ the set of equilibria (fixed points) of the replicator dynamics (eq. 5). We then have*

$$\mathbb{A} = \mathbb{N}^* \subseteq \mathbb{N} \subseteq \mathbb{F}.$$

The above theorem's purpose is mainly to allow us making statements about the replicator dynamics by using easier arguments based on static payoff comparisons. A simple proposition that we can now prove is the following one.

Proposition 1 (Necessary conditions for global asymptotic stability of plasticity). *If there exists a $\bar{k}$ such that, for all $0 < k < \bar{k}$, $\lim_{t \to \infty} \phi_p(t, \mathbf{x}_0) = 1$, for all $\mathbf{x}_0 \in \text{int}(\Delta^{n+1})$, then there exists $\mathbf{z} \in (\Delta\mathcal{A})^{n+1}$ such that:*
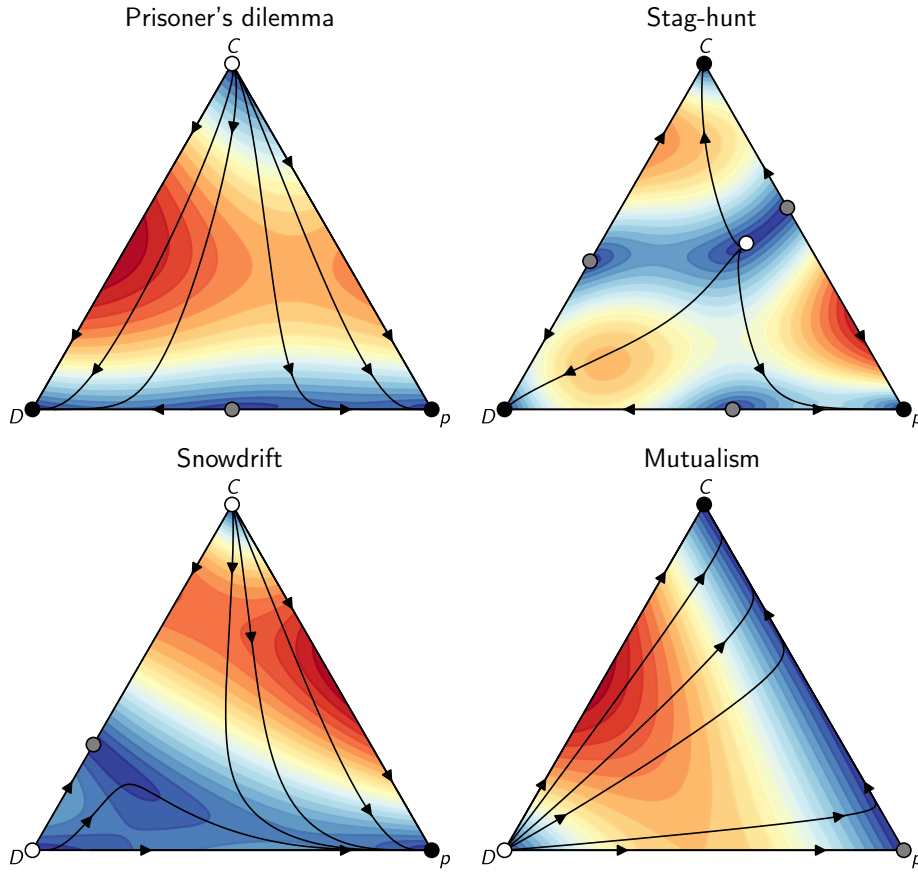
Fig. 2: Replicator dynamics with added plastic types in 4 games of cooperation. Each triangle represents a projection of the 3-simplex on the 2-dimensional plane, where at each vertex the corresponding strategy (C, D, or p) is at frequency 1. On the face opposed to a vertex, the corresponding strategy is absent, i.e., it is at frequency 0. These plots are displayed for the best possible outcome that can be achieved by a plastic type. Parameter values: Prisoner's Dilemma: $(R, S, T, P) = (3, -2, 5, 1)$, and $(z_C, z_D, z_p) = (0, 0, 1)$; Stag-hunt: $(R, S, T, P) = (8, 0, 4, 3)$, and $(z_C, z_D, z_p) = (0, 1, 1)$; Snowdrift: $(R, S, T, P) = (2.5, 0.1, 5, -1)$, and $(z_C, z_D, z_p) = (0.03, 0.45, 0.9)$; Mutualism: $(R, S, T, P) = (5, 6, 0, 0)$, and $(z_C, z_D, z_p) = (1, 1, 1)$; in all games we used a cost of $k = 0.3$. Contours indicate speed with blue corresponding to slower dynamics and red corresponding to faster dynamics. Black dot: Asymptotically stable equilibrium (a "sink") of the replicator dynamics; gray dot: an unstable saddle; white dot: an unstable source.

