

Evolutionary Game Dynamics, Cooperation, and Costly Punishment

Joseph J. Armao IV
Advisor: Prof. Martin Nowak

March 30, 2009

*Senior Thesis submitted in partial fulfillment of the honors
requirements for the degree of Bachelor of Arts in Mathematics*

Harvard University

Contents

1	Introduction	4
1.1	The Prisoner's Dilemma	4
1.2	Evolution in Games	5
1.3	Note to the Reader	7
2	Static Game Analysis	9
2.1	Symmetric Game Theory and Nash Equilibria	9
2.2	Evolutionarily Stable Strategies	16
3	Game Dynamics in Continuum Populations	23
3.1	Game Dynamics	23
3.2	Equilibrium States	26
3.2.1	Dynamical Systems Formalism	26
3.2.2	A Brief Note on Computation	29
3.3	Dynamic and Static Equilibria	29
3.3.1	Weak Payoff Positivity	30
3.3.2	Monotonicity	32
4	Specific Dynamics	37
4.1	Fitness and Selection Strength	37
4.2	Replicator Dynamics	38
4.2.1	Motivation and Definition	39
4.2.2	Asymptotic Stability and the ESS	40
4.3	Viability Updating	42
4.3.1	Motivation and Definition	43
4.3.2	Dynamical Classification	45
4.4	Two-Strategy Dynamics	45
4.4.1	The Replicator Equation	47
4.4.2	The Viability Equation	48
4.4.3	Nash and ESS Analogues	52
5	Cooperation with Costly Punishment	54
5.1	Game-Theoretic Analysis	54
5.1.1	Nash Equilibria and ESS	55
5.2	Replicator Dynamics Analysis	56
5.2.1	Dynamical Equilibria	56

5.2.2	Edge Dynamics	58
5.3	Viability Updating Analysis	60
5.3.1	Weak Selection Equilibria	60
5.3.2	Edge Dynamics	61
5.4	Discussion of Results	63
6	Conclusion	65
A	Normal-Form Games and Nash Existence	66
B	8-Strategy Cooperation / Punishment Game	70
B.1	Payoff Matrix	70
B.2	Replicator and Viability Edge Dynamics	71
C	Viability Updating: Weak-Selection Limit	73
D	Statement of Auxiliary Theorems	74
D.1	Picard-Lindelöf Theorem	74
D.2	Hartman Linearization Theorem	74

Acknowledgements

I would like to thank Professor Martin Nowak for advising and supervising my work. I am fortunate to be among the many students he has inspired to pursue research in evolutionary dynamics. I am indebted to Dave Rand and Corina Tarnita for their invaluable guidance and insight. I would like to thank Professor Peter Kronheimer for his ongoing counsel, which helped shape my thesis throughout the process. I thank my parents for providing the encouragement I needed throughout this endeavor (and many other endeavors, along the way). Finally, I would like to acknowledge the immeasurable learning experience the Harvard math department grants its students – an experience for which I will be forever grateful.

1 Introduction

The goal of evolutionary game theory is to explain behavior in strategic settings – typically biological or economic – from the perspective of Darwinian evolution. An evolutionary system involves three basic elements: a strategic game, a population of players, and some mathematical conception of the evolution of strategic game-play throughout the population. Traditional economic game theory is easily adapted to a population of players; evolutionary game theory, then, imposes on this population some form of mathematical dynamics.

This thesis is motivated by a specific topic in evolutionary game theory: the coevolution of cooperation and costly punishment. This is a recent variation on a classic conundrum of economic game theory, namely the famous “Prisoner’s Dilemma,” and the evolution of cooperation. The aim of this thesis is to introduce the major results of evolutionary game theory for infinite populations, and to apply these results to a new problem stemming from recent evolutionary models of cooperation and punishment.

1.1 The Prisoner’s Dilemma

The Prisoner’s Dilemma is perhaps the most well-known problem in game theory. This famous moniker comes from the backstory of two apprehended partners-in-crime, each of whom must decide whether to keep quiet (cooper-

ate, “C”) or confess (defect, “D”). Two cooperating criminals get away with just a small fine, while two defectors both serve short prison sentences; but if just one of the two defects, the defector gets off “scot-free,” while the cooperator is sentenced to many years in prison. Put into the language of economic game theory, the Prisoner’s Dilemma is a two-player game specified by the following *payoff matrix*:¹

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} b-c & -c \\ b & 0 \end{pmatrix} \end{array} \quad \text{with } b > c > 0. \quad (1.1)$$

Intuitively, a selfish rational player should defect, since this always gives a higher payoff than cooperating, regardless of the opponent’s strategy. The game-theoretic concept of Nash equilibrium (discussed in §2.1) formalizes this intuition, and in general indicates the possible strategic outcomes of games played by rational individuals who have perfect knowledge of each other’s rationality. Thus, economic game theorists have long been fascinated by the implications of the Prisoner’s Dilemma – why does cooperation arise among individuals, groups, and societies, despite the fact that defection is optimal for selfish rational individuals?

A popular economic explanation is that cooperation is self-reinforced by the possibility for reciprocation, either through repeated game-play (Fudenberg and Maskin, 1990; Trivers, 1971) or the development of social reputations (Levine and Pesendorfer, 2007; Sigmund et al., 2001). Another explanation is the mechanism of costly punishment, whereby cooperators can pay a cost to punish defectors (for references, see §5). The mechanism of costly punishment as a means of stabilizing cooperation has drawn much attention in recent years, and serves as the motivation behind the evolutionary theory developed in this paper.

1.2 Evolution in Games

Evolutionary game theory offers us a different perspective from traditional economic game theory. Evolution, too, is a selfish process, but unlike economics it presupposes no degree of rationality, nor any form of conscious

¹A player employing a row- i strategy against a column- j opponent receives a payoff equal to matrix element (i, j) .

decision-making for that matter. This is not to say that evolutionary game theory obsoletes classical game theory. Evolution takes root in the payoffs of the underlying game, and thus builds upon, rather than replaces, the results of classical game theory.

Darwinian evolution presumes that members of a population are endowed with heritable genes that determine physical and behavioral traits. Traits determine varying levels of individual fitness, and natural selection favors the reproduction of the fit. Random mutations create genetic diversity in a population, which natural selection then narrows, so that the most successful traits prevail over time. The Darwinian mechanisms of selection and mutation make sense in a game-theoretic context whenever strategic behavior is thought to be inherent or immutable, for instance if we believe that a prisoner's likelihood to cooperate or defect is determined by some fixed personality traits or moral values.

Thus, evolutionary game theory assumes that individuals' strategies are in some way fixed, programmed, or otherwise rooted in heritable traits. Typically, a two-player game is played throughout the population in some manner – often in random pairs or in small groups – and thus the payoff to playing a given strategy is determined by the probability of facing each possible opposing strategy, i.e. the proportion of the population playing each strategy. The fitness of a strategy is defined as a function of its payoff (typically some simple monotonic map), and is thus also dependent on the strategic state of the population.

The selection mechanism of evolution can be formulated in many ways. One common mathematical setting for evolutionary theory is a continuous-time dynamical system acting on a “continuum population,” or an infinite population normalized to a unit mass. A selection mechanism is incorporated in the updating rule of the dynamical system. This selection can be thought of in a biological setting as the traditional Darwinian notion of reproduction/death of the fit/unfit, or in an economic setting as the growth/contraction of wealth or influence for the strategically successful/unsuccessful. Here we prefer a third interpretation, which better suits the social setting of human cooperation games: selection acts through strategic imitation of the fit by the unfit. Mutation is indirectly incorporated into infinite-population dynamics via the possibility of small perturbations in the strategic composition of the popu-

lation. Mutation is also addressed by the static notion of an evolutionary stable strategy, due to Maynard Smith and Price (1973), which in some sense refines the Nash equilibrium of economic game theory.

There are many alternative approaches to evolutionary modeling, including discrete-time difference equations (Nowak, 2006; Fudenberg and Levine, 1998), finite-population stochastic models (Traulsen and Hauert, 2008; Traulsen et al., 2008, 2007), as well as spatial models that structure finite populations using lattices, graphs, or sets (Nakamaru and Iwasa, 2006; Nowak, 2006; Nakamaru and Iwasa, 2005). In the scope of our study we will only deal with continuous-time infinite-population dynamics, but the compelled reader is encouraged to investigate alternative models, as evolutionary game theory is a deep and, so to speak, ever-evolving field of study.

1.3 Note to the Reader

We here discuss our reliance on specific elements of economic game theory, and briefly outline the evolutionary material to follow.

The bulk of game theory is formulated for *normal-form* (or *strategic-form*) games, namely those games with a finite set of players, a set of strategies for each player, and a well-defined function² that assigns to each player a payoff based on the strategic action of all players. Evolutionary game theorists almost always restrict their analysis to normal-form games, so that they may focus on the dynamics of evolution rather than the eccentricities of non-normal games.

Furthermore, it often makes sense to assume symmetry in an evolutionary context: a population that is *homogeneous*, in the sense that all individuals have the same choice of strategies and contingent payoffs, is consistent with a symmetric underlying game. Evolutionary theory for asymmetric games, such as bimatrix games³ and multipopulation models⁴ are well-studied, but here we are solely concerned with symmetric games like the Prisoner's Dilemma. Also, as mentioned in the previous section, it is com-

²This payoff function is typically required to be a viable *von Neumann-Morgenstern utility function*; see Osborne and Rubinstein (2002) for further description.

³See Hofbauer and Sigmund (1998), Ch. 10-11.

⁴See Weibull (1995), Ch. 5.

mon to analyze two-player games, which can then be played throughout a population in various ways (Hofbauer and Sigmund, 1998; Fudenberg and Levine, 1998; Weibull, 1995).

Thus, we begin in §2.1 by establishing the language and notation of game theory for two-player symmetric games. We center our discussion on the Nash equilibrium, which remains a key notion in evolutionary game theory, and prove the existence theorem for Nash equilibria.⁵ A reader well-versed in game theory may wish to skim this section. §2.2 introduces a static refinement of the Nash equilibrium known as the evolutionary stable strategy, and proves some results that will be used throughout the paper.

§3.1 and §3.2 introduce the formalism of game dynamics, solution mappings, and equilibria in dynamical systems, some of which will be rote for the reader familiar with the theory of differential equations. §3.3 is the crux of our investigation of the link between static and dynamic equilibria, and we here prove two major theorems of evolutionary game theory.

§4 introduces the two specific evolutionary models that we will use to analyze the game of cooperation and punishment in §5. §4.4.3 also presents a useful form of evolutionary equilibrium analysis different from the more formal methods described in §3.1.

§5 is the central application of the theory in §2-4. We here study the statics and dynamics of a game of cooperation and costly punishment that builds upon the Prisoner's Dilemma. This analysis is the original work of David G. Rand, Mayuko Nakamaru, Hisashi Ohtsuki, and myself, under the supervision of Martin Nowak and his Program for Evolutionary Dynamics. Our analysis is of interest in response to recent results, which have supported the stable evolution of cooperation when there is a strategic option to punish defectors. Using the tools of static game theory as well as the dynamics of two different evolutionary models, we show that when defectors are also allowed to “antisocially” punish cooperators, cooperation is no longer favored by evolution.

⁵For a more rigorous introduction to general game theory, including the Nash existence proof for general finite games, see Appendix A.

2 Static Game Analysis

There are certain evolutionary traits intrinsic to a game which we can analyze before imposing any dynamical mathematical structure. The identification of static equilibria – namely, Nash equilibria and evolutionary stable strategies – offers insight into the nature of a game, and establishes a benchmark for comparison with evolutionary analysis.

2.1 Symmetric Game Theory and Nash Equilibria

In a normal-form game between m players, each player i is associated with a pure strategy space \mathbb{S}_i and a payoff function $\pi_i : \times_{i=1}^m \mathbb{S}_i \rightarrow \mathbb{R}$, where the product space $\times_{i=1}^m \mathbb{S}_i$ contains all possible combinations of players' strategies. Thus a game is fully specified by the triple $\mathcal{G} = (m, \{\mathbb{S}_i\}, \{\pi_i\})$.

In a symmetric game, all m players share a pure strategy space \mathbb{S} , and each player's payoff function π_i must be invariant under any permutation δ of players and players' strategies: $\pi_i(s_1, \dots, s_m) = \pi_{\delta(i)}(s_{\delta(1)}, \dots, s_{\delta(m)})$. In other words, all payoffs depend only on which strategies are played – independent of which opponent plays each strategy – allowing us to ignore subscripts on symmetric payoff functions.

Supposing there are two players and a finite strategy space,⁶ say $|\mathbb{S}| = n$, the payoff function π is given by a real $n \times n$ payoff matrix A , where $\pi(s_i, s_j) = a_{ij}$, as in the Prisoner's Dilemma (1.1). A two-player symmetric game is then denoted $\mathcal{G} = (2, \mathbb{S}, A)$.

In game \mathcal{G} , a player can also employ a *mixed strategy* by playing each pure strategy $s_i \in \mathbb{S}$ with fixed probability. This creates a mixed-strategy space, or the space of all probability weightings x_1, \dots, x_n across n pure strategies. This mixed-strategy space is then isomorphic to the simplex $S_n \subset \mathbb{R}^n$ of dimension $(n - 1)$:

$$S_n = \{\mathbf{x} \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1, x_i \geq 0 \ \forall i\}. \quad (2.1)$$

⁶The existence of Nash equilibria is only guaranteed for games with finite strategy spaces, so this assumption will suffice for our purposes.

Henceforth we will treat this equivalence not as an isomorphism, but as a definition: mixed strategies are column-vectors in the simplex S_n , whose vertices are the pure strategies $\mathbb{S} = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$, the standard basis vectors of \mathbb{R}^n . Let us define, for each strategy \mathbf{x} , the subset of the simplex consisting of the vertices which span \mathbf{x} .

Definition. The *support* of a strategy $\mathbf{x} \in S_n$ is the set of pure strategies to which \mathbf{x} assigns nonzero probability weights:

$$\mathcal{S}(\mathbf{x}) = \{\mathbf{e}_i \in S_n \mid x_i > 0\}. \quad (2.2)$$

The support of a strategy is a fundamental notion with frequent applications throughout game theory.

The expected payoff of an \mathbf{x} -strategist playing against a \mathbf{y} -strategist is given by the probabilistic sum of payoffs over all pure strategies:

$$\sum_{i=1}^n \sum_{j=1}^n Pr(s_i)Pr(s_j)\pi(s_i, s_j) = \sum_{i,j} x_i y_j a_{ij} = \mathbf{x} \cdot \mathbf{A}\mathbf{y} \quad (2.3)$$

where (\cdot) denotes the standard inner product on \mathbb{R}^n . Note that (2.3) is consistent with our initial definition of pure-strategy payoffs, since playing strategy \mathbf{e}_i against \mathbf{e}_j gives payoff $\mathbf{e}_i \cdot \mathbf{A}\mathbf{e}_j = a_{ij}$.

It is important that the payoff function is a bilinear form on \mathbb{R}^n , i.e. that it is linear in each player's mixed strategy.⁷ We will use this fact in our analysis of the properties of Nash equilibria, as well as in the Nash existence proof. We now introduce the best-reply function, which lays the foundation for the Nash equilibrium.

Definition. Given an opponent's strategy $\mathbf{y} \in S_n$, a player's *best-reply* is given by

$$\beta(\mathbf{y}) = \{\mathbf{x} \in S_n \mid \mathbf{x} \cdot \mathbf{A}\mathbf{y} \geq \mathbf{e}_i \cdot \mathbf{A}\mathbf{y}, \forall \mathbf{e}_i \in \mathbb{S}\}. \quad (2.4)$$

Note that this definition permits multiple best-replies to a given strategy. Thus a player's best reply β is not a proper function, but rather a correspondence on S_n , or equivalently a set-valued function $\beta : S_n \rightarrow \mathcal{P}(S_n)$.⁸

⁷It is, in fact, always the case in any normal-form game that payoffs are multilinear in the players' mixed strategies; see (A.1) in Appendix A.

⁸Where $\mathcal{P}(A)$ denotes the set of all subsets of A .

A best reply is thus a strategy which maximizes a player’s payoff, taking the opponent’s strategy as a given. A Nash equilibrium occurs when each player’s strategy is a best reply to his opponent’s strategy: $(\mathbf{x}, \mathbf{y}) \in \beta(\mathbf{y}) \times \beta(\mathbf{x}) \subset S_n^2$. In other words, a Nash equilibrium is a set of strategies from which no player can profitably deviate. In an evolutionary context, we are particularly interested in those *symmetric* Nash equilibria (\mathbf{x}, \mathbf{x}) , in which each player employs the same equilibrium strategy. A symmetric Nash equilibrium can be represented by a single strategy, and thus corresponds to a single point in the simplex, which will translate nicely into a population state when we introduce evolutionary dynamics in §3. We now formalize this discussion:

Definition. A strategy $\mathbf{x} \in S_n$ is a *symmetric Nash equilibrium* if $\mathbf{x} \in \beta(\mathbf{x})$, i.e. if

$$\mathbf{x} \cdot A\mathbf{x} \geq \mathbf{e}_i \cdot A\mathbf{x} \quad \forall i \in \{1, \dots, n\}. \quad (2.5)$$

We say that the Nash equilibrium is *strict* if $\{\mathbf{x}\} = \beta(\mathbf{x})$ (i.e. if the above inequality holds strictly).

Note that the Nash equilibrium quite elegantly requires that each player check only for profitable pure-strategy deviations, because the payoff $\mathbf{y} \cdot A\mathbf{x}$ to a mixed-strategy \mathbf{y} is linear in the probability y_i assigned to each pure strategy. As another consequence of this linearity, a mixed-strategy Nash equilibrium \mathbf{x} must be indifferent among the pure strategies in its support $\mathcal{S}(\mathbf{x})$ (otherwise, a profitable deviation is achieved by decreasing the probability weight of an inferior support strategy). It follows that a mixed strategist in Nash equilibrium can always deviate to any probabilistic combination of his support strategies without altering his payoff, and thus a mixed strategy cannot be a strict Nash equilibrium (Fudenberg and Levine, 1998; Fudenberg and Tirole, 1991).

As mentioned in the introduction, Nash equilibria tell us which stable outcomes we can expect in a game played by perfectly rational agents with perfect knowledge of each other’s rationality. Even if perfect rationality is too much to ask of real-world game participants, it is nonetheless valuable to be able to identify these “ideal” equilibria in a game.

Example 2.1. Recall the Prisoner’s Dilemma game (1.1) from the introduction, where we discussed the fact that defection is the only strategy that

can arise in rational game-play. We can now justify this logic formally, by observing that defection (D) is the unique Nash equilibrium of the game. Furthermore, D is a strict Nash equilibrium, since $0 > -b$. \square

Example 2.2. We now introduce to the Prisoner’s Dilemma a second round of game-play, in which a player can choose to incur a cost $\alpha > 0$ in order to punish a defector by an amount $\beta > 0$. Adding a player’s payoffs from the two rounds gives the overall payoff matrix

$$\begin{array}{cc} & \begin{array}{c} CP \\ DN \end{array} \\ \begin{array}{c} CP \\ DN \end{array} & \begin{pmatrix} b - c & -c - \alpha \\ b - \beta & 0 \end{pmatrix} \end{array}$$

where “CP” denotes a cooperator who punishes defectors, and “DN” a defector who does not punish other defectors. While DN is still a strict Nash equilibrium (since $0 > -c - \alpha$), we see that CP is also strict Nash if punishment is sufficiently effective, i.e. if $\beta > c$ (and weak Nash if $\beta = c$). Thus, cooperation can become a rational strategy when there is the chance of punishment. We can also calculate any potential mixed-strategy Nash equilibrium: supposing a Nash strategist plays CP with probability x (hence DN with probability $1 - x$), we use the Nash condition of indifference among support strategies to compute,

$$\begin{aligned} x(b - c) + (1 - x)(-c - \alpha) &= x(b - \beta) \\ \implies x &= \frac{\alpha + c}{\alpha + \beta} \end{aligned}$$

which is between 0 and 1 as long as $\beta > c$. In other words, a mixed Nash equilibrium $\left(\frac{\alpha+c}{\alpha+\beta}, 1 - \frac{\alpha+c}{\alpha+\beta}\right)$ exists as long as CP itself is Nash. \square

Defection in the traditional Prisoner’s Dilemma of Example 2.1 is in some sense the purest form of strict Nash equilibrium, since cooperation is “strictly dominated,” i.e. never a best reply to an opponent’s strategy. The notion of strategic dominance is useful in both static and evolutionary game theory, and thus we briefly introduce dominance relations before proceeding to the existence theorem for Nash equilibria.⁹

⁹As a prelude to the evolutionary results of §3.3, we need only introduce strict strategic dominance, although weak dominance is also an important concept in classical game theory. See Fudenberg and Tirole (1991) for a thorough treatment of strategic dominance.

