

10-15-2018

Coordination and Evolutionary Dynamics: When are Evolutionary Models Reliable?

Daniel Graydon Stephenson
Chapman University, dstephenson@chapman.edu

Follow this and additional works at: https://digitalcommons.chapman.edu/esi_pubs



Part of the [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Stephenson, D. G. (2018). Coordination and evolutionary dynamics: When are evolutionary models reliable? *Games and Economic Behavior*. doi: 10.1016/j.geb.2018.10.002

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Publications by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Coordination and Evolutionary Dynamics: When are Evolutionary Models Reliable?

Comments

This article was originally published in *Games and Economic Behavior* in 2018. DOI: [10.1016/j.geb.2018.10.002](https://doi.org/10.1016/j.geb.2018.10.002)

Creative Commons License



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Copyright

The author



ELSEVIER

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb


Coordination and evolutionary dynamics: When are evolutionary models reliable? ☆

Daniel Stephenson

ARTICLE INFO

Article history:

Received 1 September 2017

Available online xxxx

JEL classification:

C62

C73

C91

C92

Keywords:

Evolutionary game theory

Sign preservation

Experiment

Coordination

Stability

Attacker defender

ABSTRACT

This study reports a continuous-time experimental test of evolutionary models in coordinated attacker–defender games. It implements three experimental treatment conditions: one with strong coordination incentives, one with weak coordination incentives, and one with zero coordination incentives. Each treatment exhibits identical equilibrium predictions but distinct evolutionary predictions. Observed behavior was tightly clustered around equilibrium under both the zero coordination treatment and the weak coordination treatment but widely dispersed from equilibrium under the strong coordination treatment. This result was anticipated by explicitly dynamic models but not by conventional stability criteria. In contrast to the widely maintained assumption of sign-preservation, subjects frequently switched to lower earning strategies, suggesting that non-sign-preserving evolutionary models may provide a more accurate characterization of human behavior.

© 2018 The Author. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Classical game theorists often describe Nash equilibrium as the outcome of introspective reasoning prior to play.¹ In contrast, evolutionary game theory explains Nash equilibrium as the long run outcome of an adaptive process.² This paper tests the predictive validity of evolutionary stability criteria, evolutionary dynamics, and the key assumption of sign-preservation that underlies many evolutionary models. A clean separation between these theoretical models is obtained in coordinated attacker–defender games that yield distinct predictions from evolutionary models but identical predictions from equilibrium models.

In coordinated attacker defender games, each agent faces conflict with members of an opposing population and opportunities for coordination with members of their own population. The experiment implements three treatment conditions:

☆ The author thanks the Texas A&M Humanities and Social Science Enhancement of Research Capacity Program and the College of Liberal Arts at Texas A&M and the Chapman University Economic Science Institute for their generous financial support of this research. This research was supported by NSF grant SES-1458541. The author is grateful for the insightful comments from three anonymous reviewers, attendees of the 2015 North American Meeting of the Economic Science Association, and attendees of the 2016 Texas Experimental Economics Symposium. The author would also like to thank Alexander L. Brown, Catherine C. Eckel, Daniel E. Fragiadakis, Ryan Oprea, and the experimental research group at Texas A&M University for their continual support and feedback.

E-mail address: dstephenson@chapman.edu.

URL: <http://www.danielgstephenson.com>.

¹ A thorough discussion of the epistemic conditions for Nash equilibrium in classical game theory can be found in Aumann and Brandenburger (1995).

² This interpretation of Nash equilibrium as the long run outcome of an adaptive adjustment process is not a recent innovation. Notably, it was employed by Cournot (1838) and Nash (1950).

<https://doi.org/10.1016/j.geb.2018.10.002>

0899-8256/© 2018 The Author. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

one with strong coordination incentives, one with weak coordination incentives, and one with zero coordination incentives. Consistent with the continuous-time structure of dynamic evolutionary models, subjects could adjust their strategy continuously and earned continuous flow payoffs throughout each experimental session. These continuous-time experimental protocols allow for the observation of dynamic behavioral phenomena that can be difficult to identify under conventional discrete-time experimental protocols.³

Many theoretical models make identical predictions for all three treatments, including Nash equilibrium, logit quantal response equilibrium (McKelvey and Palfrey, 1995), level- k (Stahl and Wilson, 1994), and the time average of the Shapley polygon (Benaïm et al., 2009). In contrast, evolutionary models make distinct predictions for each treatment. Evolutionary stability criteria (Taylor, 1979) and positive definiteness conditions (Hopkins and Seymour, 2002) predict weak stability under zero coordination incentives but strong instability under positive coordination incentives.