(1) $\pi(j, j) \leq \pi(\mathbf{z}_j, j)$, $\forall j \in \mathcal{A}$,

(2) $\pi(j, \mathbf{z}_j) < \pi(\mathbf{z}_p, \mathbf{z}_p)$, $\forall j \in \mathcal{A}$,

(3) $\sum_{j=1}^{n} x_{j,NE} \pi(\mathbf{z}_j, j) > \pi(\mathbf{x}_{NE})$, $\forall \mathbf{x}_{NE} \in \Delta\mathcal{A}$.

Proof. We can prove statements (1) and (3) at the same time. Statement (1) means that no pure strategy constitutes a strict NE in the extended game with plasticity, hence that plasticity can invade any monomorphic population of pure types. Statement (3) says that any $\mathbf{x}_{NE} \in \Delta\mathcal{A}$ is locally unstable against the invasion by plasticity. Indeed, if there were other NE than plasticity, then they would be at least stable (and asymptotically stable if strict). But since plasticity is globally asymptotically stable, no other pure or mixed strategy is stable (and a fortiori asymptotically stable), hence there are no other NE to the extended game with plasticity, so (1) and (3) follow.

Statement (2) means that the plastic type constitutes a pure strict NE of the extended game with plasticity. But we also know that being a strict NE is equivalent to being asymptotically stable by the above theorem. Since we assumed that plasticity is globally asymptotically stable, (2) follows.

□

The three statements of Prop. 1 have several implications and interpretations. First, statements (1) and (2) taken together imply that the task of achieving global stability is not trivial for the plastic type. If we were to take the standpoint of a plastic type that can choose its response $\mathbf{z}$ freely, then in order to maximize the chances of invading a monomorphic population consisting of type $i \in \mathcal{A}$ the plastic type would like to choose

$$\arg \max_{\mathbf{z}_i} \{\pi(\mathbf{z}_i, i)\} = \beta(i), \tag{13}$$

where we used $\beta(\cdot)$ to denote the best response function. However, statement (2) of Prop. 1 (in addition to our examples of games of cooperation studied above) shows that the best response is not the strategy that maximizes the chances of a plastic type resisting invasion by the pure types. Indeed, to achieve such stability against pure types, the plastic type should choose

$$\arg \min_{\mathbf{z}_i} \{\pi(i, \mathbf{z}_i)\}. \tag{14}$$

The strategy $\mathbf{z}_i$ in eq. 14 is different from the one in eq. 13, except in zero-sum games (we investigate zero-sum games in more detail below), which means that optimizing invasion ability vs. stability of plasticity can be two conflicting tasks, as we illustrated in the Stag-hunt game above. The inequality in statement (2) of Prop. 1 also implies that in order to maximize its chances of resisting invasion, the plastic type must further take the strategy

$$\arg \max_{\mathbf{z}_p} \{\pi(\mathbf{z}_p, \mathbf{z}_p)\}. \tag{15}$$

We have seen that in games of cooperation choosing such a $\mathbf{z}_p$ is equivalent to the plastic type cooperating with itself. However, if the plastic type establishes cooperation in order to be a stable strategy against pure types, one can expect that this opens the door to plastic free riders who defect against the established nice plastic types (in the Prisoner's Dilemma game and the Snowdrift game). Hence the cooperation established by our plastic types seems to be only a transient state, but to to delineate the cases where this intuition holds true, an exhaustive study of the case where the response $\mathbf{z}$ of the plastic is under selection would be necessary (this is beyond the scope of the current paper).

Another consequence of Prop. 1 and eq. 13 is that a plastic type cannot invade populations at a monomorphic pure NE. As we have seen in the Snowdrift game example above, this does not mean that every game without pure NE allows for the global stability of plasticity. However, we show in the next proposition that in games without pure NE, there always exists a strategy that guarantees at least co-existence of plasticity and the pure types.