Definition. A strategy $\mathbf{x} \in S_n$ is *dominated* if there exists some $\mathbf{x}' \in S_n$ such that

$$\mathbf{x}' \cdot A\mathbf{y} \geq \mathbf{x} \cdot A\mathbf{y} \quad \forall \mathbf{y} \in S_n. \quad (2.6)$$

We say that \mathbf{x} is *strictly* dominated if the above inequality holds strictly; otherwise, \mathbf{x} is *weakly* dominated.

This definition leads to a process of narrowing a player's rational strategy space by iteratively “deleting” strictly dominated pure strategies (we will see shortly why weakly dominated strategies are not eliminated). For instance, beginning with a game $\mathcal{G} = (2, \mathbb{S}, A)$, a player first rules out all strictly dominated strategies to obtain a rational subgame $\mathcal{G}^1 = (2, \mathbb{S}^1, A^1)$. Then it is possible that there are strategies in \mathcal{G}^1 , undominated in the original game \mathcal{G} , which are now strictly dominated in \mathcal{G}^1 by way of the removal of certain strategies from \mathcal{G} . This leads to a twice-trimmed game \mathcal{G}^2 , and so on, until $\mathcal{G}^{T+1} = \mathcal{G}^T$ for some T (which occurs when no remaining strategies are strictly dominated). The strategies \mathbb{S}^T of this final subgame are the pure strategies which are not *iteratively strictly dominated* in the game \mathcal{G} .

Note that any pure-strategy Nash equilibrium is undominated, hence contained in the set \mathbb{S}^T . Moreover, a Nash strategy cannot contain any strictly dominated strategy in its support – by the linearity of the payoff function, it is profitable for a mixed-strategist to transfer probability weight from a dominated strategy to one of its dominators. Thus, we have proven the following result:

Proposition 2.1. No strategy in the support of a symmetric Nash equilibrium is iteratively strictly dominated.

Proposition 2.1 allows us to search for Nash equilibria after reducing the game \mathcal{G} to its final trimmed subgame \mathcal{G}^T , without losing any potential equilibria. This result does not extend to weak domination: a weak Nash equilibrium can contain in its support (or can itself be) a weakly dominated strategy.¹⁰ The evolutionary equilibrium we introduce in §2.2 is a refinement of the Nash equilibrium, and cannot contain any weakly dominated strategies in its support.

¹⁰We will see an example of this, as well as of strictly dominated strategies, in our analysis in §5.

We are now ready to prove existence of symmetric Nash equilibrium in symmetric games.¹¹ The proof hinges on the Kakutani Fixed Point Theorem, a generalization of the Brouwer Fixed Point Theorem of topology.¹²

Theorem (Kakutani Fixed Point Theorem). *Let S be a nonempty, convex, and compact subset of a Euclidean space, and $\beta : S \rightarrow \mathcal{P}(S)$ a set-valued function. If β has convex nonempty images and a closed graph,¹³ then it has a fixed point $x \in \beta(x)$.*

Theorem (Existence of Symmetric Nash Equilibria). *Every finite symmetric two-player game has a symmetric Nash equilibrium.*

Proof. By definition (2.5), a fixed point of the best reply map (2.4) is a symmetric Nash equilibrium of the game. Thus it suffices to show that for a finite game, $\beta : S_n \rightarrow \mathcal{P}(S_n)$ satisfies the conditions of Kakutani's theorem.

- (a) *The domain of β is a nonempty, convex, and compact subset of a Euclidean space.*

It is clear from (2.1) that the simplex S_n is a nonempty, convex, and compact (closed and bounded) subset of the Euclidean space \mathbb{R}^n . Thus the domain of β meets the requirements.

- (b) *$\beta(\mathbf{y})$ is nonempty, $\forall \mathbf{y} \in S_n$.*

Since a player's payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{y}$ is a linear (hence continuous) function of his strategy $\mathbf{x} \in S_n$, it attains a maximum on the compact set S_n . Namely, there is some $\mathbf{x}^* \in S_n$ such that $\mathbf{x}^* \cdot \mathbf{A}\mathbf{y} \geq \mathbf{x} \cdot \mathbf{A}\mathbf{y}$ for all $\mathbf{x} \in S_n$. Thus, by definition (2.4), $\mathbf{x}^* \in \beta(\mathbf{y})$.

- (c) *$\beta(\mathbf{y})$ is convex, $\forall \mathbf{y} \in S_n$.*

For a given $\mathbf{y} \in S_n$, let $\mathbf{x}^1, \dots, \mathbf{x}^k \in \beta(\mathbf{y})$ be best replies, i.e. for all $j \in \{1, \dots, k\}$, \mathbf{x}^j achieves the maximal payoff against \mathbf{y} : $\mathbf{x}^j \cdot \mathbf{A}\mathbf{y} = \lambda \geq \mathbf{x} \cdot \mathbf{A}\mathbf{y}$, $\forall \mathbf{x} \in S_n$. Then for any list of nonnegative weights $\alpha_1, \dots, \alpha_k$ with unit sum, the corresponding convex combination of best replies also achieves this maximal payoff, and is thus itself a best reply: $(\sum_{j=1}^k \alpha_j \mathbf{x}^j) \cdot \mathbf{A}\mathbf{y} = \sum_j \alpha_j (\mathbf{x}^j \cdot \mathbf{A}\mathbf{y}) = \sum_j \alpha_j \lambda = \lambda$.

¹¹See Appendix A for proof of the Nash Existence Theorem for arbitrary finite games, as in Nash (1951).

¹²We here state but not prove Kakutani's theorem; see Kakutani (1941) for proof.

¹³ β has a *closed graph* if: $(\{\mathbf{x}^j\}, \{\mathbf{y}^j\}) \rightarrow (\mathbf{x}, \mathbf{y})$, with $\mathbf{x}^j \in \beta(\mathbf{y}^j) \Rightarrow \mathbf{x} \in \beta(\mathbf{y})$.

(d) β has a closed graph.

Recall that β has a closed graph if: $(\{\mathbf{x}^j\}, \{\mathbf{y}^j\}) \rightarrow (\mathbf{x}, \mathbf{y})$, with $\mathbf{x}^j \in \beta(\mathbf{y}^j) \Rightarrow \mathbf{x} \in \beta(\mathbf{y})$. We prove this implication by contradiction, using the following “three-epsilon” argument. Suppose that $(\{\mathbf{x}^j\}, \{\mathbf{y}^j\}) \rightarrow (\mathbf{x}, \mathbf{y})$ but $\mathbf{x} \notin \beta(\mathbf{y})$. Then there exists an $\varepsilon > 0$ and an $\mathbf{x}' \in S_n$ such that $\mathbf{x}' \cdot \mathbf{A}\mathbf{y} > \mathbf{x} \cdot \mathbf{A}\mathbf{y} + 3\varepsilon$. By the continuity of the bilinear payoff function, and the convergence of $(\{\mathbf{x}^j\}, \{\mathbf{y}^j\}) \rightarrow (\mathbf{x}, \mathbf{y})$, we also have for j sufficiently large: $\mathbf{x}' \cdot \mathbf{A}\mathbf{y}^j > \mathbf{x}' \cdot \mathbf{A}\mathbf{y} - \varepsilon$, and $\mathbf{x} \cdot \mathbf{A}\mathbf{y} > \mathbf{x}^j \cdot \mathbf{A}\mathbf{y}^j - \varepsilon$. Combining the three inequalities, we get:

$$\mathbf{x}' \cdot \mathbf{A}\mathbf{y}^j > \mathbf{x}' \cdot \mathbf{A}\mathbf{y} - \varepsilon > \mathbf{x} \cdot \mathbf{A}\mathbf{y} + 2\varepsilon > \mathbf{x}^j \cdot \mathbf{A}\mathbf{y}^j + \varepsilon.$$

This contradicts the fact that $\mathbf{x}^j \in \beta(\mathbf{y}^j)$, and so it must be the case that β has a closed graph.

Thus, by the Kakutani fixed point theorem, the best-reply mapping of any finite two-player symmetric game has a fixed point $\mathbf{x} \in \beta(\mathbf{x})$, which constitutes a symmetric Nash equilibrium of the game. □

Note that the existence theorem does not guarantee the existence of pure-strategy equilibria in an arbitrary game (although there are pure-strategy equilibria in the cooperation games considered in this paper). For instance, the proverbial “Rock-Paper-Scissors” game has a unique mixed-strategy Nash equilibrium, but no pure equilibria (Fudenberg and Tirole, 1991).

The beauty of the Nash equilibrium lies in its simple formulation, its universality, and its ease of computation. There have been many attempts at modifying the Nash equilibrium, but none have displaced or undermined its original formulation. That being said, there is certainly room for refinement. One obvious shortcoming of the Nash equilibrium is that it tells us nothing of how it arises. Evolutionary game theory makes strides in addressing this question, as well as in addressing the reciprocal question of how systems evolve away from Nash equilibria – for instance, toward cooperation in the Prisoner’s Dilemma.

To consider another shortcoming of the Nash equilibrium, recall that only a pure strategy can be a strict Nash equilibrium. This forces us to turn to the weak Nash equilibrium as our central notion of mixed equilibrium. However,

in an evolutionary setting, a population of individuals employing a weak Nash strategy is vulnerable to invasion by a mutant strategy that earns the same payoff against the weak Nash strategy as the weak Nash strategy earns from itself. Then, if the mutant strategy outperforms the Nash strategy against other mutants, natural selection will favor the mutant strategy and cause it to take over the population. We now introduce a key form of evolutionary equilibrium, which permits a mixed strategy to be employed throughout a stable, uninvadeable population – a powerful compromise between the strict and weak Nash conditions.

2.2 Evolutionarily Stable Strategies

The notion of evolutionarily stable strategies in game theory was first introduced by Maynard Smith and Price (1973), and soon after formalized by Maynard Smith (1974). A strategy is said to be evolutionarily stable if, when a population uniformly adopts the strategy, any small group of deviant strategists fares worse than the stable strategic majority.

Definition. $\mathbf{x} \in S_n$ is an *evolutionarily stable strategy* (ESS) if for every strategy $\mathbf{y} \in S_n$, $\mathbf{y} \neq \mathbf{x}$, there exists an *invasion barrier* $\bar{\varepsilon}(\mathbf{y}) \in (0, 1)$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$:

$$\mathbf{x} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) > \mathbf{y} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}). \quad (2.7)$$

It is a well-known result (which we will soon prove) that an ESS must in fact have a uniform invasion barrier, so that the above definition holds for some minimal $\bar{\varepsilon}(\mathbf{y}) = \bar{\varepsilon} > 0$ throughout the simplex. In other words, there is some nonzero threshold share of the population, below which any proportion of invaders employing a foreign strategy is repelled.

Using the linearity of A to rearrange (2.7), we get:

$$(1 - \varepsilon)(\mathbf{x} \cdot A\mathbf{x} - \mathbf{y} \cdot A\mathbf{x}) + \varepsilon(\mathbf{x} \cdot A\mathbf{y} - \mathbf{y} \cdot A\mathbf{y}) > 0. \quad (2.8)$$

For sufficiently small ε , (2.8) holds whenever either the first term is positive, or the first term is zero and the second term is positive. Thus we have constructed an equivalent definition of ESS:

Proposition 2.2. A strategy $\mathbf{x} \in S_n$ is an ESS if and only if the following two conditions are satisfied.

(a) *Equilibrium Condition:*

$$\mathbf{x} \cdot A\mathbf{x} \geq \mathbf{e}_i \cdot A\mathbf{x} \quad \forall i \in \{1, \dots, n\}. \quad (2.9)$$

(b) *Stability Condition:*

$$\text{If } \mathbf{y} \neq \mathbf{x} \text{ and } \mathbf{y} \cdot A\mathbf{x} = \mathbf{y} \cdot A\mathbf{y}, \text{ then } \mathbf{x} \cdot A\mathbf{y} > \mathbf{y} \cdot A\mathbf{y}. \quad (2.10)$$

As was the case when we defined symmetric Nash equilibria in (2.5), the definition of ESS by Proposition 2.2 only requires comparison with pure strategies (again, since the payoff function (2.3) is linear the probabilities a mixed strategy assigns to the pure strategies). While definition (2.7) is more directly related to our original conception of the ESS, Proposition 2.2 is of greater use in computing the ESS's of a game, as it requires finitely many comparisons between entries of the payoff matrix A .

Notice that the Equilibrium Condition (2.9) implies that any ESS \mathbf{x} is a Nash equilibrium. Moreover, if a strategy \mathbf{x} is strict Nash in the sense of (2.5), then the inequality in (2.9) holds strictly (making Stability Condition (2.10) superfluous), so \mathbf{x} is an ESS. This gives us Proposition 2.3.

Proposition 2.3. The following chain of implications holds for a strategy in a symmetric two-player game:

$$\text{Strict Nash Equilibrium} \Rightarrow \text{ESS} \Rightarrow \text{Nash Equilibrium}.$$

The intuition behind these implications is clear, if we observe that a population of strategic composition $(1 - \varepsilon)\mathbf{x} + \varepsilon\mathbf{y}$ is mathematically equivalent to a mixed strategy which plays \mathbf{x} with probability $(1 - \varepsilon)$ and \mathbf{y} with probability ε – i.e., a slight strategic deviation from \mathbf{x} . Then, an ESS can be thought of as a strategy from which “small” strategic deviations are detrimental; this is clearly weaker than a strict Nash strategy, from which any deviation is detrimental. The fact that ESS implies weak Nash is best intuited by contrapositive: if a strategy is not weak Nash, then there exists some profitable strategic alternative, in which case a “small” strategic deviation towards this profit cannot be detrimental, and thus the strategy is not an ESS.

Recalling the discussion of strategic dominance from §2.1, we observe that the ESS further distinguishes itself from a Nash equilibrium as a strategy that

cannot be weakly dominated. Thus, for the ESS we can prove a stronger version of Proposition 2.1: while a Nash equilibrium cannot have any iteratively strictly dominated support strategies, an ESS cannot even possess any weakly dominated support strategies.

Proposition 2.4. No strategy in the support of an ESS is weakly dominated.

Proof. Let $\mathbf{x} \in S_n$ be an ESS. Suppose some support strategy $\mathbf{e}_i \in \mathcal{S}(\mathbf{x})$ is weakly dominated (cf. (2.6)), i.e. there exists some $\mathbf{x}' \in S_n$ such that $\mathbf{x}' \cdot A\mathbf{y} \geq \mathbf{e}_i \cdot A\mathbf{y}$, $\forall \mathbf{y} \in S_n$. We now construct a strategy that will help lead us to a contradiction: let $\mathbf{z} = \mathbf{x} - x_i \mathbf{e}_i + x_i \mathbf{x}'$, which clearly lies in the simplex. It follows from this definition that \mathbf{z} weakly dominates \mathbf{x} , since for any strategy \mathbf{y} ,

$$\begin{aligned} \mathbf{z} \cdot A\mathbf{y} - \mathbf{x} \cdot A\mathbf{y} &= \mathbf{x} \cdot A\mathbf{y} - x_i(\mathbf{e}_i \cdot A\mathbf{y}) + x_i(\mathbf{x}' \cdot A\mathbf{y}) - \mathbf{x} \cdot A\mathbf{y} \\ &= x_i(\mathbf{x}' \cdot A\mathbf{y} - \mathbf{e}_i \cdot A\mathbf{y}) \geq 0 \end{aligned}$$

with the final inequality following from the weak domination of \mathbf{e}_i by \mathbf{x}' . Letting $\mathbf{y} = \mathbf{x}$ above, we have $\mathbf{z} \cdot A\mathbf{x} \geq \mathbf{x} \cdot A\mathbf{x}$. Since ESS condition (2.9) implies $\mathbf{x} \cdot A\mathbf{x} \geq \mathbf{z} \cdot A\mathbf{x}$, we have that in fact $\mathbf{x} \cdot A\mathbf{x} = \mathbf{z} \cdot A\mathbf{x}$. Now, ESS condition (2.10) gives $\mathbf{x} \cdot A\mathbf{z} > \mathbf{z} \cdot A\mathbf{z}$. But, letting $\mathbf{y} = \mathbf{z}$ above, we also have that $\mathbf{z} \cdot A\mathbf{z} \geq \mathbf{x} \cdot A\mathbf{z}$, which is a contradiction. Thus, our ESS \mathbf{x} cannot have any weakly dominated strategy in its support. \square

We now introduce another useful property of the support of an ESS.

Proposition 2.5. If \mathbf{x} is an ESS and $\mathcal{S}(\mathbf{y}) \subset \mathcal{S}(\mathbf{x})$ for some strategy $\mathbf{y} \neq \mathbf{x}$, then \mathbf{y} is not a Nash equilibrium.

Proof. If \mathbf{x} is an ESS, then it is Nash by Proposition 2.3. Recall from the discussion in §A that a mixed strategy Nash equilibrium must be indifferent to all linear combinations of its support strategies. Thus $\mathbf{y} \cdot A\mathbf{x} = \mathbf{x} \cdot A\mathbf{x}$. Then, by the ESS Stability Condition (2.10), $\mathbf{x} \cdot A\mathbf{y} > \mathbf{y} \cdot A\mathbf{y}$, i.e. \mathbf{y} is not a Nash equilibrium. \square

Proposition 2.5 offers two interesting corollaries that shed light on the possible sets of ESS in a game.

Corollary 2.5.1. A finite game has finitely many ESS's.

Corollary 2.5.2. If a game has an ESS in the interior of the simplex, then it is the unique Nash equilibrium (hence ESS) of the game.

Corollary 2.5.1 relies on the fact that there are finitely many support sets (cf. (2.2)) in a finite game, while Corollary 2.5.2 follows directly from the original proposition.

A generalization of the global ESS of Corollary 2.5.2 is the notion of local superiority, which occurs when an ESS dominates some region of neighboring strategies in the simplex. We will prove that, in fact, local superiority is a defining property of all ESS's. But first, we must establish the existence of uniform invasion barriers (cf. (2.7)) for any ESS in symmetric two-player games. This result was first established by Vickers and Cannings (1987), although we also derive elements of our proof here from Weibull (1995) and Hofbauer and Sigmund (1998).

Theorem 2.1. *Every ESS has a uniform nonzero invasion barrier.*

Proof. Given an ESS \mathbf{x} , we let $\bar{\varepsilon}(\mathbf{y})$ denote the least upper bound invasion barrier for a mutant strategy \mathbf{y} , namely, the least upper bound on all initial proportions of \mathbf{y} such that \mathbf{y} is eliminated from a population predominated by \mathbf{x} . We seek to prove that $\bar{\varepsilon}(\mathbf{y})$ attains a positive minimal value on the set of all possible mutants. The difficulty of this proof results from the fact that this set of mutants is not compact, since \mathbf{x} can invade itself. Thus, we must prove that $\bar{\varepsilon}(\mathbf{y})$ has a positive minimum on the ‘‘punctured simplex,’’ $S'_n = S_n \setminus \{\mathbf{x}\}$.

We begin by converting the ESS definition (2.7) into a map $f_{\mathbf{x}} : [0, 1] \times S_n \rightarrow \mathbb{R}$ given by

$$\begin{aligned} f_{\mathbf{x}}(\varepsilon, \mathbf{y}) &= \mathbf{x} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) - \mathbf{y} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) \\ &= (\mathbf{x} - \mathbf{y}) \cdot A\mathbf{x} - \varepsilon(\mathbf{x} - \mathbf{y}) \cdot A(\mathbf{x} - \mathbf{y}). \end{aligned}$$

By definition, (2.7) holds exactly when $f_{\mathbf{x}}(\varepsilon, \mathbf{y}) > 0$. Thus, we can formulate the invasion barrier $\bar{\varepsilon}(\mathbf{y})$ of \mathbf{x} in terms of $f_{\mathbf{x}}$:

$$\bar{\varepsilon}(\mathbf{y}) = \sup\{\delta \in [0, 1] \mid f_{\mathbf{x}}(\varepsilon, \mathbf{y}) > 0, \forall \varepsilon \in (0, \delta)\}.$$

Observe that, if $f_{\mathbf{x}}(\cdot, \mathbf{y})$ has a unique zero ε_0 , then $\bar{\varepsilon}(\mathbf{y}) = \varepsilon_0$ by definition; if $f_{\mathbf{x}}(\cdot, \mathbf{y})$ is positive-definite, then (2.7) holds for all $\varepsilon \in [0, 1]$, so we have the

maximal invasion barrier $\bar{\varepsilon}(\mathbf{y}) = 1$. We now proceed to prove that $\bar{\varepsilon}(\mathbf{y})$ has a minimal value on the boundary of the punctured simplex, before showing that this minimum holds on all of S'_n .