This experiment tests three widely employed evolutionary dynamics: the best response dynamic (Gilboa and Matsui, 1991), the Smith dynamic (Smith, 1984), and the logit dynamic (Fudenberg and Levine, 1998). The widely employed assumption of sign-preservation states that agents switch exclusively from lower performing strategies to higher performing strategies. The best response dynamic and the Smith dynamic satisfy sign-preservation. They predict convergence to equilibrium under zero coordination incentives but convergence to a stable limit cycle under nonzero coordination incentives. The logit dynamic does not impose sign preservation. It predicts convergence to equilibrium under sufficiently weak coordination incentives but convergence to a stable limit cycle under sufficiently strong coordination incentives. All three evolutionary dynamics predict that limit cycles will exhibit greater deviation from equilibrium under stronger coordination incentives.

In line with these predictions, subjects exhibited significant cyclical behavior under all three experimental treatment conditions. Behavior remained tightly clustered around equilibrium in the zero coordination treatment and the weak coordination treatment but exhibited significantly greater deviation from equilibrium in the strong coordination treatment. This result is consistent with dynamic evolutionary models, but not with equilibrium models or evolutionary stability criteria, since these models classify the strong coordination treatment and the weak coordination treatment identically. In contrast to the assumption of sign-preservation, subjects frequently switched from higher earning strategies to lower earning strategies, suggesting that non-sign-preserving evolutionary dynamics may provide a more accurate characterization of human behavior.

The remainder of this paper is organized as follows: Section 2 discusses the related literature. Section 3 describes the theoretical framework and Section 4 describes the experimental design. Section 5 covers the hypotheses, Section 6 presents the main results, and Section 7 concludes.

2. Related literature

This paper contributes to a growing body of experimental research testing evolutionary game theory in laboratory experiments. Early experimental investigation of evolutionary game theory focused on testing evolutionary models of equilibrium selection. Van Huyck et al. (1990) observe strong convergence to inefficient pareto dominated equilibria in minimum effort games with multiple equilibria. The authors suggested that the emergence of these particular equilibria may have been driven by the presence of strategic uncertainty rather than previously traditional equilibrium refinement methods. Crawford (1991) later formalized this conjecture via evolutionary stability criteria.

Oprea et al. (2011) employ continuous-time laboratory procedures in testing evolutionary models in Hawk–Dove population games with two asymmetric Nash equilibria and one symmetric Nash equilibrium. In accordance with evolutionary stability criteria, they found that subjects approximate one of the asymmetric equilibria under two population matching protocols, but approximate the symmetric equilibrium under the one population matching protocols. Their experiment provides empirical evidence for the ability of evolutionary models to help solve the equilibrium selection problem in games with multiple Nash equilibria.

Cason et al. (2013) implemented laboratory procedures where subjects select mixed strategies and continuously adjust strategies over time. They observed cyclical behavior in three rock-paper-scissors population games: one with a stable equilibrium and two with unstable equilibria. In these games, they tested predictions from the time average of the Shapley polygon (TASP) (Benaïm et al., 2009) against predictions from Nash equilibrium. The TASP describes the time-average of long run cyclical behavior under the best response dynamic. Their experiment provides empirical evidence for the ability of the TASP to characterize deviations from Nash equilibrium.

Lim and Neary (2016) conducted discrete-time experiments investigating stochastic adjustment dynamics in multi-population coordination games with tension in preferences between populations. They find that subject frequently select myopic best responses to population behavior in the previous period. They find that the probability of deviations from this myopic best response was decreasing in the payoff from the best-response but independent of the payoff to other actions.

Benndorf et al. (2016) conducted continuous-time experiments investigating hawk–dove games that feature a mixture of both inter-population and intra-population interactions. If payoffs are primarily determined by inter-population interactions the evolutionary replicator dynamics predict convergence to the asymmetric pure strategy Nash equilibrium. Conversely, if payoffs are primarily determined by intra-population interactions, then behavior is predicted to converge to the symmetric

³ See Friedman et al. (2015) and Oprea et al. (2011) for examples of such phenomena.

mixed strategy Nash equilibrium. The authors also identify an intermediate parameter region under which behavior is predicted to exhibit a hybrid equilibrium where one population utilizes a pure strategy and the other a mixture. Their experiment tests the equilibrium selection predictions in each of these three distinct parameter regions.