Proposition 2 (Coexistence of plasticity and pure types in games with mixed NE). *In games with no pure symmetric Nash equilibrium, there always exists a strategy $\mathbf{z} \in (\Delta\mathcal{A})^{n+1}$ such that plasticity persists in the population in the long run, i.e., there exists a $\bar{k}$ such that, for all $0 < k < \bar{k}$, any $\mathbf{x} \in \Delta\mathcal{A}$ (i.e., such that $x_p = 0$) is unstable.*

Proof. To show that any state $\mathbf{x} \in \Delta\mathcal{A}$ such that $x_p = 0$ is unstable, it suffices to show that there exists a cost $k > 0$ such that $w_p(\mathbf{x}) > \bar{w}(\mathbf{x})$, for all $\mathbf{x} \in \Delta\mathcal{A}$, where we recall that $\bar{w}(\mathbf{x})$ denotes the average fitness in the population at state $\mathbf{x} \in \Delta\mathcal{A}$. Note that the fitness of any type $i \in \mathcal{A}$ at state $\mathbf{x} \in \Delta\mathcal{A}$ is

$$w_i(\mathbf{x}) = \sum_{j \in \mathcal{A}} x_j \pi(i, j). \tag{16}$$

The fitness at state $\mathbf{x} \in \Delta\mathcal{A}$ of a mutant plastic type is

$$w_p(\mathbf{x}) = \sum_{j \in \mathcal{A}} x_j \pi(\mathbf{z}_j, j) - k. \tag{17}$$

Since the game has no pure NE, there exists for each $j \in \mathcal{A}$ and $i \in \mathcal{A}$ a strategy $\mathbf{z}_j \in \Delta\mathcal{A}$ such that $\pi(\mathbf{z}_j, j) \geq \pi(i, j)$ (moreover the inequality is strict for some $\mathbf{z}_j \in \mathcal{A}$). Let the plastic type adopt such a strategy $\mathbf{z}$. We then have that $w_p(\mathbf{x}) > w_i(\mathbf{x})$. Since the average fitness $\bar{w}(\mathbf{x})$ is a convex combination

of the $w_i(\mathbf{x})$, this implies $w_p(\mathbf{x}) > \bar{w}(\mathbf{x})$. Hence the mutant plastic type increases in frequency at any population state $\mathbf{x} \in \Delta\mathcal{A}$, which makes such states unstable under the replicator dynamics (eq. 5) and leads to the desired result. □

We just defined the conditions for plasticity to persist in populations of pure types, but these do not guarantee local stability of plasticity. The following proposition establishes a necessary condition for the *inability* of plasticity to be immune against invasion by pure types.

Proposition 3 (Impossible stability of plasticity). *If there exists $i \in \mathcal{A}$ such that plasticity cannot resist invasion by i, i.e. such that $w_p(\mathbf{e}_p) < w_i(\mathbf{e}_p)$, then (i, i) is the symmetric Pareto-efficient equilibrium of the game.*

Proof. If plasticity cannot resist invasion by $i \in \mathcal{A}$, this means that

$$w_p(\mathbf{e}_p) < w_i(\mathbf{e}_p)$$

$$\pi(\mathbf{z}_p^{\mathrm{S}}, \mathbf{z}_p^{\mathrm{S}}) \leq \pi(i, \mathbf{z}_i), \qquad \forall \mathbf{z}_i \in \Delta\mathcal{A}. \tag{18}$$

By way of contradiction, suppose that (i, i) is not the symmetric Pareto-efficient equilibrium of the game. Then there exists $\bar{\mathbf{z}}_p \in \Delta\mathcal{A}$ such that $\pi(\bar{\mathbf{z}}_p, \bar{\mathbf{z}}_p) > \pi(i, i)$. Hence a strategy of the plastic type with $\mathbf{z}_p = \bar{\mathbf{z}}_p$ and $\mathbf{z}_i = i$ would guarantee local stability of plasticity. But, by eq. 18, this contradicts our assumption that plasticity is not locally stable against the invasion by i. This proves the result. □

The converse is not true. If (i, i) is Pareto-efficient, it might still be possible to find strategies of the plastic type such that it resists invasion by i. Taking any of the three social dilemmas above, cooperation is Pareto-efficient, but it is always possible to find a strategy that guarantees local stability of plasticity.