Let $\mathcal{S} \equiv \mathcal{S}(\mathbf{x})$ be the support of \mathbf{x} (cf. (2.2)), recall that \mathbf{x} must be indifferent among support strategies, since an ESS is Nash by Proposition 2.3. Then, for any $\mathbf{y} \neq \mathbf{x}$ on the boundary face containing \mathbf{x} , we have $\mathcal{S}(\mathbf{y}) \subset \mathcal{S}$, so (2.7) holds (i.e. $f_{\mathbf{x}}(\varepsilon, \mathbf{y}) > 0$) for all ε . Thus, $\bar{\varepsilon}(\mathbf{y}) = 1$ on the boundary face containing \mathbf{x} .

We now observe that, for all \mathbf{y} with $\mathcal{S}(\mathbf{y}) \not\subset \mathcal{S}$, $f_{\mathbf{x}}(\varepsilon, \mathbf{y})$ is not the zero map. To see this, first let

$$\mathbf{e}_i \cdot A\mathbf{y} = (A\mathbf{y})_i \begin{cases} = \lambda & \text{for } \mathbf{e}_i \in \mathcal{S} \\ \leq \lambda - \mu < \lambda & \text{for } \mathbf{e}_i \notin \mathcal{S} \end{cases}$$

since \mathbf{x} is an ESS, which must prefer its own support strategies. Then:

$$\begin{aligned} (\mathbf{x} - \mathbf{y}) \cdot A\mathbf{x} &= \sum_{\mathbf{e}_i \in \mathcal{S}} (x_i - y_i)(A\mathbf{x})_i - \sum_{\mathbf{e}_i \notin \mathcal{S}} y_i(A\mathbf{x})_i \\ &\geq (1 - \sum_{\mathbf{e}_i \in \mathcal{S}} y_i)\lambda - \sum_{\mathbf{e}_i \notin \mathcal{S}} y_i(\lambda - \mu) = \mu \sum_{\mathbf{e}_i \notin \mathcal{S}} y_i > 0. \end{aligned}$$

Thus, for fixed \mathbf{y} with $\mathcal{S}(\mathbf{y}) \not\subset \mathcal{S}$, $f_{\mathbf{x}}(\varepsilon, \mathbf{y})$ has at most one zero in ε . Let us now consider all $\mathbf{y} \in S'_n$ which lie in the union of all boundary faces that do not contain \mathbf{x} : $\mathcal{Z}_{\mathbf{x}} = \{\mathbf{y} \in S'_n \mid y_i = 0 \text{ for some } \mathbf{e}_i \in \mathcal{S}\}$. For fixed $\mathbf{y} \in \mathcal{Z}_{\mathbf{x}}$, if $f_{\mathbf{x}}(\varepsilon, \mathbf{y})$ does not have a zero, then (2.7) holds for all ε , so once again $\bar{\varepsilon}(\mathbf{y}) = 1$. Otherwise, $f_{\mathbf{x}}(\varepsilon_o, \mathbf{y}) = 0$ for $\varepsilon_o = (\mathbf{x} - \mathbf{y}) \cdot A\mathbf{x} / (\mathbf{x} - \mathbf{y}) \cdot A(\mathbf{x} - \mathbf{y}) > 0$. We have herein fully specified $\bar{\varepsilon} : \mathcal{Z}_{\mathbf{x}} \rightarrow \mathbb{R}$ as

$$\bar{\varepsilon}(\mathbf{y}) = \begin{cases} \varepsilon_o & \text{if } \varepsilon_o \in (0, 1) \\ 1 & \text{otherwise.} \end{cases}$$

Thus, $\bar{\varepsilon}(\mathbf{y})$ is a positive-definite continuous function on compact domain $\mathcal{Z}_{\mathbf{x}}$, so it has a positive minimum value: $\min_{\mathcal{Z}_{\mathbf{x}}} \bar{\varepsilon}(\mathbf{y}) = \bar{\varepsilon} > 0$. Then, the minimum of $\bar{\varepsilon}(\mathbf{y})$ on $bd(S'_n)$ is just $\min\{\bar{\varepsilon}, 1\}$.

We now extend this result to all of the punctured simplex: since \mathbf{x} and

\mathcal{Z}_x together span S_n , we can write any point $\mathbf{y}' \in S'_n$ as $\mathbf{y}' = \alpha\mathbf{y} + (1 - \alpha)\mathbf{x}$ for some $\mathbf{y} \in \mathcal{Z}_x$ and $\alpha \in (0, 1]$. But then

$$\begin{aligned} f_x(\varepsilon, \mathbf{y}') &= \mathbf{x} \cdot A(\alpha\varepsilon\mathbf{y} + (1 - \alpha\varepsilon)\mathbf{x}) - (\alpha\mathbf{y} - (1 - \alpha)\mathbf{x}) \cdot A(\alpha\varepsilon\mathbf{y}' + (1 - \alpha\varepsilon)\mathbf{x}) \\ &= \alpha f_x(\alpha\varepsilon, \mathbf{y}). \end{aligned}$$

Therefore, for any $\mathbf{y}' \in S'_n$, $\bar{\varepsilon}(\mathbf{y}') = \min\{\bar{\varepsilon}/\alpha, 1\} \geq \bar{\varepsilon} > 0$, which proves the existence of a positive invasion barrier that holds uniformly for all mutant strategies. □

We are now equipped to formulate the ESS in yet another way: as a locally superior strategy (Vickers and Cannings, 1987; Hofbauer et al., 1979). This characterization of the ESS will prove useful in relating static and dynamical equilibria in §4.2.¹⁴

Definition. A strategy $\mathbf{x} \in S_n$ is *locally superior* if there exists some neighborhood U of \mathbf{x} such that: $\mathbf{x} \cdot A\mathbf{y} > \mathbf{y} \cdot A\mathbf{y}$, $\forall \mathbf{y} \in U \cap S_n$, $\mathbf{y} \neq \mathbf{x}$.

Proposition 2.6. A strategy is an ESS if and only if it is locally superior.

Proof. Let $\mathbf{x} \in S_n$ be an ESS. From Theorem 2.1, there exists an $\bar{\varepsilon} > 0$ so that condition (2.7) holds uniformly for all $\varepsilon \in (0, \bar{\varepsilon})$ and all $\mathbf{y} \in S_n$, $\mathbf{y} \neq \mathbf{x}$. We now multiply (2.7) by ε and add $(1 - \varepsilon)\mathbf{x} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x})$ to both sides, to get:

$$\mathbf{x} \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) > (\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) \cdot A(\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}).$$

Thus, \mathbf{x} is locally dominant in some neighborhood $U = \{\varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x} \mid \forall \mathbf{y} \in S_n, \forall \varepsilon \in (0, \bar{\varepsilon})\}$. Since all steps in this proof are reversible, we have proven the desired result. □

Proposition 2.6 is of interest even beyond its application to static and dynamic equilibria. It offers an alternative interpretation of the ESS as a best reply to neighboring strategies in the simplex. We once again invoke the mathematical equivalence of strategic population states and mixed strategies to gain intuition: playing a game against a random opponent in a population

¹⁴While we only need to prove this result for symmetric two-player games here, a more general proof can be found in Zeeman (1980).

of ε \mathbf{y} -invaders and $(1 - \varepsilon)$ \mathbf{x} -incumbents gives an expected payoff equivalent to playing against a fixed opponent whose strategy is shifted an ε th of the way from \mathbf{x} to \mathbf{y} in the simplex, and who is thus the incumbent's "neighbor" in strategy space.

Now having established the key notions of Nash equilibrium and ESS from static game theory, we introduce the foundations of evolutionary dynamics.

3 Game Dynamics in Continuum Populations

A static game becomes dynamic when it is mapped to a system of differential equations, to form a dynamical system. The specific form of the mapping derives from the assumptions made on how the game is played throughout the population – in random pairs, in small groups, etc. In this section, we introduce the necessary mathematical framework of dynamical systems, before exploring the connection between such a system and an underlying two-player symmetric game, $\mathcal{G} = (2, \mathbb{S}, A)$.

3.1 Game Dynamics

Before defining game dynamics, we must define a population on which the dynamics will operate. As discussed in the introduction, we here consider an infinite population, which we treat as a continuum of unit mass. A standard choice for the population’s strategy space is the set of pure strategies \mathbb{S} of the underlying game \mathcal{G} , as in Weibull (1995). Thus *population states* are analogous to (pure or mixed) strategies of the underlying game: a state of the population is specified by a vector \mathbf{x} in the unit simplex, where x_i is the share of the population playing the pure strategy \mathbf{e}_i . The notion of the support of a strategy (cf. (2.2)) translates directly to the support of a population state, now representing the set of pure strategies being employed in the population state.

Hofbauer and Sigmund (1998) takes a different approach, choosing for the population’s strategy space a fixed number of possibly mixed strategies \mathbf{x}^i from the original game. While this framework allows for more general strategy spaces, it is in fact mathematically equivalent to Weibull’s pure-strategy population: simply redefine the underlying game payoff matrix as $A' = (\mathbf{x}^i \cdot A\mathbf{x}^j)$, and return to the simplex by mapping $\mathbf{x}^i \rightarrow \mathbf{e}_i$. Thus, without loss of generality, we can take Weibull’s pure-strategy approach.

Definition. A *game dynamics* is defined by a system of n time-independent ordinary differential equations given by

$$\dot{x}_i = x_i g_i(x_1, \dots, x_n) \quad i = 1, \dots, n \quad (3.1)$$

where $\mathbf{x} \in S_n$ is a population state, and $\mathbf{g} : X \rightarrow \mathbb{R}^n$ is the *growth rate function*, defined on some open domain $X \subset \mathbb{R}^n$ which contains S_n .

Note that each component $g_i(\mathbf{x}) = \dot{x}_i/x_i$ is just the growth rate of strategy \mathbf{e}_i in population state \mathbf{x} . Presumably, \mathbf{g} is related to the payoff matrix A of \mathcal{G} (though we will not explore possible functional forms of \mathbf{g} until we first prove some general results).

Definition. A *solution mapping* for the game dynamics (3.1) is a function $\boldsymbol{\xi} : \mathbb{R} \times X \rightarrow X$ such that, for all $\mathbf{x} \in X$,

- (a) $\boldsymbol{\xi}(0, \mathbf{x}) = \mathbf{x}$
- (b) $\dot{\boldsymbol{\xi}}_i(t, \mathbf{x}) = \xi_i(t, \mathbf{x})g_i(\boldsymbol{\xi}(t, \mathbf{x}))$, $\forall t \in \mathbb{R}$, for $i = 1, \dots, n$.

Observe that in our case $\boldsymbol{\xi}$ must be time-invariant, in the sense that:

$$\boldsymbol{\xi}(t, \boldsymbol{\xi}(t', \mathbf{x})) = \boldsymbol{\xi}(t + t', \mathbf{x}) \quad \forall \mathbf{x} \in S_n, \forall t, t' \in \mathbb{R} \quad (3.2)$$

because of the fact that (3.1) is time-independent. In evolutionary theory, we are more concerned with the long-run locality of a solution than with its specific trajectory, for it is the outcome – not the pathway – of evolution that is of chief interest.¹⁵ Thus, our dynamical aim is analogous to finding Nash equilibria in static game theory: we seek to identify stable, sustainable behaviors. In a dynamical system, stable behaviors arise in the form of asymptotic states, which rely on the existence, uniqueness, and continuity of solution trajectories. To guarantee these conditions, we must impose the following condition of “regularity.”

Definition. A game dynamics is *regular* if its growth-rate function $\mathbf{g} : X \rightarrow \mathbb{R}^n$ is Lipschitz continuous¹⁶ and satisfies

$$\mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = 0 \quad (3.3)$$

for all $\mathbf{x} \in S_n \subset X$.

Note that continuous differentiability implies Lipschitz continuity, since a continuous derivative is bounded on a compact set, and thereby satisfies the condition for Lipschitz continuity. Nearly all well-studied evolutionary models have growth rates which are at least continuously differentiable, and thus the key condition for regularity is typically (3.3), not Lipschitz continuity.

¹⁵Our preference for asymptotics over trajectories is in part a “de facto” determination, since evolutionary models for cooperation are more conceptual than rigorously scientific.

¹⁶Namely: $\forall C \subset X$ compact, $\exists \lambda \in \mathbb{R}$ s.t. $\|\mathbf{g}(\mathbf{x}) - \mathbf{g}(\mathbf{y})\| \leq \lambda \|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y} \in C$.

Proposition 3.1. Any regular game dynamics allows a solution mapping $\boldsymbol{\xi}$ which defines a unique path $\boldsymbol{\xi}(\cdot, \mathbf{x}) : \mathbb{R} \rightarrow S_n$ through each initial state $\mathbf{x} \in S_n$. Moreover, $\boldsymbol{\xi}$ is continuous in $t \in \mathbb{R}$ and in $\mathbf{x} \in S_n$.

Proposition 3.1 follows immediately from the classical result of the *Picard-Lindelöf Theorem*.¹⁷ Lipschitz continuity allows us to invoke the theorem, and condition (3.3) allows us to constrain our solutions to the simplex. In fact, the simplex is not the only set left invariant by our game dynamics, as the following proposition will illustrate.

Proposition 3.2. The simplex S_n is invariant under regular game dynamics, in the sense that if $\mathbf{x}^\circ \in S_n$ then $\boldsymbol{\xi}(t, \mathbf{x}^\circ) \in S_n$ for all $t \in \mathbb{R}$. Moreover, the interior of the simplex is invariant, as is each boundary face, and each vertex.

Proof. We begin by proving that the simplex is invariant for arbitrary initial state $\mathbf{x}^\circ \in S_n$. It is clear that regular game dynamics (3.1) preserves the sum of the coordinates of any vector $\boldsymbol{\xi}(t, \mathbf{x}^\circ) = \mathbf{x} \in X$:

$$\frac{d}{dt} \left(\sum_{i=1}^n \xi_i(t, \mathbf{x}^\circ) \right) = \sum_{i=1}^n \dot{x}_i = \sum_{i=1}^n x_i g_i(\mathbf{x}) = \mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = 0.$$

Thus it remains to show that each coordinate $\xi_i(t, \mathbf{x}^\circ)$ is nonnegative for all times $t \in \mathbb{R}$. Suppose that $\boldsymbol{\xi}(\cdot, \mathbf{x}^\circ)$ leaves the simplex at some time $t^* \in \mathbb{R}$, i.e. $\xi_i(t^*, \mathbf{x}^\circ) < 0$ for some i . Then, by continuity of $\boldsymbol{\xi}$, $\xi_i(t^\circ, \mathbf{x}^\circ) = 0$ for some $t^\circ \in \mathbb{R}$. But there is a solution $\boldsymbol{\xi}^1(\cdot, \mathbf{x}^1)$ through the point $\mathbf{x}^1 = \boldsymbol{\xi}(t^\circ, \mathbf{x}^\circ)$, with $\xi_i^1(t, \mathbf{x}^1) = 0$ for all $t \in \mathbb{R}$ (since $x_i = 0 \Rightarrow \dot{x}_i = 0$). Then $\boldsymbol{\xi}(\cdot, \mathbf{x}^\circ)$ and $\boldsymbol{\xi}^1(\cdot, \mathbf{x}^1)$ are distinct solutions (they disagree at $t = t^*$) which both pass through the state \mathbf{x}^1 . This contradicts the uniqueness of the solution mapping, as given by Proposition 3.1, and thus we have shown that the simplex S_n is invariant under regular game dynamics.

The invariance of $bd(S_n) = \{\mathbf{x} \in S_n \mid x_i = 0 \text{ for at least one } i \in \{1, \dots, n\}\}$ follows immediately: $\xi_i(0, \mathbf{x}) = 0 \Rightarrow \dot{\xi}_i(0, \mathbf{x}) = 0$, and we have already shown that $\boldsymbol{\xi}(\cdot, \mathbf{x})$ does not leave the simplex for any $\mathbf{x} \in S_n$.

Now suppose that $int(S_n)$, the complement of the $bd(S_n)$, is not invariant. Then there exists some $\mathbf{x} \in int(S_n)$ and $t \in \mathbb{R}$ such that $\boldsymbol{\xi}(t, \mathbf{x}) \in$

¹⁷Stated in Appendix D.1; see Hirsch and Smale (1974) for proof.

$S_n \setminus \text{int}(S_n) = \text{bd}(S_n)$. But $\boldsymbol{\xi}(-t, \boldsymbol{\xi}(t, \mathbf{x})) = \boldsymbol{\xi}(-t + t, \mathbf{x}) = \boldsymbol{\xi}(0, \mathbf{x})$ from (3.2), and $\boldsymbol{\xi}(0, \mathbf{x}) = \mathbf{x} \in \text{int}(S_n)$, which contradicts the invariance of the simplex boundary.

Because $\text{int}(S_n)$ is invariant, $\xi_i(0, \mathbf{x}) \neq 0 \Rightarrow \xi_i(t, \mathbf{x}) \neq 0$ for all $t \in \mathbb{R}$. Thus, not only is the simplex boundary invariant as a whole, but also each individual boundary face $\text{bd}(S_n)_i = \{\mathbf{x} \in S_n \mid x_j = 0 \text{ only for } j = i\}$ is invariant. Finally, each vertex is trivially invariant, since from the game dynamics equations, $x_i = 0$ implies that $\xi_i(t, \mathbf{x}) = 0$ for each i and any time t . □

As an important note, the invariant subsets of the simplex are not invariant over an infinite time horizon. For example, while an interior solution cannot reach the boundary of the simplex in finite time, it can possibly converge to a boundary state as $t \rightarrow +\infty$. Furthermore, these invariances only hold under the evolutionary dynamics of selection without mutation; if mutation is introduced in the form of perturbations in state space, then a population can evolve from, for instance, a boundary state to an interior state.

We have now shown that every regular game dynamics constitutes a *dynamical system* \mathcal{D} , with time range \mathbb{R} , state space S_n , and continuous time-invariant solution mapping $\xi : \mathbb{R} \times S_n \rightarrow S_n$, as in (3.2). We can thus denote a regular game dynamics by the triple $\mathcal{D} = (\mathbb{R}, S_n, \xi)$, in analogy to our notation for normal-form games. Many of the ensuing definitions and results can be formulated for irregular game dynamics, but for our purposes we need only consider regular dynamics henceforth.

3.2 Equilibrium States

3.2.1 Dynamical Systems Formalism

In attempt to understand the long-run behavioral outcomes of evolutionary models, we review the fundamental asymptotic features of dynamical systems. Note that, while the following statements are phrased in terms of our dynamical system $\mathcal{D} = (\mathbb{R}, S_n, \xi)$, these results in fact hold for any compact state space $C \subset \mathbb{R}^n$. Accordingly, we will invoke no property of the simplex but that it is compact.

Definition. A state $\mathbf{x} \in S_n$ is a *stationary state* (also “steady state” or “fixed point”) of \mathcal{D} if $\boldsymbol{\xi}(t, \mathbf{x}) = \mathbf{x}$ for all $t \in \mathbb{R}$.

Definition. A state $\mathbf{x} \in S_n$ is *stable* (also “Lyapunov stable”) in \mathcal{D} if every neighborhood U of \mathbf{x} contains a neighborhood U' of \mathbf{x} such that

$$\boldsymbol{\xi}(t, \mathbf{x}') \in U \quad \forall \mathbf{x}' \in U' \cap S_n, \quad \forall t \geq 0.$$

Definition. A state $\mathbf{x} \in S_n$ is *asymptotically stable* in \mathcal{D} if it is stable and there exists a neighborhood U' of \mathbf{x} such that

$$\lim_{t \rightarrow \infty} \boldsymbol{\xi}(t, \mathbf{x}') = \mathbf{x} \quad \forall \mathbf{x}' \in U' \cap S_n.$$

By definition, an asymptotically stable state is stable; we now show that a stable state is stationary.

Proposition 3.3. If a state is stable, then it is stationary.

Proof. Suppose $\mathbf{x} \in S_n$ is not stationary. Then at some time t , a solution path leads to a state $\mathbf{y} \neq \mathbf{x}$: $\boldsymbol{\xi}(t, \mathbf{x}) = \mathbf{y}$. Letting $d(\mathbf{x}, \mathbf{y}) = \delta$ be the Euclidean distance between the states, we observe that the solution $\boldsymbol{\xi}(t, \mathbf{x})$ leaves the neighborhood $B_{\delta/2}(\mathbf{x})$ in finite time, so \mathbf{x} is not stable. □

Thus, all stationary states fall into one of three categories, based on a state’s response to perturbations in state space: (i) small perturbations drive the system away from the state (unstable); (ii) small perturbations leave the system forever within in a small neighborhood of the state (stable); (iii) small perturbations are soon counteracted as the system always returns to the state (asymptotically stable).

Stationarity and stability are key to our assessment of the behavioral outcomes of evolution. An unstable stationary state is only robust in the unperturbed dynamics, which represent selection without mutation, and thus such a state cannot be maintained by a population in a long-term evolutionary context. A stable state that is not asymptotically stable is not evolutionarily favored, for selection does not preferentially drive nearby populations to the state. If such a state is arrived at, there can be “neutral drift,” or uncorrected mutations, which can lead to nearby (possibly unstable) states. Thus, true evolutionary stability in a dynamical system requires robustness to mutation,

which is only achieved by asymptotic stability. In an asymptotically stable population state, perturbative mutations are corrected by selection; this is the conceptual analogy of the ESS from 2.2. We will explore in §3.3 the extent to which dynamically stable states are, like the static ESS, refinements of the fundamental notion of Nash equilibrium.

But first, in accordance with our interest in long-term behavior, we define one last form of asymptotic equilibrium, which formalizes the notion of long-run dynamical outcome.

Definition. The Ω -limit of a state $\mathbf{x} \in S_n$ is the set of all accumulation points of the solution through \mathbf{x} :

$$\Omega(\mathbf{x}) = \{\mathbf{y} \in S_n \mid \lim_{t_k \rightarrow +\infty} \boldsymbol{\xi}(t_k, \mathbf{x}) = \mathbf{y}, \text{ for some sequence } \{t_k\} \subset \mathbb{R}\}.$$

Since our system \mathcal{D} leaves the simplex invariant, and every infinite sequence in a compact set has an accumulation point, it follows that every population state $\mathbf{x} \in S_n$ has a nonempty Ω -limit. With the following proposition, we prove that when the Ω -limit consists of a sole accumulation point, it is in fact stationary.

Proposition 3.4. Any state which constitutes the Ω -limit of another state is stationary.

Proof. Let $\{\mathbf{y}\} = \Omega(\mathbf{x})$, for $\mathbf{x}, \mathbf{y} \in S_n$. Then $\lim_{t \rightarrow +\infty} \boldsymbol{\xi}(t, \mathbf{x}) = \mathbf{y}$, so that for any neighborhood U of \mathbf{y} , there exists some $t_U \geq 0$ with $\boldsymbol{\xi}(t, \mathbf{x}) \in U$, $\forall t \geq t_U$. Suppose that \mathbf{y} is not stationary. Then $\boldsymbol{\xi}(t^*, \mathbf{y}) \neq \mathbf{y}$, for some time $t^* \geq 0$. Thus, since $\boldsymbol{\xi}(t^*, \mathbf{y})$ is continuous in \mathbf{y} (by Proposition 3.1), there exists a neighborhood V of \mathbf{y} such that every trajectory beginning in V is elsewhere at time $t = t^*$: $\boldsymbol{\xi}(t^*, \mathbf{y}') \notin V$, $\forall \mathbf{y}' \in V$.

Now, letting $U = V$, we have that $\boldsymbol{\xi}(t, \mathbf{x}) \in V$, $\forall t \geq t_V$, with $\mathbf{y}' \equiv \boldsymbol{\xi}(t_V, \mathbf{x}) \in V$. By (3.2) we can specify any trajectory beginning at \mathbf{y}' as $\boldsymbol{\xi}(t, \mathbf{y}') = \boldsymbol{\xi}(t, \boldsymbol{\xi}(t_V, \mathbf{x})) = \boldsymbol{\xi}(t + t_V, \mathbf{x}) \in V$, $\forall t \geq 0$. Setting $t = t^*$, we then have that $\boldsymbol{\xi}(t^*, \mathbf{y}') \in V$. But $\mathbf{y}' \in V$, which is a contradiction. □

Proposition 3.4 establishes the direct link between long-term dynamical behavior and stationarity. Thus, we have verified that finding the stationary states of a game dynamics sufficiently determines the long-run evolutionary outcomes.

3.2.2 A Brief Note on Computation

While our classification of stationary states thus far is useful in developing an intuition for dynamical equilibria, as well as in proving the ensuing theoretical results, it does not immediately allow us to determine the stable states of a given model. Stationary states are often analytical solvable, from the system of equations: $\dot{x}_i = 0$, $i = 1, \dots, n$. Otherwise, numeric integration can be used to simulate solution trajectories from many random initial states, yielding stationary states in the form of Ω -limits.

Once stationary states are obtained, we typically determine stability in a game dynamics by analyzing the eigenvalues of the system's Jacobian matrix $\mathcal{J}(\mathbf{x})$ evaluated at each stationary state.¹⁸ Given a stationary state \mathbf{x} , an eigenvalue of $\mathcal{J}(\mathbf{x})$ with positive real part indicates a positive velocity away from \mathbf{x} in the direction of the corresponding eigenvector. In evolutionary terms, given a mutation toward this critical eigenstate, selection drives the population away from \mathbf{x} . Thus, for a stable state \mathbf{x} , no eigenvalue of $\mathcal{J}(\mathbf{x})$ can have positive real part; for an asymptotically stable state \mathbf{x} , all eigenvalues of $\mathcal{J}(\mathbf{x})$ must have negative real part. The linearization theorem of Hartman (1960)¹⁹ formally justifies this Jacobian-based stability analysis. Dynamic stability analysis cannot always be performed via linearization, but this method is commonly used in evolutionary dynamics, and it will suffice for our purposes in §5.

3.3 Dynamic and Static Equilibria

We now explore the connection between long-run game-dynamical behavior and static game theory. Imposing some basic requirements on the relationship between the dynamical system \mathcal{D} and the underlying game \mathcal{G} , we can prove strong results connecting dynamic equilibria and Nash equilibria.

¹⁸Denoting the game dynamics (3.1) as $f_i(\mathbf{x}) = x_i g_i(\mathbf{x})$, the Jacobian of the system at a state \mathbf{x} is: $\mathcal{J}(\mathbf{x}) = \left(\frac{\partial f_i}{\partial x_j}(\mathbf{x}) \right)$.

¹⁹Which we state, but not prove, in Appendix D.2.

3.3.1 Weak Payoff Positivity

We begin by defining for any point $\mathbf{x} \in S_n$ a corresponding set

$$\mathcal{B}(\mathbf{x}) = \{\mathbf{e}_i \in S_n \mid \mathbf{e}_i \cdot A\mathbf{x} > \mathbf{x} \cdot A\mathbf{x}\}. \quad (3.4)$$

Interestingly, the meaning of this set can be interpreted in several ways, since a point in the simplex can be taken to represent either a mixed strategy in the underlying game \mathcal{G} , or a population state in the dynamical system \mathcal{D} . Here we choose the interpretation of Weibull (1995), that $\mathcal{B}(\mathbf{x})$ is the (possibly empty) set of pure strategies which earn above average payoff if \mathcal{G} is played by random pairs of individuals throughout a population in state \mathbf{x} .

We now introduce a game dynamics condition stronger than regularity.

Definition. A regular game dynamics (3.1) is *weakly payoff positive* if, for all $\mathbf{x} \in S_n$,

$$\mathcal{B}(\mathbf{x}) \neq \emptyset \Rightarrow g_i(\mathbf{x}) > 0 \text{ for some } \mathbf{e}_i \in \mathcal{B}(\mathbf{x}). \quad (3.5)$$

Weak payoff positivity is not so much to ask – it only requires that if any pure strategies earn above the population average payoff, at least one such strategy grows in the population. Indeed, weakly payoff positive dynamics is the largest class of game dynamics for which we can prove the Folk Theorem of Evolutionary Game Theory, a powerful result that classifies fundamental game-dynamical equilibria as Nash equilibria of the underlying game (Hofbauer and Sigmund, 1998; Fudenberg and Levine, 1998; Weibull, 1995).

Theorem 3.1 (Folk Theorem of Evolutionary Game Theory). *In weakly payoff positive game dynamics, given by dynamical system $\mathcal{D} = (\mathbb{R}, S_n, \xi)$ and underlying game $\mathcal{G} = (2, \mathbb{S}, A)$, the following implications hold:*

- (a) *If the state $\mathbf{x} \in \text{int}(S_n)$ is stationary in \mathcal{D} , then the corresponding strategy \mathbf{x} is a symmetric Nash equilibrium of \mathcal{G} .*
- (b) *If the state $\mathbf{x} \in S_n$ is stable in \mathcal{D} , then the corresponding strategy \mathbf{x} is a symmetric Nash equilibrium of \mathcal{G} .*
- (c) *If the state $\{\mathbf{x}\} = \Omega(\mathbf{y})$ for some $\mathbf{y} \in \text{int}(S_n)$, then the corresponding strategy \mathbf{x} is a symmetric Nash equilibrium of \mathcal{G} .*

Proof. First, observe that $\mathcal{B}(\mathbf{x}) = \emptyset \Rightarrow \mathbf{e}_i \cdot A\mathbf{x} \leq \mathbf{x} \cdot A\mathbf{x} \forall i \Rightarrow \mathbf{x}$ is Nash (cf. (2.5)).

- (a) Let $\mathbf{x} \in \text{int}(S_n)$ be stationary in \mathcal{D} . Then $\dot{x}_i = x_i g_i(\mathbf{x}) = 0$ for all i , because \mathbf{x} is stationary, and $x_i \neq 0$ for all i because $\mathbf{x} \in \text{int}(S_n)$. Thus $g_i(\mathbf{x}) = 0$ for all i , which implies $\mathcal{B}(\mathbf{x}) = \emptyset$, so \mathbf{x} is Nash in \mathcal{G} .
- (b) Let $\mathbf{x} \in S_n$ be stable (hence stationary) in \mathcal{D} . By similar logic as (a), $g_i(\mathbf{x}) = 0$ whenever $\mathbf{e}_i \in \mathcal{S}(\mathbf{x})$. Now suppose that \mathbf{x} is not Nash in \mathcal{G} . Then $\mathcal{B}(\mathbf{x}) \neq \emptyset$, so by weak payoff positivity, there is some $\mathbf{e}_j \in \mathcal{B}(\mathbf{x})$ for which $g_j(\mathbf{x}) > 0$, and also $\mathbf{e}_j \notin \mathcal{S}(\mathbf{x})$, i.e. $x_j = 0$.

Since \mathbf{g} is continuous in a regular dynamics, there exists some $\delta > 0$ and neighborhood $U \subset S_n$ of \mathbf{x} such that $g_j(\mathbf{x}') > \delta$ for any $\mathbf{x}' \in U$. Then, recalling the game dynamics (3.1), we have

$$\dot{\xi}_j(t, \mathbf{x}') = \xi_j(t, \mathbf{x}') g_j(\boldsymbol{\xi}(t, \mathbf{x}')).$$

Then, if $\boldsymbol{\xi}(t, \mathbf{x}') \in U$ whenever $t \in [0, t^*)$,

$$\log \left(\frac{\xi_j(t, \mathbf{x}')}{\xi_j(0, \mathbf{x}')} \right) = \int_0^t g_j(\boldsymbol{\xi}(\tau, \mathbf{x}')) d\tau > \int_0^t \delta d\tau$$

which implies:

$$\xi_j(t, \mathbf{x}') > x'_j e^{\delta t}.$$

Thus, while $x_j = 0$, $\xi_j(t, \mathbf{x}')$ grows exponentially as long as $\boldsymbol{\xi}(t, \mathbf{x}') \in U$. In particular, for any neighborhood $U' \subset U$ of \mathbf{x} , the solution $\boldsymbol{\xi}(t, \mathbf{x}')$ leaves U' as $t \rightarrow t^*$, contradicting the fact that \mathbf{x} is stable. Thus, \mathbf{x} is Nash in \mathcal{G} .

- (c) Let $\{\mathbf{x}\}$ be the Ω -limit of \mathbf{y} for some $\mathbf{y} \in \text{int}(S_n)$, namely: $\lim_{t \rightarrow +\infty} \boldsymbol{\xi}(t, \mathbf{y}) = \mathbf{x}$. Then by Proposition 3.4, \mathbf{x} is stationary. Thus we begin in the same manner as (b): supposing \mathbf{x} is not Nash, we have that $\mathcal{B}(\mathbf{x}) \neq \emptyset$, so that there is some $\mathbf{e}_j \in \mathcal{B}(\mathbf{x})$ for which $g_j(\mathbf{x}) > 0$ and $x_j = 0$.

Since \mathbf{g} is continuous, there exists some neighborhood $U \subset S_n$ of \mathbf{x} such that $g_j(\mathbf{x}') > 0$ for any $\mathbf{x}' \in U$. Also, since $\boldsymbol{\xi}(t, \mathbf{y}) \rightarrow \mathbf{x}$ as

$t \rightarrow +\infty$, there is a time $t^* > 0$ such that $\boldsymbol{\xi}(t, \mathbf{y}) \in U$ for all $t \geq t^*$. But since $x_j = 0$, we must have $\xi_j(t, \mathbf{y}) < 0$ for some $t \geq t^*$, which contradicts the fact that g_j is positive on U . Thus \mathbf{x} is Nash in \mathcal{G} .

□

The Folk Theorem asserts that, in a weakly payoff positive game dynamics, certain forms of dynamical equilibrium are at least as strong as the static Nash equilibrium. This alone does not qualify these dynamical equilibria as meaningful refinements of the Nash equilibrium. However, as discussed previously, the asymptotically stable state is our best candidate for Nash refinement, since it is robust to invasion by mutant strategies in the dynamical setting – much like the ESS in static game theory. Thus, the Folk Theorem, together with our evolutionary intuition, establishes asymptotic stability as a refinement of the Nash equilibrium, under certain game dynamics.

3.3.2 Monotonicity

We can develop even stronger connections between dynamic and static equilibria for a subclass of weakly payoff positive dynamics known as monotonic game dynamics.

Definition. A regular game dynamics (3.1) is *convex monotonic* if, for any states $\mathbf{x}, \mathbf{y} \in S_n$, the ranking of pure- versus mixed-strategy growth rates respects their difference in payoff:

$$\mathbf{e}_i \cdot A\mathbf{x} > \mathbf{y} \cdot A\mathbf{x} \iff g_i(\mathbf{x}) > \mathbf{y} \cdot \mathbf{g}(\mathbf{x}) \quad \forall i \in \{1, \dots, n\}. \quad (3.6)$$

We say that the game dynamics is simply *monotonic* if the above implication is only guaranteed for any pure strategy $\mathbf{y} = \mathbf{e}_j$.

Thus convex monotonicity implies monotonicity by definition, and the following proposition tells us that monotonicity in turn implies weak payoff positivity.

Proposition 3.5. A monotonic game dynamics is weakly payoff positive.

Proof. Assuming monotonicity, we aim to show that condition (3.5) holds for an arbitrary population state $\mathbf{x} \in S_n$. Since \mathbf{x} is a vector of nonnegative weights, and $\mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = 0$ in any regular dynamics, it is always true that either $g_i(\mathbf{x}) = 0$ for all i , or $g_i(\mathbf{x}) > 0$ for some i . In the first case, condition

(3.5) holds trivially. In the latter case, we can choose the strategy \mathbf{e}_i which maximizes payoff $\mathbf{e}_i \cdot \mathbf{A}\mathbf{x}$, and monotonicity then guarantees that $g_i(\mathbf{x})$ will be maximal, hence positive. Thus, $\mathbf{e}_i \in \mathcal{B}(\mathbf{x})$, so (3.5) holds. \square

We now prove that any monotonic dynamics eliminates from the population any pure strategy that is iteratively strictly dominated (cf. (2.6)) in the underlying game.²⁰ This is a useful property of monotonic dynamics that is not guaranteed in the broader class of weakly payoff positive dynamics. It turns out that monotonicity suffices to eliminate only those pure strategies which are iteratively strictly dominated by some other pure strategy; we need convex monotonicity to guarantee the elimination of pure strategies dominated by any other strategy (Hofbauer and Weibull, 1996).

Theorem 3.2. *Let \mathbf{e}_i be a pure strategy that is iteratively strictly dominated in a game \mathcal{G} . Then, under any convex monotonic game dynamics \mathcal{D} , the proportion of the population playing \mathbf{e}_i converges to zero from any interior initial state:*

$$\lim_{t \rightarrow +\infty} \xi_i(t, \mathbf{x}) = 0 \quad \forall \mathbf{x} \in \text{int}(S_n).$$

Proof. Recall from §2.1 the discussion of iterative strategic dominance, with $\mathbb{S}^1, \mathbb{S}^2, \dots, \mathbb{S}^T$ denoting the iteratively trimmed strategy sets of $\mathcal{G} = (2, \mathbb{S}, \pi)$. Let $\mathbb{S}_\circ \subseteq \mathbb{S}$ be the set of pure strategies which are iteratively strictly dominated by pure strategies, but do not converge to zero under the dynamics \mathcal{D} . Also, for notational convenience in this proof, we will let $\mathbb{S} = \{1, \dots, n\}$, and refer to a pure strategy \mathbf{e}_i simply as “strategy i .”

Suppose that $\mathbb{S}_\circ \neq \emptyset$. For any $i \in \mathbb{S}_\circ$, let $\tau(i)$ denote the number of iterative deletions necessary to eliminate strategy i , namely $i \in \mathbb{S}^{\tau(i)} \setminus \mathbb{S}^{\tau(i)+1}$. Let $k \in \mathbb{S}_\circ$ minimize $\tau(i)$ on \mathbb{S}_\circ , and $\tau = \tau(k)$ (i.e., of the strategies which do not converge to zero in \mathcal{D} , strategy k is the first to be eliminated, in the τ th trimmed subgame). Then there is a strategy $\mathbf{y} \in S_n$ such that $\mathbf{e}_k \cdot \mathbf{A}\mathbf{e}_j < \mathbf{y} \cdot \mathbf{A}\mathbf{e}_j$ for all $j \in \mathbb{S}^\tau$. Note that because $\tau = \tau(k)$ is minimal, $\lim_{t \rightarrow +\infty} \xi_j(t, \mathbf{x}) = 0$ for any $j \notin \mathbb{S}^\tau$, since any such strategy j is strictly

²⁰See Samuelson and Zhang (1992) for a proof of this result for all (not necessarily symmetric) two-player normal-form games.

dominated and not in \mathbb{S}_o . Thus, as $t \rightarrow +\infty$,

$$\begin{aligned} \mathbf{e}_k \cdot A\boldsymbol{\xi}(t, \mathbf{x}) - \mathbf{y} \cdot A\boldsymbol{\xi}(t, \mathbf{x}) &= \sum_{j \in \mathbb{S}^r} (\mathbf{e}_k \cdot A\mathbf{e}_j - \mathbf{y} \cdot A\mathbf{e}_j) \xi_j(t, \mathbf{x}) \\ &\quad + \sum_{j \notin \mathbb{S}^r} (\mathbf{e}_k \cdot A\mathbf{e}_j - \mathbf{y} \cdot A\mathbf{e}_j) \xi_j(t, \mathbf{x}) \\ &\rightarrow \sum_{j \in \mathbb{S}^r} (\mathbf{e}_k \cdot A\mathbf{e}_j - \mathbf{y} \cdot A\mathbf{e}_j) \xi_j(t, \mathbf{x}) \leq 0. \end{aligned}$$

Strict inequality is guaranteed if $\mathbf{x} \in \text{int}(S_n)$, so that $\xi_j(t, \mathbf{x}) \neq 0 \ \forall j$.

Then, by convex monotonicity, $g_k(\boldsymbol{\xi}(t, \mathbf{x})) - \mathbf{y} \cdot \mathbf{g}(\boldsymbol{\xi}(t, \mathbf{x})) < 0$ for $\mathbf{x} \in \text{int}(S_n)$. By regularity, \mathbf{g} is continuous, so there exists an $\varepsilon > 0$ and a $t^* > 0$ such that:

$$g_k(\boldsymbol{\xi}(t, \mathbf{x})) - \mathbf{y} \cdot \mathbf{g}(\boldsymbol{\xi}(t, \mathbf{x})) = \frac{\dot{\xi}_k(t, \mathbf{x})}{\xi_k(t, \mathbf{x})} - \sum_{i=1}^n y_i \frac{\dot{\xi}_i(t, \mathbf{x})}{\xi_i(t, \mathbf{x})} \leq -\varepsilon \quad \forall t > t^*.$$

Integrating from t^* to t , we get:

$$\frac{\xi_k(t, \mathbf{x})}{\prod_{i=1}^n \xi_i(t, \mathbf{x})^{y_i}} \leq \frac{\xi_k(t^*, \mathbf{x})}{\prod_{i=1}^n \xi_i(t^*, \mathbf{x})^{y_i}} e^{-\varepsilon(t-t^*)} \quad \forall t > t^*.$$

Since $\prod_i \xi_i(t, \mathbf{x})^{y_i} \leq 1$, this implies that $\lim_{t \rightarrow +\infty} \xi_k(t, \mathbf{x}) = 0$, which contradicts the fact that $k \in \mathbb{S}_o$. Hence $\mathbb{S}_o = \emptyset$, as desired. \square

The corresponding result for a general monotonic dynamics is an immediate corollary to Theorem 3.2: taking \mathbf{y} to be a pure strategy in the above proof, we see that a monotonic game dynamics guarantees that any pure strategy strictly dominated by another pure strategy converges to zero in the dynamical population.