4.1 Zero-sum games

From the discussion of the previous sections, it seems that it is not easy to find classes of games where one can guarantee the global stability of plasticity. Looking more closely at our previous results, however, one might infer that if the two tasks of maximizing one's own payoff and minimizing the other's payoff are two compatible endeavours, plastic types can achieve global stability. Hence, the class of *strictly competitive games* would seem to favor plastic types. A game is strictly competitive if, for all $\ell \in \mathcal{A}$ and any pair of strategies $i \in \mathcal{A}$, $j \in \mathcal{A}$, the inequality $\pi(i, j) > \pi(\ell, j)$ implies that $\pi(j, i) < \pi(j, \ell)$. In particular, maximizing one's own payoff is equivalent to minimizing the other's payoff. Note that in a strictly competitive game, all strategy pairs are Pareto-efficient since switching from any action that induces a higher payoff for a focal player induces a loss for his opponent. It has been shown that strictly competitive games are affine transformations of zero-sum games (Adler et al., 2009). We next show that in the class of symmetric zero-sum games without pure NE, there always exists a strategy of the plastic type that guarantees its global stability under the replicator dynamics.

Proposition 4 (Global stability of costly plasticity in zero-sum games). *In zero-sum games with no pure symmetric NE, there always exists a strategy $\mathbf{z} \in (\Delta\mathcal{A})^{n+1}$ of the plastic type such that it is globally asymptotically stable under the replicator dynamics.*

Proof. In order to show that plasticity is globally asymptotically stable, we show a stronger result, i.e., that at any state $\mathbf{x} \in \mathrm{int}(\Delta^{n+1})$, we have $w_p(\mathbf{x}) > \bar{w}(\mathbf{x})$. This implies that $\dot{x}_p > 0$ at any state $\mathbf{x} \in \mathrm{int}(\Delta^{n+1})$ and since the simplex is a bounded set, this in turn implies that $\lim_{t\to\infty} \phi_p(t, \mathbf{x}_0) = 1$ for all $\mathbf{x}_0 \in \mathrm{int}(\Delta^{n+1})$.

348 Now, to show that $w_p(\mathbf{x}) > \bar{w}(\mathbf{x})$ we will show that $w_p(\mathbf{x}) > w_i(\mathbf{x})$ for all $i \in \mathcal{A}$. To do so, we write the
349 fitness of a pure type i as

$$w_i(\mathbf{x}) = x_i \pi(i,i) + x_p \pi(i, \mathbf{z}_i) + \sum_{\substack{j \neq i \\ j \neq p}} x_j \pi(i,j). \tag{19}$$

351 Similarly, we write the fitness of the plastic type as

$$w_p(\mathbf{x}) = x_i \pi(\mathbf{z}_i, i) + x_p \pi(\mathbf{z}_p, \mathbf{z}_p) + \sum_{\substack{j \neq i \\ j \neq p}} x_j \pi(\mathbf{z}_j, j) - k. \tag{20}$$

353 In order to compare the part of the fitness that is under the summation symbol, note that since there
354 is no pure NE, there exists for each $j \in \mathcal{A}$ a strategy $\mathbf{z}_j$ of the plastic type such that $\pi(\mathbf{z}_j, j) \geq \pi(i,j)$.
355 Hence there is a strategy $\mathbf{z}$ such that the term under the summation symbol is larger for the plastic type
356 than the pure type, or

$$\sum_{\substack{j \neq i \\ j \neq p}} x_j \pi(\mathbf{z}_j, j) > \sum_{\substack{j \neq i \\ j \neq p}} x_j \pi(i,j). \tag{21}$$

The remaining terms in the fitness are also larger for the plastic type than for type $i \in \mathcal{A}$. Indeed, we
have $\pi(\mathbf{z}_p, \mathbf{z}_p) = 0$ for any $\mathbf{z}_p \in \Delta\mathcal{A}$. At the same time since the game is zero-sum with no pure NE, there
exists $\mathbf{z}_i$ such that $\pi(i, \mathbf{z}_i) < 0$, thus

$$\pi(\mathbf{z}_p, \mathbf{z}_p) > \pi(i, \mathbf{z}_i).$$

Also, since $\pi(i, \mathbf{z}_i) = -\pi(\mathbf{z}_i, i)$, and $\pi(i,i) = 0$, we also have

$$\pi(\mathbf{z}_i, i) > \pi(i,i).$$

358 Taken together the last three inequalities imply that there exists a k such that $w_p(\mathbf{x}) > w_i(\mathbf{x})$ for all
359 $i \in \mathcal{A}$ at any state $\mathbf{x} \in \text{int}(\Delta^{n+1})$, hence $w_p(\mathbf{x}) > \bar{w}(\mathbf{x})$, which completes the proof. $\qquad\square$

360 Note that the requirement that the game possesses no pure symmetric NE excludes symmetric 2×2
361 zero-sum games, since these games necessarily have a dominant action. Another remark is that the proof
362 is also valid for strictly competitive games since our argument is independent of affine transformations
363 to the payoffs of the game. This remark allows us to apply our results to the famous *good* Rock-Paper-
364 Scissors game, where winning yields a payoff of a, while losing induces a loss of b, and $a > b$ (Sandholm,
365 2011), which entails the payoff matrix

$$
\begin{array}{c}
\\
R \\
P \\
S
\end{array}
\begin{array}{c}
\begin{array}{ccc}
R & P & S
\end{array} \\
\left[
\begin{array}{ccc}
0 & -b & a \\
a & 0 & -b \\
-b & a & 0
\end{array}
\right].
\end{array} \tag{22}
$$

367 In such a game, the interior NE, $\mathbf{x}_{\text{NE}} = (1/3, 1/3, 1/3)$, is globally stable under the replicator dynamics,
368 with solution trajectories displaying damped oscillations around this equilibrium. In Fig. 3, we show the
369 effect of adding a plastic type with a strategy described in the proof of Prop. 4 to the good RPS game. In
370 this game, this creates a plastic type that plays the best response to any pure type and plays any $i \in \mathcal{A}$
371 against itself such that $\mathbf{z}_i = \beta(i)$ for any $i \in \{R, P, S\}$ and $\mathbf{z}_p = (1, 0, 0)$ [the choice of always playing R
372 is arbitrary, in accordance with the proof of Prop. 4]. Even though the plastic type cannot invade on the
373 faces of the simplex of dimension 2 (because symmetric 2×2 zero-sum games necessarily have a dominant
374 action, and hence a pure symmetric NE), it achieves a larger payoff than any fully mixed population.
375 This explains why it is globally stable in the interior of the simplex.
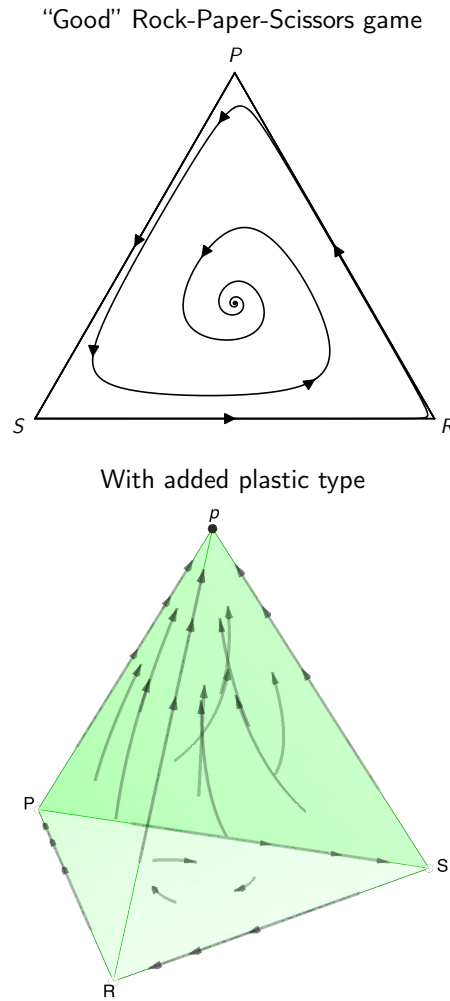
Fig. 3: Replicator dynamics in the original good RPS (top), and with added plastic type (bottom). The 3-strategy simplex was produced using our own code (available), while the 4-strategy simplex was produced using Dynamo (Franchetti and Sandholm, 2013). Parameter values: Payoff matrix of eq. 22 with $a = 2$ and $b = 1$; $\mathbf{z}_R = (0, 1, 0)$; $\mathbf{z}_P = (0, 0, 1)$; $\mathbf{z}_S = (1, 0, 0)$; $\mathbf{z}_p = (1, 0, 0)$.

5 Discussion

In this work, we considered the introduction of plastic types in populations consisting of simple individuals adopting pure strategies in evolutionary symmetric games. We showed that despite having the capacity of responding in any possible way to other types in the population, there does not exist a general strategy for plasticity to eliminate pure types from the population. There are several explanations to this. The first one is that in games with pure symmetric Nash equilibria, costly plasticity cannot invade an equilibrium population. Moreover, we demonstrated that there is often a conflict between optimizing invasion success and optimizing stability for plastic types, the best illustration of this result being the Stag-Hunt game, where these two tasks (optimizing invasion ability versus stability) lead to totally different behavioral responses of the plastic types. We also saw that the success of plasticity is dependent on the class of games played in the population. There are games, such as zero-sum games without pure symmetric NE, where there always exists a strategy that grants plasticity global stability under the replicator dynamics. On the other hand, there are games, such as certain Mutualism games, where plasticity can neither invade an equilibrium population nor be stable against the invasion by pure types.