Theorem 3.2 (and its corollary result) establishes a meaningful property of monotonic dynamics, in that any asymptotically *reachable* state must be “rational,” in the sense that no portion of the population plays a strategy which is strictly dominated (hence “irrational”) in the underlying game. This property can be seen as both a strongpoint and a limitation of monotonic dynamics: on the one hand, the dynamics can perhaps help to refine our

notion of stable rational behavior in a game; at the same time, the dynamics is so intimately tied to the underlying game that economically irrational behaviors are not allowed to evolve.

We can further bolster our understanding of monotonic dynamics with a precise characterization of its stationary states in terms of underlying game payoffs (Weibull, 1995):

Proposition 3.6. Under any monotonic game dynamics, the set of stationary states is given by

$$S_n^\circ = \{\mathbf{x} \in S_n \mid \mathbf{e}_i \cdot A\mathbf{x} = \mathbf{x} \cdot A\mathbf{x}, \forall \mathbf{e}_i \in \mathcal{S}(\mathbf{x})\}. \quad (3.7)$$

Proof. By definition of monotonicity, $\mathbf{x} \in S_n^\circ$ if and only if $g_i(\mathbf{x}) = c$ for some $c \in \mathbb{R}$ and any $\mathbf{e}_i \in \mathcal{S}(\mathbf{x})$. This is equivalent to $\mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = c$, and we require that $c = 0$ in any regular dynamics. Thus \mathbf{x} is a stationary state. Since the preceding implications are all bidirectional, we have proven that S_n° is precisely the set of stationary states in a monotonic game dynamics. \square

Proposition 3.6 allows for explicit computation of stationary states in a finite game: there are finitely many support sets,²¹ and for each support set there are an equal number of linear equations ($(A\mathbf{x})_i = \mathbf{x} \cdot A\mathbf{x}, \forall \mathbf{e}_i \in \mathcal{S}(\mathbf{x})$) and unknowns ($x_i, \forall \mathbf{e}_i \in \mathcal{S}(\mathbf{x})$).

Also, note that S_n° is precisely the set of Nash equilibria of all subgames of the underlying game \mathcal{G} (where a subgame is simply a restriction of \mathcal{G} to some subset of its original pure strategy set). To see this: any $\mathbf{x} \in S_n^\circ$ is Nash in the subgame restricted to the strategy space $\mathcal{S}(\mathbf{x})$, and conversely any Nash strategy \mathbf{x} lies in S_n° because \mathbf{x} is indifferent among the strategies in $\mathcal{S}(\mathbf{x})$.

While Proposition 3.6 fully determines stationarity in a monotonic game dynamics, it tells us nothing of stability. As discussed in §3.2, an unstable stationary state is only robust in the unperturbed dynamics, i.e. under selection without mutation, and thus cannot be maintained by a population in a long-term evolutionary context.

²¹Namely, an n -strategy game has $2^n = |\mathcal{P}(\{1, \dots, n\})|$ possible support sets.

There are other classifications of dynamics which lead to more nuanced connections between static and dynamic equilibria, but the results proven thus far will suffice for our purposes. Further investigation is left to the compelled reader.

4 Specific Dynamics

Now having established some key results connecting static games and evolutionary dynamics in general, we develop two particular dynamical models. The discussion of their specific dynamics will serve to illustrate the theoretical results of §3, as well as to lay the foundations for our study of cooperation and punishment in §5.

4.1 Fitness and Selection Strength

To move from the general game dynamics (3.1) to a specific evolutionary model, we must specify each strategy’s growth rate as a function of population state and underlying game payoffs. Typically a fitness function is first defined from payoffs, and then a game-dynamical growth rate is derived from fitness. While this intermediary fitness mapping may seem a tedious formality, it is in fact both conceptually and mathematically advantageous to embed a manipulable fitness function within the rigid structure of a game dynamics.

Traulsen et al. (2008) discuss the choice of functional form for the payoff-to-fitness mapping, which determines individuals’ “evolutionary fitness” in a population state, based on their expected payoffs from playing the underlying game against random opponents. While in most cases changing the fitness mapping alters only the rate of selection, and not the critical evolutionary outcomes of the dynamics, there are cases when the choice of fitness is crucial and relevant.

Since the individuals in our population are distinguishable only by the strategies they employ, it suffices to define a fitness function for each pure strategy. The traditional fitness of strategy \mathbf{e}_i is given by the linear map

$$f_i(\mathbf{x}) = 1 - w + w(\mathbf{e}_i \cdot A\mathbf{x}) \quad (4.1)$$

where $w \in [0, 1]$ gives the strength of selection, i.e. the extent to which evolutionary fitness depends on game payoff. Under *strong selection*, or $w = 1$, individuals’ evolutionary fitness is entirely determined by their payoffs in the specified underlying game. The opposite extreme, $w = 0$, corresponds to *neutral drift*, when evolution pays no heed to the outcome of the underlying game, thus barring any nontrivial evolutionary analysis.

An alternative formulation of fitness is the exponential mapping

$$f_i(\mathbf{x}) = e^{\omega(\mathbf{e}_i \cdot A\mathbf{x})} \quad (4.2)$$

where $\omega > 0$ is the selection strength, analogous to w in the linear map. The exponential fitness is in some sense more general than linear fitness, since there is a greater range of selection strength achievable through ω than through w .

We now consider the effects of varying the selection strength parameter in both fitness mappings. It is reasonable to suppose that, in general, the evolutionary fitness of individuals in a population is determined by many factors – including other games – and that the particular underlying game in consideration only contributes in part to this overall fitness (Traulsen et al., 2007). The notion of *weak selection* achieves this aim: when $w \ll 1$ in (4.1), or ω sufficiently small²² in (4.2), game payoffs are just slight perturbations of individual fitnesses, which are all very close to 1. In the weak selection limit $\omega = w \rightarrow 0$, the two fitness mappings are approximately equivalent:

$$f_i(\mathbf{x}) \approx 1 + w(\mathbf{e}_i \cdot A\mathbf{x}) \quad (4.3)$$

This limit can be illuminating, especially in the analysis of stochastic dynamics for finite populations (Nowak et al., 2004), or when the game dynamical system of equations is too complex for analytical solutions (as we will witness in §5.3).

4.2 Replicator Dynamics

The original paradigm for evolutionary dynamics, first introduced by Taylor and Jonker (1978), is the now-ubiquitous *replicator dynamics*. The replicator dynamics have many widespread motivations and justifications, for instance as a form of imitation or learning (Traulsen and Hauert, 2008), as utility-based economic competition (Fudenberg and Levine, 1998), or simply as biological survival-of-the-fittest (Hofbauer and Sigmund, 1998).²³ Keeping in

²²We seek ω small enough that we may take the first-order Taylor approximation of the fitness function.

²³In particular, Hofbauer and Sigmund (1998) present a differentiable, invertible map that proves the equivalence of the replicator dynamics and the famous Lotka-Volterra equations governing predator-prey population dynamics in ecology.

mind this notion of imitation, which is perhaps best suited for the evolution of cooperation, we proceed in deriving the replicator dynamics from the general game dynamics (3.1).

4.2.1 Motivation and Definition

The replicator dynamics stems from, in some sense, the “simplest” possible growth rate function $\mathbf{g}(\mathbf{x})$ in the game dynamics (3.1). We first suppose that the growth rate of individuals playing strategy \mathbf{e}_i in a population is directly proportional to the strategy’s linear fitness (4.1). This leads to a (not necessarily simplex-constrained) continuous-time dynamical system,

$$\dot{y}_i = y_i f_i(\mathbf{x})$$

where y_i is the number of individuals in the population playing strategy \mathbf{e}_i , and $\mathbf{x} \in S_n$ is the strategic population state, as usual. This system induces a game dynamics in the population shares $x_i = y_i / \sum_{j=1}^n y_j$,

$$\begin{aligned} \dot{x}_i &= \frac{\dot{y}_i \sum_{j=1}^n y_j - y_i \sum_{j=1}^n \dot{y}_j}{(\sum_{j=1}^n y_j)^2} \\ &= x_i f_i(\mathbf{x}) - x_i \sum_{j=1}^n x_j f_j(\mathbf{x}) \\ &= x_i (f_i(\mathbf{x}) - \phi(\mathbf{x})) \end{aligned} \tag{4.4}$$

where $\phi(\mathbf{x}) = \sum_j x_j f_j(\mathbf{x})$ is the average fitness of a population in state \mathbf{x} , or equivalently, the expected payoff in a game between two randomly chosen members of the population. The dynamical system (4.4) is known as the *replicator dynamics*, first introduced by Taylor and Jonker (1978). By substituting the linear fitness expression (4.1) into (4.4), we can easily verify that the selection strength parameter w has no effect on the dynamics beyond a universal “velocity” factor, which alters the speed of evolution uniformly for all strategies. Thus, without loss of generality, we can take the strong selection limit $w = 1$, to achieve a more concrete form of the replicator dynamics,

$$\dot{x}_i = x_i (\mathbf{e}_i \cdot \mathbf{A} \mathbf{x} - \mathbf{x} \cdot \mathbf{A} \mathbf{x}) \quad i = 1, \dots, n. \tag{4.5}$$

The preceding derivation illustrates our initial claim that the replicator dynamics is the “simplest” form of regular game dynamics: the mapping from

payoff to fitness to growth rate is the minimal deviation from the identity map necessary to leave the unit simplex invariant. Intuitively, this growth rate is given by the difference between the expected payoff of playing a randomly chosen opponent, and the population average expected payoff $\phi(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x}$ (as mentioned above).

We now observe that the replicator dynamics is a regular game dynamics (cf. (3.3)). Regularity follows from the fact that \mathbf{g} is polynomial (hence Lipschitz continuous) in \mathbf{x} , and (4.5) is invariant on S_n by construction:

$$\mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = \sum_{i=1}^n x_i (\mathbf{e}_i \cdot A\mathbf{x} - \mathbf{x} \cdot A\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x} - (\mathbf{x} \cdot A\mathbf{x}) \sum_{i=1}^n x_i = 0.$$

Furthermore, we can establish the stronger property of convex monotonicity, in the sense of (3.6).

Proposition 4.1. The replicator dynamics (4.5) is convex monotonic.

Proof. For any strategy i , the following holds for all $\mathbf{x}, \mathbf{y} \in S_n$:

$$\begin{aligned} \mathbf{e}_i \cdot A\mathbf{x} > \mathbf{y} \cdot A\mathbf{x} &\iff \mathbf{e}_i \cdot A\mathbf{x} - \mathbf{x} \cdot A\mathbf{x} > \mathbf{y} \cdot A\mathbf{x} - \mathbf{x} \cdot A\mathbf{x} \\ &\iff g_i(\mathbf{x}) > \sum_{i=1}^n y_i (\mathbf{e}_i \cdot A\mathbf{x}) - \left(\sum_{i=1}^n y_i\right) (\mathbf{x} \cdot A\mathbf{x}) \\ &\iff g_i(\mathbf{x}) > \mathbf{y} \cdot \mathbf{g}(\mathbf{x}). \end{aligned}$$

This satisfies the definition (3.6) of convex monotonicity. □

Thus, recalling the results of §3.3, Proposition 4.1 allows us to use the Folk Theorem 3.1, as well as Theorem 3.2, to describe the asymptotic behavior of the replicator dynamics. Moreover, the next section will establish an even stronger result, which intimately relates the asymptotics of the replicator dynamics to the evolutionarily stable state from static game theory.

4.2.2 Asymptotic Stability and the ESS

Taylor and Jonker (1978) first established the connection between a stable state in the replicator dynamics and an ESS of the underlying game. Hofbauer et al. (1979) offer a refinement of the original result, which is of greater

interest to us. Its proof is an application of Lyapunov’s second stability theorem, which we here state but not prove.²⁴

Theorem (Lyapunov Stability Theorem). *Let $\mathcal{D} = (\mathbb{R}, C, \xi)$ be a dynamical system defined on some compact subset C of \mathbb{R}^n . A point $\mathbf{x} \in C$ is asymptotically stable if and only if there exists a neighborhood $U \subset C$ of \mathbf{x} and a continuous function $v : U \rightarrow \mathbb{R}$ such that the following conditions hold:*

$$v(\mathbf{x}) \text{ is the unique maximum of } v \text{ on } U. \quad (4.6)$$

$$v(\xi(t, \mathbf{x}')) > v(\mathbf{x}') \text{ if } \mathbf{x}' \neq \mathbf{x}, t > 0, \text{ and } \xi(t', \mathbf{x}') \in U \forall t' \in [0, t]. \quad (4.7)$$

In words, Lyapunov’s theorem guarantees the asymptotic stability of a state in a dynamical system, if we find a real-valued function on the state space that meets conditions (4.6) and (4.7). Note that condition (4.7) mandates that the Lyapunov function be strictly increasing on every trajectory in a neighborhood of the asymptotically stable state. If the function is differentiable, this translates to having a positive-definite derivative along every such trajectory.

Theorem 4.1. *Any ESS of an underlying game is an asymptotically stable state in the replicator dynamics.*

Proof. We will prove the theorem by invoking Lyapunov’s theorem for the replicator dynamics, using the well-known “entropy function”:

$$\mathcal{E}_{\mathbf{x}}(\mathbf{y}) = \prod_{i=1}^n y_i^{x_i}.$$

It is well-known that $\mathcal{E}_{\mathbf{x}}(\mathbf{y})$ has a unique maximum on the simplex at $\mathbf{y} = \mathbf{x}$.²⁵ To compute this explicitly, we maximize the logarithm of the monotonic function $\mathcal{E}_{\mathbf{x}}$ using the technique of Lagrange multipliers:

$$\frac{\partial \log(\mathcal{E}_{\mathbf{x}}(\mathbf{y}))}{\partial y_i} = x_i/y_i = \lambda.$$

²⁴See Weibull (1995) for proof and references. Note that Lyapunov’s theorem is formulated more generally for asymptotically stable closed sets; we here state a special case of the theorem, when the set is a singleton. Also, the inequality is typically reversed in condition (4.7), which is equivalent to multiplying our Lyapunov function by a factor of -1 .

²⁵For instance, $\mathcal{E}_{\mathbf{x}}(\mathbf{y})$ is the likelihood function for random sampling from a multinomial distribution, and the maximum likelihood estimate for a sample \mathbf{y} is proportional to the sampling probabilities \mathbf{x} .

Constraining \mathbf{y} to the simplex then gives $\lambda = 1$, i.e. $\mathbf{y} = \mathbf{x}$.

Recall from the proof of Theorem 2.1 that the set of all boundary faces of S_n which do not contain \mathbf{x} is a closed set. Thus, its complement in S_n is an open neighborhood of \mathbf{x} :

$$U_{\mathbf{x}} = \{\mathbf{y} \in S_n \mid y_i > 0 \text{ whenever } \mathbf{e}_i \in \mathcal{S}(\mathbf{x})\}.$$

It is clear that $\mathcal{E}_{\mathbf{x}}(\mathbf{y}) > 0$ for $\mathbf{y} \in U_{\mathbf{x}}$. Furthermore, under the replicator dynamics (4.5), we have:

$$\begin{aligned} \frac{\dot{\mathcal{E}}_{\mathbf{x}}}{\mathcal{E}_{\mathbf{x}}} &= \frac{d}{dt} \log \mathcal{E}_{\mathbf{x}} = \frac{d}{dt} \sum_{i=1}^n x_i \log y_i = \sum_{\mathbf{e}_i \in \mathcal{S}(\mathbf{x})} x_i \frac{\dot{y}_i}{y_i} \\ &= \sum x_i ((A\mathbf{y})_i - \mathbf{y} \cdot A\mathbf{y}) \\ &= \mathbf{x} \cdot A\mathbf{y} - \mathbf{y} \cdot A\mathbf{y}. \end{aligned}$$

Proposition 2.6 guarantees that any ESS is locally superior, i.e. $\mathbf{x} \cdot A\mathbf{y} - \mathbf{y} \cdot A\mathbf{y} > 0$ in some neighborhood $V_{\mathbf{x}}$ of \mathbf{x} . Thus $\dot{\mathcal{E}}_{\mathbf{x}}(\mathbf{y}) > 0$ on the neighborhood $V_{\mathbf{x}} \cap U_{\mathbf{x}}$ of \mathbf{x} , so $\mathcal{E}_{\mathbf{x}}$ is a strict local Lyapunov function, and \mathbf{x} is asymptotically stable in the replicator dynamics. □

Van Damme (1987) shows by counterexample that the converse of Theorem 4.1 does not hold in general. Thus, while the Folk Theorem 3.1 tells us that (in any weakly payoff positive dynamics) an asymptotically stable state is a refinement of the symmetric Nash equilibrium, we find that in the replicator dynamics, it is a weaker refinement than the ESS. The replicator dynamics can then be used to find game equilibria that are compromises between Nash and ESS conditions, but it does not offer us a refined notion of equilibrium that is stronger than (or unrelated to) the ESS. In §4.4.1 our analysis of two-strategy replicator systems demonstrates this inescapable link between stability in the replicator dynamics and ESS's of the underlying game.

4.3 Viability Updating

We now introduce a game dynamics which operates by a selection mechanism quite different from that of the replicator dynamics. We have seen that the replicator dynamics is convex monotonic, and thus firmly supports

the notions of strategic dominance and rationality in the underlying game. Viability updating, on the other hand, is not even weakly payoff positive, and thus it may allow conventionally irrational strategies to persist through evolution. In particular, our viability model operates through random interactions of small groups of individuals in a population, while the replicator dynamics is premised on random pairwise interactions. Thus viability updating captures the gregarious aspect of evolution, which favors strategies that are self-reinforced by “ganging up” on other strategies.

4.3.1 Motivation and Definition

Viability updating, like the replicator dynamics, begins with the reasonable fundamental assumption that individuals with lower fitnesses are more likely to change strategy (or “die” in a biological context). Furthermore, following the derivation of Ohtsuki (2008), we develop viability updating in a z -mixed population, in which individual fitness is determined from playing an underlying game \mathcal{G} in random groups of z individuals.

Thus, at any instant in time, we successively randomly choose: (1) an individual to possibly change strategy (the imitator), (2) $z - 1$ other individuals to form an interaction group (the z -group), and (3) a member of the z -group whose strategy is possibly adopted by the imitator (the imitatee). Observe that the probabilistic constituents of this random z -group is governed by random sampling from the multinomial distribution,

$$M(\mathbf{z}; z, \mathbf{x}) \equiv \frac{z!}{z_1! \cdots z_n!} x_1^{z_1} \cdots x_n^{z_n}. \quad (4.8)$$

In words, $M(\mathbf{z}; z, \mathbf{x})$ gives the probability that exactly z_i individuals of each strategy \mathbf{e}_i appear in a random group of $z = \sum_{i=1}^n z_i$ individuals.

The payoff an \mathbf{e}_i -individual receives in a z -group of composition \mathbf{z} is just $\mathbf{e}_i \cdot A\mathbf{z}$. Thus, recalling our original updating assumption, the imitator adopts the imitatee’s strategy with probability inversely proportional to the fitness of the imitator, where fitness $f_i(\mathbf{z})$ derives from game-play within the z -group. To keep this probability nonnegative, we take f_i to be the exponential payoff-to-fitness mapping given by (4.2). Then, the probability that an imitator of

strategy \mathbf{e}_i changes strategy is just

$$\begin{aligned} p_i(\mathbf{z}) &= \frac{\gamma}{f_i(\mathbf{z})} \\ &= \gamma e^{-\omega(\mathbf{e}_i \cdot A\mathbf{z})} \end{aligned} \quad (4.9)$$

where γ is chosen so that $p_i(\mathbf{z})$ is less than 1.²⁶ Recall from (4.2) that ω is the selection strength, and that the weak selection limit corresponds to $\omega \rightarrow 0$.

We are now equipped to construct the differential equations which specify the dynamics of viability updating. At any instant in time, a player of strategy \mathbf{e}_i is imitated by a player of strategy \mathbf{e}_j with probability

$$x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_j(\mathbf{x}) \frac{z_i}{z} \quad (4.10)$$

where x_j is the probability that an \mathbf{e}_j -player is chosen as imitator, $M(\mathbf{z}; z, \mathbf{x})$ is the probability of interaction in a z -group of composition \mathbf{z} , z_i/z is the probability that an \mathbf{e}_i -player in the z -group is chosen as imitatee, and $p_j(\mathbf{x})$ is the probability that the chosen imitator changes strategy. Similarly, the probability that an \mathbf{e}_i player imitates another strategy is

$$x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_i(\mathbf{x}). \quad (4.11)$$

Thus, from (4.10) and (4.11), we get the viability updating differential equations:

$$\dot{x}_i = \sum_{j=1}^n x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_j(\mathbf{x}) \frac{z_i}{z} - x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_i(\mathbf{x}) \quad (4.12)$$

for $i = 1, \dots, n$. Recall that \mathbf{z} ranges over all n -tuples of nonnegative integers that sum to z . The form of (4.12) can be somewhat simplified by using vector notation, but the simplification is inconsequential since the system of equations is analytically intractable for n strategies and arbitrary selection strength ω . Indeed, we will see in our analysis of an eight-strategy game in §5 that we need to invoke the weak selection limit $\omega \rightarrow 0$ in order to achieve analytical results on the interior of the population simplex.

²⁶For instance, if the minimum possible payoff achieved in a z -group is some negative value $P < 0$, it suffices to let $\gamma = e^{\omega P}$.

4.3.2 Dynamical Classification

Unlike the replicator dynamics, which we derived by setting each strategy's growth rate equal to its fitness, the viability model is not so easily defined in terms of the game-dynamical growth rate function in (3.1). The viability equations (4.12) implicitly define the growth rates $\dot{x}_i/x_i = g_i(\mathbf{x})$, which are well-behaved functions because each \dot{x}_i has a zero at $x_i = 0$.²⁷ Moreover, each growth rate $g_i(\mathbf{x})$ is a continuously differentiable function on S_n because it is composed of rational and exponential functions that have no poles on the simplex. Thus, it remains to show that $\mathbf{x} \cdot \mathbf{g}(\mathbf{x}) = \sum_{i=1}^n \dot{x}_i = 0$ to verify that the viability model (4.12) is a regular game dynamics, in the sense of (3.3). We see that, as desired,

$$\begin{aligned} \sum_{i=1}^n \dot{x}_i &= \sum_{i=1}^n \sum_{j=1}^n x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_j(\mathbf{x}) \frac{z_i}{z} - \sum_{i=1}^n x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_i(\mathbf{x}) \\ &= \sum_{j=1}^n x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_j(\mathbf{x}) - \sum_{i=1}^n x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) p_i(\mathbf{x}) \\ &= 0. \end{aligned}$$

Thus, the viability model (4.12) is a regular game dynamics, and we can now ask the question of whether viability updating satisfies any stronger conditions, such as weak payoff positivity (or monotonicity) from §3.3. In fact, we will prove by a simple two-strategy counterexample in §4.4.2 that the viability model does not guarantee weak payoff positivity (nor monotonicity). Thus, viability updating may lead to evolutionary outcomes very different from the replicator dynamics, and from static game analysis.

4.4 Two-Strategy Dynamics

A brief study of two-strategy dynamics is useful both as a simple example of the replicator and viability equations, and as a means of analyzing the *edge dynamics* of multi-strategy evolutionary games. For instance, we will use these results in §5 in order to compute the dynamics of an 8-strategy game on the edges of the population simplex.

²⁷To see this, observe that when $x_i = 0$, $z_i = 0$ with probability 1.

We begin with an arbitrary two-strategy game, given by the general pay-off matrix

$$A = \begin{matrix} & \begin{matrix} 1 & 2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} a & b \\ c & d \end{pmatrix} \end{matrix}. \quad (4.13)$$

Let $x = x_1$ and $1 - x = x_2$, so that a game dynamics is specified by a single differential equation \dot{x} , defined on the simplex $S_1 = [0, 1]$. The pure population states $x = 0$ and $x = 1$, or the vertices of S_1 , are stationary by Proposition 3.2. Thus, we can classify the evolutionary behavior of two-strategy systems based on the dynamical response to perturbations from the two stationary pure states. Such a perturbation represents an invasion by mutants of one strategy in a population predominated by the other strategy, and the stability of each pure population state can be expressed concisely in terms of the population's response to invaders, who constitute an arbitrarily small population share $\varepsilon > 0$.

Proposition 4.2. For two-strategy game dynamics, there are four critical types of evolutionary behavior, based on each strategy's possible response to a ε -invasion by the other strategy:

- (a) *Dominance*: a strategy is dominant if it is asymptotically stable, or if it is stable and the other strategy is unstable. For instance, strategy 1 is dominant if

$$\dot{x} \geq 0 \text{ for } x \geq 1 - \varepsilon \quad \text{and} \quad \dot{x} \geq 0 \text{ for } x \leq \varepsilon$$

with $\dot{x} \geq 0$ holding strictly in at least one of the two cases. (Reversing this inequality gives the condition for strategy 2 dominating.)

- (b) *Bistability*: strategies 1 and 2 are bistable if they are each asymptotically stable:

$$\dot{x} > 0 \text{ for } x \geq 1 - \varepsilon \quad \text{and} \quad \dot{x} < 0 \text{ for } x \leq \varepsilon.$$

- (c) *Coexistence*: strategies 1 and 2 coexist if they are both unstable:

$$\dot{x} < 0 \text{ for } x \geq 1 - \varepsilon \quad \text{and} \quad \dot{x} > 0 \text{ for } x \leq \varepsilon.$$

- (d) *Neutrality*: strategies 1 and 2 are neutral if they are both stable, but not asymptotically so:

$$\dot{x} = 0 \text{ for } x \geq 1 - \varepsilon, x \leq \varepsilon.$$

□

It is easy to interpret the fate of invaders in each scenario in Proposition 4.2, using the fact that invaders grow, stagnate, or die, if the invaded population state is respectively unstable, stable, or asymptotically stable. A two-strategy game dynamics can be quite complex on the entirety of S_1 , with the possibility of many interior stationary states; however, Proposition 4.2 elegantly reduces all such dynamics to their evolutionary essence, in terms of simple invasion conditions. We now derive these invasion conditions explicitly in terms of game payoffs, for the replicator and viability models.

4.4.1 The Replicator Equation

When applied to the two-strategy game (4.13), the replicator dynamics (4.5) reduces to a single differential equation:

$$\begin{aligned}\dot{x} &= x[(ax + b(1 - x)) - (ax^2 + (b + c)x(1 - x) + d(1 - x)^2)] \\ &= x(1 - x)[(a - b - c + d)x + b - d]\end{aligned}\tag{4.14}$$

where $x_1 = x$ and $x_2 = 1 - x$. Thus, in addition to the stationary pure states $x = 0$, $x = 1$, the replicator equation (4.14) exhibits a stationary state,

$$x^* = \frac{d - b}{(a - c) + (d - b)}\tag{4.15}$$

which is a population state in the simplex S_1 if $d - b > 0$ and $a - c > 0$. Thus, the force of selection can change direction (\dot{x} can change sign) across at most one population state. In particular, selection must drive the population toward a pure population state, or else toward a globally attractive interior state. Indeed, we can use this potential interior equilibrium (4.15) to fully classify the two-strategy dynamics – not only in terms of the invasion condition of Proposition 4.2, but also along the interior of S_1 (Traulsen and Hauert, 2008; Nowak, 2006).

Proposition 4.3. We here fully classify the two-strategy replicator dynamics (4.14) based on the payoffs of the underlying game (4.13).

- (a) *Dominance*: a strategy is dominant in the two-strategy replicator dynamics if it is asymptotically stable on the interior of the simplex. For instance, strategy 1 is dominant if:

$$a \geq c \quad \text{and} \quad b \geq d$$

with one of the two above inequalities holding strictly. (Reversing the above inequalities gives the condition for dominance by strategy 2.)

- (b) *Bistability*: strategies 1 and 2 are bistable if they are each asymptotically stable:

$$a > c \quad \text{and} \quad b < d.$$

In this case, the population state $x^* \in \text{int}(S_1)$ given by (4.15) is an unstable interior stationary state.

- (c) *Coexistence*: strategies 1 and 2 coexist if there is a global asymptotically stable interior state, given by x^* from (4.15). This occurs if:

$$a < c \quad \text{and} \quad b > d.$$

- (d) *Neutrality*: strategies 1 and 2 are neutral if the dynamics is uniformly stationary on S_1 , which occurs if:

$$a = c \quad \text{and} \quad b = d.$$

□

Recalling the definition of ESS by Proposition 2.2, it is clear that a dominant strategy in the replicator equation (4.14) is a pure-strategy ESS of the underlying game, and that a stable state of coexistence corresponds to a mixed-strategy ESS. Proposition 4.3 thus proves that the converse of Theorem 4.1 holds for the two-strategy replicator equation: a population state is asymptotically stable if and only if the corresponding (pure or mixed) strategy is an ESS of the underlying game. The unfortunate implication of this result is that the two-strategy replicator equation does not refine our notion of game-theoretic equilibrium beyond the static ESS. However, we can also use this to our advantage, under certain circumstances, by reducing complex multi-strategy dynamics to simple pairwise interactions on the edges of the population simplex (to be discussed further in §4.4.3).

4.4.2 The Viability Equation

We now analyze viability updating for a generic two-strategy underlying game (4.13). We once again let $x_1 = x$ and $x_2 = 1 - x$, so that the viability

dynamics (4.12) reduces to:

$$\dot{x} = (1-x) \sum_{z_1=0}^z B(z_1; z, x) p_2(z_1) \frac{z_1}{z} - x \sum_{z_1=0}^z B(z_1; z, x) p_1(z_1) \frac{z-z_1}{z} \quad (4.16)$$

where $p_i(\cdot)$ is given by (4.9), and the multinomial distribution (4.8) is reduced to the binomial distribution $B(z_1; z, x)$ ²⁸ since we now sample from just two population strategies.

As in any game dynamics, the pure population states $x = 0$ and $x = 1$ are stationary in (4.16). Ohtsuki (2008) shows that, in the weak selection limit $\omega \rightarrow 0$, \dot{x} is of a cubic form, similar to (4.14). However, this is not the case for all selection strengths. For general selection strength we can analyze the behavior of the viability equation (4.16) only in terms of invasion conditions of Proposition 4.2, which we now derive in terms of the underlying game payoffs (4.13).

Assuming one strategy as invader and the other as incumbent, we can significantly simplify the viability equation (4.16). For instance, suppose we begin at a state $x = 1 - \varepsilon$, for some small population share of invaders $\varepsilon > 0$. Then the binomial sampling distribution becomes:

$$B(z_1; z, 1 - \varepsilon) \approx \begin{cases} 1 - z\varepsilon & \text{if } z_1 = z \\ z\varepsilon & \text{if } z_1 = z - 1 \\ 0 & \text{otherwise.} \end{cases} \quad (4.17)$$

This mathematical approximation translates to the fact that, for small enough proportions of invaders in the population, there is almost zero chance of having more than one invader in a random z -group.

We use (4.17) to simplify (4.16):

$$\begin{aligned} \dot{x} &= \varepsilon(p_2(z) - p_1(z-1)) \\ &= \varepsilon\gamma(e^{-\omega(zc)} - e^{-\omega((z-1)a+b)}). \end{aligned}$$

This tells us that, if strategy 1 is invaded, $\dot{x} \geq 0 \Leftrightarrow zc \leq (z-1)a + b$. By symmetry, if strategy 2 is invaded, we have $\dot{x} \geq 0 \iff zb \leq c + (z-1)d$.

²⁸The binomial distribution is given by: $B(z_1; z, x) = \binom{z}{z_1} x^{z_1} (1-x)^{(z-z_1)}$.

Thus, we have derived the following conditions for evolutionary behavior, in the terms of Proposition 4.2.

Proposition 4.4. We here fully classify the stable outcomes of a two-strategy viability updating dynamics (4.16), conditional on the payoffs of the underlying game (4.13).

(a) *Dominance*: strategy 1 dominates if

$$(z - 1)a + b \geq zc \quad \text{and} \quad zb \geq c + (z - 1)d$$

with one of the two inequalities holding strictly. (Reversing the above inequalities gives the condition for dominance by strategy 2.)

(b) *Bistability*: strategies 1 and 2 are bistable if

$$(z - 1)a + b > zc \quad \text{and} \quad zb < c + (z - 1)d.$$

(c) *Coexistence*: strategies 1 and 2 coexist if

$$(z - 1)a + b < zc \quad \text{and} \quad zb > c + (z - 1)d.$$

(d) *Neutrality*: strategies 1 and 2 are neutral if

$$(z - 1)a + b = zc \quad \text{and} \quad zb = c + (z - 1)d.$$

□

Each invasion condition above comes from one of two critical comparisons: the average payoff of strategy 1 in a z -group with one invader versus the payoff of the invader, $((z - 1)a + b)/z$ versus c , and likewise for strategy 2 being invaded by strategy 1, $(c + (z - 1)d)/z$ versus b . This is indeed what we might expect, given our approximation (4.17) that there is zero probability of having more than one invader in a z -group.

Another interesting observation is that the conditions in Propositions 4.3 and 4.4 are equivalent in the large- z limit of the viability model. In other words, as the size of the viability model's interaction groups approaches the entire population, the pairwise invasion conditions match those of the replicator equation. This result highlights the importance of local structure in

viability updating, as distinguished from the global selection mechanism of the replicator dynamics.

We now present a two-strategy example game which, as alluded to in §4.3, proves by counterexample that the viability model is not a weakly payoff positive game dynamics. Recalling the Folk Theorem 3.1, a weakly payoff positive game dynamics guarantees that any dynamically stable state is a Nash equilibrium of the underlying game. Thus, we seek to find game with a non-Nash strategy that is stable under viability updating.




Figure 1: A visualization of bifurcation in the viability dynamics of Example 4.1. Here we denote $x_{DN} = x$ and $x_{DP} = 1 - x$, so that the rate of change of DN in population state x is given by \dot{x} .

Example 4.1. We begin along the lines of Example 2.2, in which a Prisoner's Dilemma game is followed by a punishment round, allowing a player to incur a cost $\alpha > 0$ to harm an opponent by $\beta > 0$. We now consider a game between DP, a non-harming defector, and DP, a defector who punishes other defectors. This gives the payoff matrix:

$$\begin{array}{cc} & \begin{array}{cc} DN & DP \end{array} \\ \begin{array}{c} DN \\ DP \end{array} & \begin{pmatrix} 0 & -\beta \\ -\alpha & -\alpha - \beta \end{pmatrix}. \end{array}$$

Thus DN is the unique Nash equilibrium, and furthermore DN strictly dominates DP in rational game-play. Proposition 4.3 tells us that the replicator

equation respects these dominance relations, and thus DN is always dominant in the two-strategy dynamics. However, Proposition 4.4 gives us a very different result for the viability equation: DN dominates DP only if $\beta < z\alpha$; if $\beta > z\alpha$, DP dominates DN (see Figure 1). This proves that viability updating need not respect the dominance relations of the underlying game, and thus is not a weakly payoff positive game dynamics. \square

4.4.3 Nash and ESS Analogues

The preceding method of pairwise dynamical analysis can be used to classify the edge dynamics of any n -strategy system. If there are no interior stationary states in a game dynamics, then it is natural to next search the boundary of the simplex.²⁹ If there are no stable states on the interiors of each boundary face, then we can discuss long-run evolutionary behavior solely in terms of the edges of the simplex.

Recalling from Proposition 3.2 that the simplex vertices are always stationary, we can now test the stability of each stationary vertex. And evolutionary intuition justifies pairwise stability analysis: unless mutation is rampant, we can assume that any mutant strategy both arises and succeeds/fails to invade a population before the next mutation arises. In this case, we can define a meaningful edge-dynamical analogue of the pure-strategy Nash equilibrium:

Definition. A pure strategy is *uninvadeable* in the edge dynamics if it is not dominated by any other pure strategy – namely, it is dominant over, bistable with, or neutral with any other pure strategy.

Similarly, we can define analogues of the strict Nash equilibrium (uninvadeable with no neutralities) and of the ESS (uninvadeable, and able to invade any strategy that is a neutral invader). However, for the purposes of our analysis in §5, we will not need to invoke these stronger forms of edge equilibrium.

It is an easy exercise to verify that a pure strategy is Nash if and only if it is uninvadeable in the replicator dynamics (or, more generally, in any monotonic game dynamics (3.6)). But this equivalence does not hold in general, for instance under the viability equation, whose dynamics are not so

²⁹There can still be oscillation or chaotic motion in the simplex interior, but, as mentioned in §3.1, we are interested in stable stationary states as long-run evolutionary outcomes, rather than in specific solution trajectories.

clearly correlated with rational strategic game-play. Thus, in our classification of edge dynamics, we have introduced a notion of dynamical equilibrium very different from the stationary states discussed in §3.2 – although the two forms of equilibrium can work together, as we will witness in the following section.

5 Cooperation with Costly Punishment

This section presents the original research of David G. Rand, Mayuko Nakamaru, Hisashi Ohtsuki, and myself, in response to recent results (both theoretical and experimental) supporting the punishment of defectors as a means of fostering the evolution of cooperation (Hauert et al., 2007; Nakamaru and Iwasa, 2006, 2005; Fowler, 2005; Bowles and Gintis, 2004; Fehr et al., 2003; Boyd et al., 2003; Sigmund et al., 2001). However, experiments have also shown that there is in some cases a tendency to punish cooperators as well (Denant-Boemont et al., 2007; Nikiforakis, 2008). Thus, we here study an evolutionary game which broadens the scope of punishment to include the “antisocial” harming of cooperators, as well as the “spiteful” harming of all opponents indiscriminately.³⁰ Our static game analysis immediately justifies antisociality as a classically rational strategy, and our replicator dynamics analysis validates antisocial defection in an evolutionary context. Finally, our viability updating analysis allows for direct comparison with the results of Nakamaru and Iwasa (2006), and we again find that the harming of cooperators impedes the evolution of cooperation.

5.1 Game-Theoretic Analysis

We set up our game in the same manner as Examples 2.2 and 4.1: a one-round Prisoner’s Dilemma – in which a player can cooperate (C) or defect (D) – is followed by a punishment round, in which each player can harm defectors (P, punishment), harm cooperators (A, antisociality), harm both (S, spite), or harm neither (N). This makes for eight possible deterministic strategies, given the three independent choices of whether or not to cooperate, whether or not to harm a cooperator, and whether or not to harm a defector. Of these eight strategies, four are “illogical” in the sense that they harm their own kind: DP, DS, CA, and CS. The remaining four are then “logical”: DN, DA, CN, and CP.

³⁰The terminology of past literature poses a slight problem when we generalize the strategy set. The term “punishment” specifically refers to the harming of defectors, but we will also use it as the general term for harm, whenever a phrase (such as “Cooperation with Costly Punishment,” or “the punishment round”) has been well-established in past literature that only considered harming defectors. The meaning is intended to be clear from context.

The payoffs to the Prisoner’s Dilemma are given by the standard payoff matrix (1.1), whereby a cooperator pays a cost c to give an opponent a benefit $b > c > 0$. The payoffs to the punishment round, in which a player can incur a cost $\alpha > 0$ to decrease an opponent’s payoff by $\beta > 0$, are then added to the first-round payoffs. Thus total game payoffs are given by the 8×8 payoff matrix (B.1) in Appendix B.

5.1.1 Nash Equilibria and ESS

We employ the fundamental evolutionary assumption that players’ strategies are inherent (or otherwise fixed), and only altered through reproduction (or imitative updating). In light of this fact, and because weak Nash equilibria are invadeable in an evolutionary context, we are primarily concerned with pure-strategy Nash equilibria. Mixed-strategy ESS’s, on the other hand, are of interest, for they can be interpreted as stable states in an evolving population.

Recalling the criterion (2.5) for symmetric Nash equilibria, we see from payoff matrix (B.1) that DN and DA are always Nash strategies, and that CP is Nash if $\beta > c$. This shows immediately that antisocial defection is a rational strategy in the game, thereby validating our intuition to expand the game’s strategy space to include the harming of cooperators.

It is easy to see that these are the only three symmetric pure-strategy Nash equilibria, and so we now search the game for possible ESS’s. Recalling Proposition 2.5, an ESS cannot contain any Nash strategies in its support. Thus, there cannot be any ESS in the interior of the simplex S_8 , since we have pure-strategy Nash equilibria. Also, observe that our three Nash strategies can only be weak Nash equilibria, and not ESS’s: DN and DA, as well as CN and CP, constitute neutral strategy pairs, and neutral strategies are not robust to mutual invasion (cf. ESS definition by Proposition 2.2).³¹

Next, observe that the indiscriminate harmers, DS and CS, are strictly dominated by DN and CN (resp.) by definition (2.6). Proposition 2.1 then tells us that neither DS nor CS can appear in the support of any Nash equilibrium – hence of any ESS, by Proposition 2.3. Also, the illogical punishers, DP and

³¹DP-DS and CA-CS also constitute somewhat less noteworthy neutral strategy pairs.

CA, are weakly dominated by DN and CN (resp.), and thus cannot appear in the support of any ESS by Proposition 2.4.