In our more detailed investigation of games of cooperation, we found that in the Prisoner's Dilemma and the Snowdrift game, plastic types should cooperate with themselves with a strictly positive probability in order to be immune against the invasion by pure types. However, invasion success of pure populations

by plasticity is generally uncorrelated to playing a cooperative strategy. In the Stag-hunt game it is just impossible for costly plasticity to invade an equilibrium population. However, costless plasticity is able to do so, and can even establish cooperation (the payoff-dominant outcome) in a population that was initially at equilibrium for defection (the inferior risk-dominant outcome). This result is similar to what Robson (1990) found in an earlier work (we compare our work with that of Robson, 1990, in more detail below).

It is insightful compare our results to those of previous studies of the evolution of plasticity. In particular, previous work on frequency-independent selection (e.g., Gomulkiewicz and Kirkpatrick, 1992) has revealed that variance in fitness is a key factor driving the evolution of plasticity. In our setting, we notice that in games with pure Nash equilibria (NE), the population does not experience variance in payoffs at equilibrium. Hence, it makes sense that there is no advantage to adapt behavior to different circumstances, since there is only one circumstance which is to face the pure type that is evolutionarily stable. However, in games without pure NE, there is always variance in payoffs because one might be matched with any type in the support of the mixed NE of the game. This explains why plasticity can always invade a mixed NE, because it has the possibility to express different responses to the different existing types and potentially obtain a payoff that is greater than the NE payoff. Altogether, these results suggest that in general the ability to collect more fitness-relevant information is not always beneficial, as long as this information is costly to acquire (even if costs are very small), as other studies of the evolution of learning have previously shown (Wakano et al., 2004; Nakahashi et al., 2012; Aoki and Feldman, 2014; Dridi and Lehmann, 2015).

While our results fit well within the literature on plasticity evolution theory, there is also empirical evidence that our assumptions about plasticity correspond to some extent to real-world instances of plastic phenotypes. If the evolving behavioral trait of interest is linked to morphological or hormonal characteristics, then it could be possible for a plastic type to base their strategy on, e.g., a chemical detection mechanism, vision or smell. This is the case for many traits that indicate the quality of an individual, which will then determine the payoff of individuals interacting with it, such as size or weight in fights (Riechert, 1978; Maynard-Smith and Harper, 2004; Arnott and Elwood, 2009); this is also the case for traits that indicate the compatibility and quality of potential mates (Potts et al., 1991; Roberts and Gosling, 2003).

The ability to gather information about conspecifics or the environment is critical in evolution and, as a consequence, the proposition that individuals with more information may have an evolutionary advantage in social interactions is not new. Robson (1990) developed a model where his plastic type could only distinguish between plastic and non-plastic types and condition their action on this cue. He assumed that plasticity was costless and found in particular that plasticity can invade a population at the inferior equilibrium of a coordination game, and lead to the superior payoff-dominant equilibrium; however his plastic type could not co-exist with the superior pure type that plays the payoff-dominant action, in contrast to what we found. Banerjee and Weibull (1995) studied a very similar model to ours, but constrained the plastic type to play a best response to every pure type. We saw that this best response assumption is not the best strategy for plastic types when it comes to evolutionary stability. This phenomenon is exemplified by the Snowdrift game, where playing the best response to every pure type does not allow plasticity to be locally stable against the invasion of pure mutants; but there are Snowdrift games where playing a strategy different than the best response allows plasticity to be not only locally but globally stable. Our model is thus an extension of the ideas developed in these two earlier studies (Robson, 1990; Banerjee and Weibull, 1995).

From a broader perspective, given the recent renewed interest in the evolution of strategies in the repeated Prisoner's dilemma (Press and Dyson, 2012; Adami and Hintze, 2013; Stewart and Plotkin, 2013, 2016), our work helps bridge the gap between these recent studies and classical work on the evolution of pure non-plastic strategies in one-shot games (Hofbauer and Sigmund, 1998). Indeed, in order to be able to

express a repeated-game strategy, an organism must first have the capacity to express a plastic social phenotype. Another recent trend is the focus on environmental variation in social evolution (Ashcroft et al., 2014; Dridi and Lehmann, 2014; Weitz et al., 2016), and our study reminds that environmental variation is already embedded within social evolution through frequency-dependent payoffs (Graves and Weinreich, 2017), which is why we find that plasticity may evolve even in the absence of variation in the game structure itself.