Thus, we have narrowed any potential ESS to a single strategy: CN, the non-punishing cooperator. But CN is always invaded by DN, and therefore there is no ESS in our game. We now move from this static analysis to the evolutionary dynamics, which will perhaps reveal some invasion-robust equilibrium that our static game analysis did not reveal.

5.2 Replicator Dynamics Analysis

Here we analyze the replicator dynamics (4.5) for our underlying 8-strategy cooperation and punishment game, with payoff matrix $A = (B.1)$.

5.2.1 Dynamical Equilibria

Before any explicit computation, we can predict much of the behavior of the replicator dynamics from our static analysis of last section, as well as the theoretical results of §3 and §4. Recall that the Folk Theorem offers us three dynamical refinements of the Nash equilibrium: interior stationary states, stable states, and Ω -limits of interior states. From Theorem 3.2, the population shares x_{DS} and x_{CS} must always converge to zero, since DS and CS are strictly dominated in rational game-play, so immediately we know that there can be no interior stationary states. By Proposition 3.4, this also implies that if the Ω -limit of an interior state is a single limit point, the limit point lies on the boundary of the simplex. Thus, the replicator dynamics can only yield boundary equilibria for our game.

Recall that Proposition 3.6 gives the stationary states of any monotonic game dynamics, such as the replicator dynamics:

$$S_n^\circ = \{\mathbf{x} \in S_n \mid \mathbf{e}_i \cdot A\mathbf{x} = \mathbf{x} \cdot A\mathbf{x}, \forall \mathbf{e}_i \in \mathcal{S}(\mathbf{x})\}.$$

As noted after the proof of this proposition in §3.3, S_n° can be explicitly computed by solving finitely many systems of linear equations. In particular, we seek non-edge stationary states (i.e. population states involving more than two pure strategies), since the edge dynamics are described in much greater detail by Proposition 4.3, which we will invoke in the next section. Thus, we

compute the non-edge stationary states of our system.³² As expected, none lie in the interior of the population simplex:

If $\beta > \alpha + 2c$:

$$(1) \quad x_{DP} = \frac{1}{4} + \frac{\alpha+2c}{4\beta}, \quad x_{DA} = \frac{1}{4} - \frac{\alpha+2c}{4\beta}, \quad x_{CS} = \frac{1}{2}.$$

$$(2) \quad x_{CP} = \frac{1}{4} + \frac{\alpha+2c}{4\beta}, \quad x_{CA} = \frac{1}{4} - \frac{\alpha+2c}{4\beta}, \quad x_{DN} = \frac{1}{2}.$$

If $\beta > |\alpha - 2c|$:

$$(3) \quad x_{DP} = \frac{1}{4} - \frac{\alpha-2c}{4\beta}, \quad x_{DA} = \frac{1}{4} + \frac{\alpha-2c}{4\beta}, \quad x_{CN} = \frac{1}{2}.$$

$$(4) \quad x_{CP} = \frac{1}{4} - \frac{\alpha-2c}{4\beta}, \quad x_{CA} = \frac{1}{4} + \frac{\alpha-2c}{4\beta}, \quad x_{DS} = \frac{1}{2}.$$

A stationary parameterized curve exists for $x \in \left[\frac{c}{2\beta}, \frac{1}{2}\right]$, if $\beta > c$:³³

$$(5) \quad x_{CP} = x, \quad x_{CA} = \frac{1}{2} - x, \quad x_{DP} = \frac{1}{2} - x + \frac{c}{2\beta}, \quad x_{DA} = x - \frac{c}{2\beta}.$$

We evaluate the stability of the stationary states (1)-(5) by the eigenvalue method discussed in §3.2.2. Since the Jacobian $\mathcal{J}(\mathbf{x})$ of our system is an 8×8 matrix, the roots of its 8th-degree characteristic polynomial cannot be solved analytically. Thus, for each stationary state \mathbf{x} we evaluate $\mathcal{J}(\mathbf{x})$ across 100,000 random sets of parameter values b, c, α , and β ,³⁴ and use Matlab to compute numerical eigenvalues. We find that all stationary states (1)-(4), as well as all points on the stationary curve (5),³⁵ have at least one eigenvalue with positive real part for every tested parameter set. Therefore, we conclude that all non-edge equilibria of our game are unstable for reasonable parameter values, and we restrict further analysis to the dynamics of the simplex edges, as justified in §4.4.3.

³²Matlab code and other computations available on request.

³³By “stationary parameterized curve,” we simply mean a parameterized curve $\gamma : \left[\frac{c}{2\beta}, \frac{1}{2}\right] \rightarrow S_n$, with $\gamma(x)$ stationary in the replicator dynamics for any $x \in \left[\frac{c}{2\beta}, \frac{1}{2}\right]$.

³⁴Each random parameter set was generated by choosing a random uniform value for each parameter on the interval $[0, 20]$, with the sole constraint that $b > c$, in accordance with Prisoner’s Dilemma payoffs.

³⁵For each parameter set, 1,000 random states on the curve were chosen for numeric stability evaluation.

5.2.2 Edge Dynamics

We can explicitly compute all pairwise replicator dynamics using Proposition 4.3, which specifies not only the invasion conditions for each strategy pair, but also the full dynamics along each simplex edge. We thus compute the following equilibria:³⁶

- (1) Stationarity along the neutral DN-DA edge.
- (2) Stationarity along the neutral CN-CP edge.
- (3) Stationarity along the neutral DP-DS edge.
- (4) Stationarity along the neutral CA-CS edge.

If $\beta > c$ (*harm is effective*):

- (5) Bistability between DN and CP, with unstable stationary state:
 $x_{DN} = \frac{\beta-c}{\alpha+\beta}$, $x_{CP} = 1 - \frac{\beta-c}{\alpha+\beta}$.
- (6) Coexistence between DP and CS.

If $\alpha > c$ (*harm is expensive*):

- (7) Bistability between DA and CN.
- (8) Bistability between DS and CA.

If $\alpha + \beta > c$ (*harm is expensive and/or effective*):

- (9) Bistability between DA and CP, with unstable stationary state:
 $x_{DA} = \frac{1}{2} \frac{\alpha+\beta-c}{\alpha+\beta}$, $x_{CP} = 1 - \frac{1}{2} \frac{\alpha+\beta-c}{\alpha+\beta}$.
- (10) Bistability between DS and CP.
- (11) Coexistence between DP and CN.
- (12) Coexistence between DP and CA.

³⁶For the sake of brevity here, we only state the specific location of an equilibrium in two special cases, which are of particular interest.

The remaining edges exhibit dominance relations, rather than dynamical equilibria. In Figure 2 (see Appendix B.2) we summarize the complete edge dynamics, including the equilibria (1)-(12) and the remaining edges. The edge dynamics change across parameter space, based on the inequality conditions in (1)-(12) above, and thus Figure 2 depicts the edge dynamics for the resulting five different parameter regions.

Figure 2 corroborates the observation in §4.4.3 that any Nash equilibrium in the underlying game is uninvadable in the replicator edge dynamics: DN and DA are uninvadable in all parameter regions, and CP is uninvadable whenever $\beta > c$. Since there is no ESS in the underlying game, there is never a dominant strategy in the edge dynamics. Nonetheless, the logical defectors DN and DA form asymptotically stable edge, in the sense that any other invading strategy is eliminated anywhere along the DN-DA edge. CP is also uninvadable, yet vulnerable by way of its neutrality with CN, which can always be invaded by DN. In the presence of mutation, neutral drift along the CN-CP edge makes the demise of cooperation inevitable, even when CP is itself uninvadable. Thus, evolution will always lead to “logical defection,” with a randomly drifting combination of non-harming defectors and antisocial defectors.

To distinguish the evolutionary success of DN and DA, we compare their relative basins of attraction with respect to the stable CP vertex – in other words, we look at the location of the unstable coexistences (5) and (9) above. When $\beta > \alpha + c$, the unstable DA-CP equilibrium is further from DA than the unstable DN-CP equilibrium is from DN. This results in DA having a larger basin of attraction than DN when playing against CP. Also, the DA-CP equilibrium is never closer to DA than $DA = \frac{1}{2}$, while the DN-CP equilibrium can go all the way to $DN = 1$, in which case CP dominates DN. Both of these facts show that if inflicting harm is inexpensive, DA can in some sense outperform DN against CP.

Thus, we have shown that the replicator dynamics leads to the evolution of antisocial defection, strengthening the result achieved by the Nash and ESS analysis of the static game. The replicator dynamics also offers further insight in showing that antisocial defectors often fare better than non-harming defectors against punishing cooperators.

5.3 Viability Updating Analysis

Nakamaru and Iwasa (2006) considers a game between DN, DP, CN, and CP, and shows that the CN-CP edge is stable under the viability updating dynamics (4.12). We now investigate whether this result holds in the more general 8-strategy game.

5.3.1 Weak Selection Equilibria

Finding analytical solutions for the viability model for general selection strength ω is problematic.³⁷ Thus, we relegate ourselves to the weak selection limit $\omega \rightarrow 0$ (cf. (4.3)), which simplifies the viability differential equations to a more manageable form:³⁸

$$\dot{x}_i = x_i(\mathbf{e}_i \cdot [zA - A^T]\mathbf{x} - \mathbf{x} \cdot [zA - A^T]\mathbf{x}) \quad (5.1)$$

where A (with transpose A^T) is the payoff matrix of the underlying game. Note that (5.1) is equivalent to the replicator dynamics under the change of payoff matrix $A \rightarrow zA - A^T$, which represents a change in fitness in the replicator equations (4.4). Thus, weak selection viability updating can be thought of as replicator dynamics up to a change in fitness: “replicator fitness” is an individual’s expected payoff from playing a random opponent; “viability fitness” under weak selection is the difference between an individual’s total expected payoff from playing z random opponents, and the expected payoff a single random opponent receives against the individual. Consequently, viability updating rewards strategies that lower opponents’ payoffs – thereby foreshadowing the evolution of costly punishment.

It is interesting to note that, just as was the case with the invasion conditions in Propositions 4.3 and 4.4, the replicator and weak-selection viability equations are equivalent in the large z limit. This once again highlights the importance of finite interaction groups in z -mixed viability updating.

Now, we can once again use Proposition 3.6 to compute the set of stationary states S_n^o of our system – only, this time we use $zA - A^T$ as our payoff matrix,

³⁷In fact, in our research efforts, we have not even succeeded as yet in running simulations by numerical integration, for the 8-strategy viability equations demand an excess of computing power.

³⁸See Appendix C for derivation.

to obtain stationary states for weak-selection viability updating. We find the following non-edge equilibria:

If $\beta > \frac{\alpha}{z}$, $2\frac{b+zc}{z+1} + \alpha$:

$$(1) \quad x_{DP} = \frac{1}{4} \frac{(z-1)(\beta+\alpha)+2(b+zc)}{z\beta-\alpha}, \quad x_{DA} = \frac{1}{4} \frac{(z+1)(\beta-\alpha)-2(b+zc)}{z\beta-\alpha}, \quad x_{CS} = \frac{1}{2}.$$

$$(2) \quad x_{CP} = \frac{1}{4} \frac{(z-1)(\beta+\alpha)+2(b+zc)}{z\beta-\alpha}, \quad x_{CA} = \frac{1}{4} \frac{(z+1)(\beta-\alpha)-2(b+zc)}{z\beta-\alpha}, \quad x_{DN} = \frac{1}{2}.$$

If $\beta > \frac{\alpha}{z}$, $\alpha - 2\frac{b+zc}{z+1}$, $2\frac{b+zc}{z-1} - \alpha$ or $\beta < \frac{\alpha}{z}$, $\alpha - 2\frac{b+zc}{z+1}$, $2\frac{b+zc}{z-1} - \alpha$:

$$(3) \quad x_{DP} = \frac{1}{4} \frac{(z+1)(\beta-\alpha)+2(b+zc)}{z\beta-\alpha}, \quad x_{DA} = \frac{1}{4} \frac{(z-1)(\beta+\alpha)-2(b+zc)}{z\beta-\alpha}, \quad x_{CN} = \frac{1}{2}.$$

$$(4) \quad x_{CP} = \frac{1}{4} \frac{(z+1)(\beta-\alpha)+2(b+zc)}{z\beta-\alpha}, \quad x_{CA} = \frac{1}{4} \frac{(z-1)(\beta+\alpha)-2(b+zc)}{z\beta-\alpha}, \quad x_{DS} = \frac{1}{2}.$$

A stationary parameterized curve exists for $x \in \left[\frac{1}{2} \frac{b+zc}{z\beta-\alpha}, \frac{1}{2} \right]$, if $\beta > \frac{\alpha+(b+zc)}{z}$:

$$(5) \quad x_{CP} = x, \quad x_{CA} = \frac{1}{2} - x, \quad x_{DP} = \frac{1}{2} - x + \frac{1}{2} \frac{b+zc}{z\beta-\alpha}, \quad x_{DA} = x - \frac{1}{2} \frac{b+zc}{z\beta-\alpha}.$$

To evaluate the stability of equilibria (1)-(5), we once again numerically compute the eigenvalues of our system's 8×8 Jacobian, for 100,000 random parameter sets. As was the case with the replicator dynamics, we find that all non-edge equilibria have at least one eigenvalue with positive real part. Therefore, all interior equilibria are unstable for reasonable parameter values, and we restrict further analysis to the edge dynamics.

5.3.2 Edge Dynamics

We now analyze the viability edge dynamics by computing all pairwise invasion conditions, as given by Proposition 4.4. Recall that for the viability model, we cannot specify the exact dynamics along each edge – as we could for the replicator dynamics – since the two-strategy viability equation (4.16) is not of a “nice” polynomial form. We compute the pairwise equilibria:

- (1) Neutrality between DN and DA.
- (2) Neutrality between CN and CP.
- (3) Neutrality between DP and DS.

(4) Neutrality between CA and CS.

If $\beta \in \left(\frac{\alpha+b}{z} + c, z(c + \alpha) + b\right)$:

(5) Bistability between DN and CP.

(6) Coexistence between DP and CS.

If $\beta \in \left(\frac{\alpha-b}{z} - c, z(\alpha - c) - b\right)$:

(7) Bistability between DA and CN.

(8) Bistability between DS and CA.

If $\beta > \frac{b+zc}{z-1} - \alpha$:

(9) Bistability between DA and CP.

(10) Bistability between DS and CP.

(11) Coexistence between DP and CN.

(12) Coexistence between DP and CA.

The remaining edges exhibit dominance relations, rather than dynamical equilibria. In Figure 3 we summarize the invadeability conditions on all edges of the population simplex, including the equilibria (1)-(12) and all other edges. We end up with 12 different parameter regions, and we can compare our results directly with the findings of Nakamaru and Iwasa (2006) by setting $\alpha = c = 1$ and $z = 4$.

With these values assumed for α, c , and z , the only feasible parameter regions (i.e. regions in which $b, \beta > 0$) are the regions which we have labeled 1-7 in Figure 3, in accordance with Nakamaru and Iwasa (2006). Our results within each region are extremely different from those of Nakamaru and Iwasa (2006): in Regions 3, 4, and 5, Nakamaru and Iwasa find that CP is the only uninvadeable strategy, whereas we find that DS is the only uninvadeable strategy. We find a relatively small portion of parameter space (regions 6 and 8) in which CP is uninvadeable, but nowhere is CP the unique uninvadeable strategy.

Moreover, in every region save region 12 (where DN is the unique uninvadeable strategy), either DA or DS is uninvadeable, so that defectors who harm cooperators are almost always evolutionarily stable. And even when CP is uninvadeable, it is nevertheless neutral with the perpetually invadeable strategy CN, so that the edge of logical cooperators is always unstable. Thus, in the presence of mutation, CP cannot be stably maintained by a population in the long-run.

5.4 Discussion of Results

Thus, using static game analysis as well as two models for evolutionary dynamics, we have shown that the possibility of harming cooperators dramatically changes the outcomes of evolution. Our initial game-theoretic analysis supports antisocial defection as a rational strategy, in the traditional sense. The replicator dynamics suggests that the only sustainable long-run evolutionary outcome involves some combination of non-harming defectors and antisocial defectors. Moreover, in the replicator dynamics antisocial defectors fare better than non-harming defectors against their most formidable opponent, punishing cooperators, thus further supporting antisocial defection in an evolutionary context.

The results of our 8-strategy viability updating are very different from Nakamaru and Iwasa's 4-strategy results, as we find that various forms of defection always prevail in the evolutionary long-run. Most surprisingly, the strategy of spiteful defection – which is both “irrational,” as a strictly dominated strategy, and “illogical,” as a strategy that harms its own kind – is found to be long-run stable when the effect of harm is severe. This counterintuitive outcome can be attributed to the finite interaction groups of viability updating: a strategy that reduces others' fitness increases opponents' chance of strategic update, and is thereby more likely to spread in small interaction groups (Hofbauer and Sigmund, 1998; Nakamaru and Iwasa, 2006).

Thus, we have achieved interesting results by expanding the strategy space of recent evolutionary models, and by applying the analytical methods of §2-4. Our analysis can be further developed by the consideration of different evolutionary dynamics, perhaps ones that involve spatially-structured populations or other forms of agent-based interactions in finite populations. While we have hopefully shed some light on the merits and drawbacks of costly pun-

ishment in a cooperation game, more can always be said, and new questions ever arise, in the quest to understand the evolution of cooperation.

6 Conclusion

In §2.1 we introduced the key notions from static game theory, in particular the Nash equilibrium and its chief evolutionary refinement, the ESS.

In §3 we developed the rudiments of evolutionary dynamics for continuum populations. We proved necessary conditions for a game dynamics to exhibit dynamical refinements of the Nash equilibrium, and to ensure the evolutionary elimination of irrational strategies. We witnessed in §4 two specific dynamical models which aptly demonstrate the advantages and disadvantages of a game dynamics firmly rooted in rational game-play. We also introduced a key method of analyzing the edge dynamics of a system when there are no interior stationary states.

§5 presented the central motivating problem, which served as both a worthwhile original evolutionary analysis and an extended application of the static and dynamic theories earlier developed. Our analysis yielded interesting results that challenge the use of costly punishment as a means of inducing the evolution of cooperation.

We thus conclude our study, and hope that the reader is encouraged to further explore the ever-growing field of evolutionary game dynamics.

Appendix

A Normal-Form Games and Nash Existence

Much of the initial complexity of general normal-form game theory, as compared with the symmetric two-player game theory presented in §2.1, derives from the complex notation needed to keep track of an arbitrary number of asymmetric strategy spaces and payoff functions. The introduction of general game theory is omitted from the main text because it offers us no immediate insight that is not already provided by the theory of symmetric games. However, the ensuing discussion is of interest to the reader who seeks to understand the rudiments of strategic interactions and Nash equilibria in their purest form – which might perhaps serve to better motivate the development of evolutionary theory in the body of the paper. Also, note that this appendix is intended to be read after §2.1, and thus we here omit the explanation of certain concepts that have already been elucidated by the introduction of symmetric game theory.

First recall that a normal-form game $\mathcal{G} = (m, \{\mathbb{S}^i\}, \{\pi_i\})$ is defined by m players, each with pure strategy space \mathbb{S}^i and payoff function $\pi_i : \times_{i=1}^m \mathbb{S}^i \rightarrow \mathbb{R}$, where $\times_{i=1}^m \mathbb{S}^i$ is the space of all *pure strategy profiles* $\mathbf{s} = (s^1, \dots, s^m)$.

Player i employs a mixed strategy σ^i by playing each pure strategy $s^i \in \mathbb{S}^i$ with fixed probability $\sigma^i(s^i)$. We define, once again, the support of a mixed strategy as the set of pure strategies with nonzero probability weights:

$$\mathcal{S}(\sigma^i) = \{s^i \in \mathbb{S}^i \mid \sigma^i(s^i) > 0\}.$$

Each player possesses a (possibly unique) mixed strategy space $\Sigma^i \supset \mathbb{S}^i$, so that the space of all strategy profiles – the potential strategic states of the game – is the product space $\Sigma = \times_{i=1}^m \Sigma^i \supset \times_{i=1}^m \mathbb{S}^i$. As a matter of notation, we say that a strategy profile $\boldsymbol{\sigma} = (\sigma^1, \dots, \sigma^m) \in \Sigma$ can be decomposed into σ^i (player i 's strategy) and σ^{-i} (all other players' strategies). Player i 's payoff to a mixed strategy profile $\boldsymbol{\sigma} \in \Sigma$ is given by the probabilistic sum of payoffs over all pure strategy profiles:

$$u_i(\boldsymbol{\sigma}) \equiv \sum_{\mathbf{s} \in \mathbb{S}} \Pr(\mathbf{s}) \pi_i(\mathbf{s}) = \sum_{\mathbf{s} \in \mathbb{S}} \left(\prod_{j=1}^m \sigma^j(s^j) \right) \pi_i(\mathbf{s}). \quad (\text{A.1})$$

The following proposition gives two important properties of the payoff function that we will employ in the ensuing discussion of the Nash equilibrium. The proposition follows immediately from (A.1), which is a generalization of the bilinearity of the two-player symmetric payoff function defined in §2.1.