Our model finally introduces new questions and calls for further research on the topic. Most notably, we ignored evolution of the plastic response itself, but it is natural to imagine that if one plastic type successfully invades, other plastic mutants with a different behavioral response might appear and displace the original plastic type; there is then no guarantee that the optimal plastic type will also be evolutionarily stable. Indeed, Robson (1990) showed that in the Prisoner's dilemma, if a first plastic mutant might successfully invade and establish cooperation, the introduction of further plastic mutants might lead to the loss of cooperation. A possible avenue for future research would be to see if this result applies to our setting as well, in the Prisoner's dilemma and other games.

One might argue that the ability to express mixed strategies is what gives plasticity its strategic advantage rather than the ability to be plastic. It is not the case for invasion ability because best response to pure types is always a pure strategy, so a plastic type only able to play pure strategies would not be disfavored with respect to our plastic types. However the ability to play mixed strategies is critical when looking for plastic types that are immune against the invasion by pure types.

Another extension of our work would be to consider the possibility that the game payoffs change as a function of time, and in this setting a plastic type would be able to condition its action not only on the type of the opponent but also on the type of game being played. Heller (2004) modelled a situation where some individuals could detect the state of the environment and the opponent's type in a context of a fluctuating game, but she assumed that plastic types (which she calls learning agents) only play the best response to pure types. It would be interesting to see what would happen if we do not constrain the plastic response to be the best response, since we saw in our model that the best response does not always yield the best outcome for plasticity. These and other extensions will help us better understand the evolution of complex strategies in realistic changing environments.

Acknowledgements

References

Adami, C., and A. Hintze. 2013. Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything. Nature Communications 4. doi:10.1038/ncomms3193.

Adler, I., C. Daskalakis, and C. H. Papadimitriou. 2009. A note on strictly competitive games. In Internet and Network Economics, Lecture Notes in Computer Science, pages 471–474. Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-10841-9_44.

André, J.-B., and T. Day. 2007. Perfect reciprocity is the only evolutionarily stable strategy in the continuous iterated prisoner's dilemma. Journal of Theoretical Biology 247:11–22. doi:10.1016/j.jtbi. 2007.02.007.

Aoki, K., and M. W. Feldman. 2014. Evolution of learning strategies in temporally and spatially variable environments: A review of theory. Theoretical Population Biology 91:3–19. doi:10.1016/j.tpb.2013.10.004.

Arnott, G., and R. W. Elwood. 2009. Assessment of fighting ability in animal contests. Animal Behaviour 77:991–1004. doi:10.1016/j.anbehav.2009.02.010.

Ashcroft, P., P. M. Altrock, and T. Galla. 2014. Fixation in finite populations evolving in fluctuating environments. Journal of The Royal Society Interface 11:20140663. doi:10.1098/rsif.2014.0663.

Axelrod, R., and W. D. Hamilton. 1981. The evolution of cooperation. Science 211:1390–1396.

Banerjee, A., and J. W. Weibull. 1995. Evolutionary selection and rational behavior. In A. Kirman, and M. Salmon, eds., Learning and Rationality in Economics, pages 344–363. Blackwell, Cambridge, MA.

Bergstrom, C. T., and L. A. Dugatkin. 2016. Evolution. 2nd edition. W. W. Norton & Company, New York, NY.

Diggle, S. P., A. S. Griffin, G. S. Campbell, and S. A. West. 2007. Cooperation and conflict in quorum-sensing bacterial populations. Nature 450:411–414. doi:10.1038/nature06279.

Dridi, S., and L. Lehmann. 2014. On learning dynamics underlying the evolution of learning rules. Theoretical Population Biology 91:20–36. doi:10.1016/j.tpb.2013.09.003.

———. 2015. A model for the evolution of reinforcement learning in fluctuating games. Animal Behaviour 104:87–114. doi:10.1016/j.anbehav.2015.01.037.

———. 2016. Environmental complexity favors the evolution of learning. Behavioral Ecology 27:842–850. doi:10.1093/beheco/arv184.

Franchetti, F., and W. H. Sandholm. 2013. An introduction to Dynamo: Diagrams for evolutionary game dynamics. Biological Theory 8:167–178. doi:10.1007/s13752-013-0109-z.

Gavrilets, S., and S. M. Scheiner. 1993a. The genetics of phenotypic plasticity. V. Evolution of reaction norm shape. Journal of Evolutionary Biology 6:31–48.

———. 1993b. The genetics of phenotypic plasticity. VI. Theoretical predictions for directional selection. Journal of Evolutionary Biology 6:49–68.

Gomulkiewicz, R., and M. Kirkpatrick. 1992. Quantitative genetics and the evolution of reaction norms. Evolution 46:390–411. doi:10.2307/2409860.

Grafen, A. 1990. Biological signals as handicaps. Journal of Theoretical Biology 144:517–546. doi:10.1016/S0022-5193(05)80088-8.

Graves, C. J., and D. M. Weinreich. 2017. Variability in fitness effects can preclude selection of the fittest. Annual Review of Ecology, Evolution, and Systematics 48:399–417. doi:10.1146/annurev-ecolsys-110316-022722.

Heller, D. 2004. An evolutionary approach to learning in a changing environment. Journal of Economic Theory 114:31–55. doi:10.1016/S0022-0531(03)00117-0.

Hilbe, C., Š. Šimsa, K. Chatterjee, and M. A. Nowak. 2018. Evolution of cooperation in stochastic games. Nature 559:246. doi:10.1038/s41586-018-0277-x.

Hofbauer, J., and K. Sigmund. 1998. Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge, UK.

Maynard-Smith, J., and D. Harper. 2004. Animal Signals. OUP Oxford, New York, NY.

McElreath, R., and R. Boyd. 2007. Mathematical Models of Social Evolution: A Guide for the Perplexed. University Of Chicago Press, Chicago, IL.

McNamara, J. M., C. E. Gasson, and A. I. Houston. 1999. Incorporating rules for responding into evolutionary games. Nature 401:368–371. doi:10.1038/43869.

Nakahashi, W., J. Y. Wakano, and J. Henrich. 2012. Adaptive social learning strategies in temporally and spatially varying environments. Human Nature 23:386–418. doi:10.1007/s12110-012-9151-y.

Nowak, M. A. 2006. Evolutionary Dynamics: Exploring the Equations of Life. Harvard University Press.

Potts, W. K., C. J. Manning, and E. K. Wakeland. 1991. Mating patterns in seminatural populations of mice influenced by MHC genotype. Nature 352:619. doi:10.1038/352619a0.

Press, W. H., and F. J. Dyson. 2012. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. Proceedings of the National Academy of Sciences 109:10409–10413. doi:10.1073/pnas.1206569109.

Riechert, S. E. 1978. Games spiders play: Behavioral variability in territorial disputes. Behavioral Ecology and Sociobiology 3:135–162. doi:10.1007/BF00294986.

Roberts, S. C., and L. M. Gosling. 2003. Genetic similarity and quality interact in mate choice decisions by female mice. Nature Genetics 35:103–106. doi:10.1038/ng1231.

Robson, A. J. 1990. Efficiency in evolutionary games: Darwin, Nash and the secret handshake. Journal of Theoretical Biology 144:379–396. doi:10.1016/S0022-5193(05)80082-7.

Sandholm, W. H. 2011. Population Games and Evolutionary Dynamics. MIT Press.

Stewart, A. J., and J. B. Plotkin. 2013. From extortion to generosity, evolution in the Iterated Prisoner's Dilemma. Proceedings of the National Academy of Sciences 110:15348–15353. doi:10.1073/pnas.1306246110.

———. 2016. Small groups and long memories promote cooperation. Scientific Reports 6:26889. doi:10.1038/srep26889.

Taylor, P. D., and L. B. Jonker. 1978. Evolutionary stable strategies and game dynamics. Mathematical Biosciences 40:145–156. doi:10.1016/0025-5564(78)90077-9.

Wakano, J. Y., K. Aoki, and M. W. Feldman. 2004. Evolution of social learning: A mathematical analysis. Theoretical Population Biology 66:249–258. doi:10.1016/j.tpb.2004.06.005.

Webb, J. N. 2007. Game theory: Decisions, Interactions and Evolution. Springer Verlag, London.

Weitz, J. S., C. Eksin, K. Paarporn, S. P. Brown, and W. C. Ratcliff. 2016. An oscillating tragedy of the commons in replicator dynamics with game-environment feedback. Proceedings of the National Academy of Sciences 113:E7518–E7525. doi:10.1073/pnas.1604096113.

Young, H. P. 1993. The evolution of conventions. Econometrica 61:57–84. doi:10.2307/2951778.

Zahavi, A. 1975. Mate selection: A selection for a handicap. Journal of Theoretical Biology pages 205–214.