Proposition A.1. Player i 's payoff $u_i(\boldsymbol{\sigma})$ is linear in the probabilities $\sigma^i(s^i)$ he assigns to the pure strategies $s^i \in \mathbb{S}^i$, and hence linear in player i 's mixed strategy $\sigma^i \in \Sigma^i$ itself.

We now define the notions of best reply and Nash equilibrium on their broadest terms.

Definition. Given a strategy profile $\boldsymbol{\sigma} \in \Sigma$, player i 's *best-reply* β_i is given by

$$\beta_i(\boldsymbol{\sigma}) = \{\hat{\sigma}^i \in \Sigma^i \mid u_i(\hat{\sigma}^i, \sigma^{-i}) \geq u_i(s^i, \sigma^{-i}), \forall s^i \in \mathbb{S}^i\}. \quad (\text{A.2})$$

Once again, there are possibly multiple best-replies to a given strategy profile, and thus player i 's best reply is a set-valued function $\beta_i : \Sigma \rightarrow \mathcal{P}(\Sigma^i)$.

We can now define a global best-reply function $\boldsymbol{\beta} : \Sigma \rightarrow \mathcal{P}(\Sigma)$, given by the cartesian product of individual best-reply functions, $\boldsymbol{\beta}(\boldsymbol{\sigma}) = \times_{i=1}^m \beta_i(\boldsymbol{\sigma})$.

Definition. A strategy profile $\boldsymbol{\sigma} \in \Sigma$ is a *Nash equilibrium* if $\boldsymbol{\sigma} \in \boldsymbol{\beta}(\boldsymbol{\sigma})$, i.e. if for all players i ,

$$u_i(\sigma^i, \sigma^{-i}) \geq u_i(s^i, \sigma^{-i}) \quad \forall s^i \in \mathbb{S}^i. \quad (\text{A.3})$$

The Nash equilibrium is *strict* if the global best reply is a singleton, $\boldsymbol{\beta}(\boldsymbol{\sigma}) = \{\boldsymbol{\sigma}\}$, i.e. if (A.3) holds strictly.

Again, a Nash equilibrium is simply a strategic state of the game from which no player can profitably deviate. The famous Nash Existence Theorem guarantees that every *finite* normal-form game has a Nash equilibrium (Nash, 1951). This is an important nuance of the theorem: a game is considered finite if each player's pure strategy space is finite. Letting $|\mathbb{S}^i| = n_i$, player i 's mixed-strategy space Σ^i becomes the space of all probability weightings across n_i pure strategies; namely, Σ^i is isomorphic to the $(n_i - 1)$ -dimensional simplex $S_{n_i} \subset \mathbb{R}^{n_i}$.

Recall from §2.1 the statement of Kakutani's theorem, which is central to the Nash existence theorem:

Theorem (Kakutani Fixed Point Theorem). *Let S be a nonempty, convex, and compact subset of a Euclidean space, and $\beta : S \rightarrow \mathcal{P}(S)$ a set-valued function. If β has convex nonempty images and a closed graph, then it has a fixed point $x \in \beta(x)$.*

We are now equipped to prove the most general version of Nash's theorem.

Theorem (Nash Existence Theorem). *Every finite normal-form game has a Nash equilibrium.*

Proof. By definition (A.3), a fixed point of a game's global best reply function $\beta : \Sigma \rightarrow \mathcal{P}(\Sigma)$ is a Nash equilibrium of the game. Thus it suffices to show that for a finite game, β satisfies the conditions of Kakutani's theorem.

- (1) *The domain of β is a nonempty, convex, and compact subset of a Euclidean space.*

Clearly, the simplex S_{n_i} is a nonempty, convex, and compact (closed and bounded) subset of the Euclidean space \mathbb{R}^{n_i} . Thus the domain of β is (via isomorphism) a nonempty, convex, and compact subset of a Euclidean space, as required: $\Sigma \simeq \times_{i=1}^m S_{n_i} \subset \times_{i=1}^m \mathbb{R}^{n_i}$.

- (2) *$\beta(\sigma)$ is nonempty, $\forall \sigma \in \Sigma$.*

By Proposition A.1, player i 's payoff $u_i(\sigma^i, \sigma^{-i})$ is a linear (hence continuous) function of his strategy $\sigma^i \in \Sigma^i \simeq S_{n_i}$. Since u_i is continuous on the compact set Σ^i , it attains a maximum for some $\hat{\sigma}^i \in \Sigma^i$. It then follows from the definition of best reply (A.2) that, for any given $\sigma \in \Sigma$, there exists some $\hat{\sigma}^i \in \beta_i(\sigma)$, for each i . Then $\hat{\sigma} = (\hat{\sigma}^1, \dots, \hat{\sigma}^m) \in \beta(\sigma)$.

- (3) *$\beta(\sigma)$ is convex, $\forall \sigma \in \Sigma$.*

For a given $\sigma \in \Sigma$, let $\sigma_1^i, \dots, \sigma_k^i \in \beta_i(\sigma)$ be best replies for player i ; that is to say, for all $j \in \{1, \dots, k\}$, $u_i(\sigma_j^i, \sigma^{-i}) = \lambda$ is the maximal payoff that player i can achieve against σ^{-i} . Then, any convex combination $\alpha_1, \dots, \alpha_k$ of best replies is also a best reply, by Proposition A.1: $u_i(\sum_{j=1}^k \alpha_j \sigma_j^i, \sigma^{-i}) = \sum_j \alpha_j u_i(\sigma_j^i, \sigma^{-i}) = \sum_j \alpha_j \lambda = \lambda$. Thus each $\beta_i(\sigma)$ is convex in Σ^i , so the image $\beta(\sigma)$ is convex in the product space Σ .

- (4) *β has a closed graph.*

Recall that β has a *closed graph* if: $(\{\sigma_\alpha\}, \{\hat{\sigma}_\alpha\}) \rightarrow (\sigma, \hat{\sigma})$, with $\hat{\sigma}_\alpha \in$

$\beta(\sigma_\alpha) \Rightarrow \hat{\sigma} \in \beta(\sigma)$. We prove by contradiction that β has a closed graph using the following “three-epsilon” argument. Suppose that $(\{\sigma_\alpha\}, \{\hat{\sigma}_\alpha\}) \rightarrow (\sigma, \hat{\sigma})$ but that $\hat{\sigma} \notin \beta(\sigma)$. Then $\hat{\sigma}^i \notin \beta_i(\sigma)$ for some player i , so that there exists an $\varepsilon > 0$ and a $\bar{\sigma}^i \in \Sigma^i$ such that $u_i(\bar{\sigma}^i, \sigma^{-i}) > u_i(\hat{\sigma}^i, \sigma^{-i}) + 3\varepsilon$. Since $u_i(\sigma)$ is continuous in σ by Proposition A.1, and $(\{\sigma_\alpha\}, \{\hat{\sigma}_\alpha\}) \rightarrow (\sigma, \hat{\sigma})$, we also have that for α sufficiently large: $u_i(\bar{\sigma}^i, \sigma_\alpha^{-i}) > u_i(\bar{\sigma}^i, \sigma^{-i}) - \varepsilon$, and $u_i(\hat{\sigma}^i, \sigma^{-i}) > u_i(\hat{\sigma}_\alpha^i, \sigma_\alpha^{-i}) - \varepsilon$. Combining all three inequalities:

$$u_i(\bar{\sigma}^i, \sigma_\alpha^{-i}) > u_i(\bar{\sigma}^i, \sigma^{-i}) - \varepsilon > u_i(\hat{\sigma}^i, \sigma^{-i}) + 2\varepsilon > u_i(\hat{\sigma}_\alpha^i, \sigma_\alpha^{-i}) + \varepsilon$$

which contradicts the fact that $\hat{\sigma}_\alpha^i \in \beta_i(\sigma_\alpha)$. Thus, it must be the case that β has a closed graph.

Thus, by the Kakutani fixed point theorem, the global best-reply function of any finite normal-form game has a fixed point $\sigma \in \beta(\sigma)$, which constitutes a Nash equilibrium of the game. □

B 8-Strategy Cooperation / Punishment Game

B.1 Payoff Matrix

$$\begin{array}{l}
 DN \\
 DP \\
 DA \\
 DS \\
 CN \\
 CP \\
 CA \\
 CS
 \end{array}
 \begin{pmatrix}
 DN & DP & DA & DS & CN & CP & CA & CS \\
 0 & -\beta & 0 & -\beta & b & b-\beta & b & b-\beta \\
 -\alpha & -\alpha-\beta & -\alpha & -\alpha-\beta & b & b-\beta & b & b-\beta \\
 0 & -\beta & 0 & -\beta & b-\alpha & b-\alpha-\beta & b-\alpha & b-\alpha-\beta \\
 -\alpha & -\alpha-\beta & -\alpha & -\alpha-\beta & b-\alpha & b-\alpha-\beta & b-\alpha & b-\alpha-\beta \\
 -c & -c & -c-\beta & -c-\beta & -c+b & -c+b & -c+b-\beta & -c+b-\beta \\
 -c-\alpha & -c-\alpha & -c-\alpha-\beta & -c-\alpha-\beta & -c+b & -c+b & -c+b-\beta & -c+b-\beta \\
 -c & -c & -c-\beta & -c-\beta & -c+b-\alpha & -c+b-\alpha & -c+b-\alpha-\beta & -c+b-\alpha-\beta \\
 -c-\alpha & -c-\alpha & -c-\alpha-\beta & -c-\alpha-\beta & -c+b-\alpha & -c+b-\alpha & -c+b-\alpha-\beta & -c+b-\alpha-\beta
 \end{pmatrix}
 \tag{B.1}$$

B.2 Replicator and Viability Edge Dynamics

In the ensuing two figures: red arrows indicate one strategy dominating another; black arrows indicate bistability; blue arrows indicate coexistence; uninvadeable strategies are colored green.




Figure 2: Replicator dynamics on the simplex edges in each parameter region.




Figure 3: Viability updating invasion conditions on the simplex edges in each parameter region.

C Viability Updating: Weak-Selection Limit

We here follow the derivation of Ohtsuki (2008) to obtain the weak-selection viability expression (5.1). We begin by taking the weak selection limit $\omega \rightarrow 0$, and substituting the resulting weak-selection fitness (4.3) into the general-selection viability equations (4.12):

$$\begin{aligned} \dot{x}_i &\approx \sum_{j=1}^n x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) (1 - \omega(\mathbf{e}_j \cdot A\mathbf{z})) \frac{z_i}{z} \\ &\quad - x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) (1 - \omega(\mathbf{e}_i \cdot A\mathbf{z})). \end{aligned}$$

The term of 0th-order in β vanishes, giving:

$$\begin{aligned} \omega^{-1} \dot{x}_i &\approx x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) (\mathbf{e}_i \cdot A\mathbf{z}) - \sum_j x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) (\mathbf{e}_j \cdot A\mathbf{z}) \frac{z_i}{z} \\ &= x_i \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) \sum_k a_{ik} z_k - \frac{1}{z} \sum_j x_j \sum_{\mathbf{z}} M(\mathbf{z}; z, \mathbf{x}) \sum_k a_{jk} z_k z_i \\ &= x_i \sum_k a_{ik} E[Z_k] - \frac{1}{z} \sum_j x_j \sum_k a_{jk} E[Z_i Z_k] \end{aligned} \quad (\text{C.1})$$

where $\mathbf{Z} = (Z_1, \dots, Z_n) \sim M(\mathbf{z}; z, \mathbf{x})$ is a multinomial random variable, and $E[\cdot]$ the expectation function. It is a standard result from probability theory that $E[Z_i] = zx_i$, and $E[Z_i Z_k] = zx_i(\delta_{ik} + (z-1)x_k)$ (Rice, 2007), where δ_{ik} is the Kronecker delta function. Thus, (C.1) simplifies to:

$$\begin{aligned} &= x_i \sum_k a_{ik} zx_k - \frac{1}{z} \sum_j x_j \sum_k a_{jk} zx_i (\delta_{ik} + (z-1)x_k) \\ &= x_i \left(z \sum_k a_{ik} x_k - \sum_{j,k} a_{jk} (\delta_{ik} x_j + (z-1)x_j x_k) \right) \\ &= x_i \left(z \sum_k a_{ik} x_k - \sum_j x_j a_{ji} - (z-1) \sum_{j,k} a_{jk} x_j x_k \right) \\ &= x_i (z \mathbf{e}_i \cdot A\mathbf{x} - \mathbf{e}_i \cdot A^T \mathbf{x} - (z-1) \mathbf{x} \cdot A\mathbf{x}) \\ &= x_i (\mathbf{e}_i \cdot [zA - A^T] \mathbf{x} - \mathbf{x} \cdot [zA - A^T] \mathbf{x}). \end{aligned} \quad (\text{C.2})$$

This gives us (5.1), up to a universal velocity factor of β^{-1} .

D Statement of Auxiliary Theorems

D.1 Picard-Lindelöf Theorem

Theorem. *If $\varphi : X \rightarrow \mathbb{R}^n$ is Lipschitz continuous on the open domain $X \subset \mathbb{R}^n$, then the system $\dot{\mathbf{x}} = \varphi(\mathbf{x})$ has a unique solution $\boldsymbol{\xi}(\cdot, \mathbf{x}) : T \rightarrow X$ through every state $\mathbf{x} \in X$. Moreover, $\boldsymbol{\xi}$ is continuous in $t \in T$ and in $\mathbf{x} \in X$.*³⁹

See Hirsch and Smale (1974) for proof.

D.2 Hartman Linearization Theorem

We must first establish the following two definitions before stating Hartman's theorem.

Definition. In a system of differential equations $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, a stationary state \mathbf{y} is *hyperbolic* if the Jacobian \mathcal{J} of \mathbf{f} evaluated at \mathbf{y} has no eigenvalues with zero real part.

Definition. Given a system of differential equations $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, a *linearization* L of \mathbf{f} at a point \mathbf{y} is the system given by: $\dot{\mathbf{x}} = \mathcal{J}(\mathbf{y}) \cdot \mathbf{x} \equiv L(\mathbf{x})$.

Theorem. *Let $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a smooth map. Let \mathbf{y} be a hyperbolic stationary state of the system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, and L a linearization of the system at \mathbf{y} . Then there exists a neighborhood U of \mathbf{y} in which \mathbf{f} is topologically conjugate to L – namely, there is a homeomorphism $h : U \rightarrow \mathbb{R}^n$ such that $h \circ f \circ h^{-1} = L$.*

See Hartman (1960) for proof.

³⁹Note that this continuity result is often stated and proved separately.

References

- Bowles, S. and H. Gintis (2004), “The evolution of strong reciprocity: cooperation in heterogeneous populations.” *Theor Popul Biol.*, 65, 17–28.
- Boyd, Robert, Herbert Gintis, Samuel Bowles, and Peter J. Richerson (2003), “The evolution of altruistic punishment.” *P.N.A.S.*, 100, 3531–3535.
- Denant-Boemont, Laurent, David Masclet, and Charles Noussair (2007), “Punishment, counterpunishment and sanction enforcement in a social dilemma experiment.” *Economic Theory*, 33, 145–167.
- Fehr, Ernst, Urs Fischbacher, and Simon Gächter (2003), “Strong reciprocity, human cooperation and the enforcement of social norms.” *Human Nature*.
- Fowler, James H. (2005), “Altruistic punishment and the origin of cooperation.” *P.N.A.S.*, 102, 7047–7049.
- Fudenberg, Drew and David K. Levine (1998), *The theory of learning in games*, volume 2 of *MIT Press Series on Economic Learning and Social Evolution*. MIT Press, Cambridge, MA.
- Fudenberg, Drew and Eric Maskin (1990), “Evolution and cooperation in noisy repeated games.” *The American Economic Review*, 80, 274–279.
- Fudenberg, Drew and Jean Tirole (1991), *Game theory*. MIT Press, Cambridge, MA.
- Hartman, Philip (1960), “A lemma in the theory of structural stability of differential equations.” *Proc. Amer. Math. Soc.*, 11, 610–620.
- Hauert, Christoph, Arne Traulsen, Hannelore Brandt, Martin A. Nowak, and Karl Sigmund (2007), “Via freedom to coercion: the emergence of costly punishment.” *Science*, 316, 1905–1907.
- Hirsch, Morris W. and Stephen Smale (1974), *Differential equations, dynamical systems, and linear algebra*. Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London. Pure and Applied Mathematics, Vol. 60.
- Hofbauer, J., P. Schuster, and K. Sigmund (1979), “A note on evolutionary stable strategies and game dynamics.” *J. Theoret. Biol.*, 81, 609–612.

- Hofbauer, Josef and Karl Sigmund (1998), *Evolutionary games and population dynamics*. Cambridge University Press, Cambridge.
- Hofbauer, Josef and Jorgen W. Weibull (1996), “Evolutionary selection against dominated strategies.” *Journal of Economic Theory*, 71, 558–573.
- Kakutani, Shizuo (1941), “A generalization of Brouwer’s fixed point theorem.” *Duke Math. J.*, 8, 457–459.
- Levine, David K. and Wolfgang Pesendorfer (2007), “The evolution of cooperation through imitation.” *Games Econom. Behav.*, 58, 293–315.
- Maynard Smith, J. (1974), “The theory of games and the evolution of animal conflicts.” *J. Theoret. Biol.*, 47, 209–221.
- Maynard Smith, J. and G. M. Price (1973), “The logic of animal conflict.” *Nature*, 246, 15–18.
- Nakamaru, Mayuko and Yoh Iwasa (2005), “The evolution of altruism by costly punishment in lattice-structured populations.” *Evolutionary Ecology Research*, 853–870. Score-dependent viability versus score-dependent fertility.
- Nakamaru, Mayuko and Yoh Iwasa (2006), “The coevolution of altruism and punishment: role of the selfish punisher.” *J. Theoret. Biol.*, 240, 475–488.
- Nash, John (1951), “Non-cooperative games.” *Ann. of Math. (2)*, 54, 286–295.
- Nikiforakis, Nikos (2008), “Punishment and counter-punishment in public goods games: Can we still govern ourselves?” *Journal of Public Economics*, 92, 91–112.
- Nowak, M. A., A. Sasaki, C. Taylor, and D. Fudenberg (2004), “Emergence of cooperation and evolutionary stability in finite populations.” *Nature*, 428, 646–650.
- Nowak, Martin A. (2006), *Evolutionary dynamics*. The Belknap Press of Harvard University Press, Cambridge, MA. Exploring the equations of life.
- Ohtsuki, Hisashi (2008), “Analytic calculations of the viability model.” Unpublished.

- Osborne, M. J. and A. Rubinstein (2002), *A Course in Game Theory*. MIT Press, Cambridge, MA.
- Rice, John A. (2007), *Mathematical Statistics and Data Analysis*. Thomson, California.
- Samuelson, Larry and Jianbo Zhang (1992), “Evolutionary stability in asymmetric games.” *J. Econom. Theory*, 57, 363–391.
- Sigmund, Karl, Christoph Hauert, and Martin A. Nowak (2001), “Reward and punishment.” *P.N.A.S.*, 98, 10757–10762.
- Taylor, Peter D. and Leo B. Jonker (1978), “Evolutionarily stable strategies and game dynamics.” *Math. Biosci.*, 40, 145–156.
- Traulsen, A. and C. Hauert (2008), “Stochastic evolutionary game dynamics.” *ArXiv e-prints*.
- Traulsen, Arne, Jorge M. Pacheco, and Martin A. Nowak (2007), “Pairwise comparison and selection temperature in evolutionary game dynamics.” *J. Theoret. Biol.*, 246, 522–529.
- Traulsen, Arne, Noam Shresh, and Martin A. Nowak (2008), “Analytical results for individual and group selection of any intensity.” *Bull. Math. Biol.*, 70, 1410–1424.
- Trivers, Robert L. (1971), “The evolution of reciprocal altruism.” *The Quarterly Review of Biology*, 46, 35–57.
- Van Damme, Eric (1987), *Stability and perfection of Nash equilibria*. Springer-Verlag, Berlin.
- Vickers, G. T. and C. Cannings (1987), “On the definition of an evolutionarily stable strategy.” *J. Theoret. Biol.*, 129, 349–353.
- Weibull, Jörgen W. (1995), *Evolutionary game theory*. MIT Press, Cambridge, MA. With a foreword by Ken Binmore.
- Zeeman, E. C. (1980), *Population dynamics from game theory*, 471–497. Springer Berlin, Heidelberg.