Benndorf and Martinez-Martinez (2017) examine systematic behavioral deviations from predictions of the replicator dynamic in mixed population hawk–dove games. They observe a bias against the hawk strategy in the symmetric mixed configuration and less polarization than predicted in the asymmetric configuration, both of which are consistent with theoretical predictions from logit dynamics but inconsistent with theoretical predictions from replicator dynamics. They also observe significant individual heterogeneity in the degree of behavioral noise across sessions.

In contrast to previous literature, this paper tests the theoretical predictions of evolutionary stability criteria against those of explicitly dynamic evolutionary models. It finds that subjects in coordinated attacker games exhibit significant cyclical behavior with significantly greater deviation from equilibrium under strong coordination incentives than weak coordination incentives. This result is consistent with dynamic evolutionary models, but not with equilibrium models or stability criteria.

3. Theory

This study tests evolutionary models in attacker–defender population games. Population games provide a theoretical framework for the analysis of strategic interaction between large numbers of agents. Subsection 3.1 introduces population games. Subsection 3.2 describes the class of coordinated attacker–defender games. Subsection 3.3 discusses stability criteria and Subsection 3.4 discusses evolutionary dynamics.

3.1. Population games

A population game is played by a *society* composed of one or more *populations* $p \in P = \{1, \dots, p\}$. Each population p contains agents who choose pure strategies from the set $S^p = \{1, \dots, n^p\}$. The proportion of population p that employs the pure strategy i is denoted by $x_i^p \in [0, 1]$. Accordingly, a *population state* $x^p = (x_1^p, \dots, x_{n^p}^p)$ indicates the proportion of population p that employs each pure strategy $i \in S^p$. Let $m = \sum_{i=1}^p n_p$, so a *social state* $x = (x^1, \dots, x^p) \in \mathbb{R}^m$ describes the state of all populations $p \in P$. The payoff to an agent in population p who employs pure strategy $i \in S^p$ is given by the payoff function $\pi_i^p(x)$. A social state x is said to be a *Nash equilibrium* if no agent in any population can increase her payoff by unilaterally adjusting her strategy. More formally, a social state x is a *Nash equilibrium* if, for every population $p \in P$ and every pair of pure strategies $i, j \in S^p$ such that $x_i^p > 0$, the payoff $\pi_i^p(x)$ earned by agents in population p who employ the pure strategy i is greater than or equal to the payoff $\pi_j^p(x)$ earned by agents in population p who employ the pure strategy j .

3.2. Coordinated attacker–defender games

This paper considers coordinated attacker–defender games played by a large population of attackers A and a large population of defenders D . Each defender chooses one of two symmetrically valued targets $S^D = \{1, 2\}$ to defend. Each attacker simultaneously chooses one of the two targets $S^A = \{1, 2\}$ to attack. Each defender is randomly matched against an opposing attacker. The attacker prefers to attack the undefended target, but the defender prefers to defend the target that is attacked. The payoff to an attacker from this conflict is proportional to the fraction of the defender population that does not defend the type of target she attacks. Conversely, the payoff to a defender from this conflict is proportional to the fraction of the attacker population that attacks the same type of target she chose to defend.

Each agent is also randomly matched with a partner from the same population. Agents prefer to coordinate their target with their partner. These intrapopulation coordination incentives represent positive externalities within the population. The payoff to a defender from these positive externalities is proportional to the fraction of other defenders who defend the same type of target she chooses to defend. Conversely, the payoff to an attacker from these positive externalities is proportional to the fraction of other attackers who attack the same type of target she attacks. Attackers prefer to coordinate their attacks on the same type of target as other attackers and defenders prefer to coordinate their defenses on the same type of target as other defenders. Formally, the payoff functions for the coordinated attacker defender game are given by

$$\begin{aligned} \pi_1^A(x) &= Mx_2^D + Cx_1^A & \pi_2^A(x) &= Mx_1^D + Cx_2^A \\ \pi_1^D(x) &= Mx_1^A + Cx_1^D & \pi_2^D(x) &= Mx_2^A + Cx_2^D \end{aligned} \quad (1)$$

where $M > C \geq 0$. The parameter M denotes the strength of the interpopulation conflict incentives and the parameter C denotes the strength of the intrapopulation coordination incentives. The case with $C = 0$ is equivalent to matching pennies or the evolutionary battle of the sexes (Hofbauer and Sigmund, 1998). To the best of the author's knowledge, this is the first paper to consider the case with $C > 0$. Since the interpopulation conflict incentives (M) are strictly stronger than the intrapopulation coordination incentives (C), Proposition 1 states that there is a unique Nash equilibrium social state under which both populations are evenly divided between their two pure strategies. Accordingly, any finite repetition of the coordinated attacker defender game has a unique subgame perfect Nash equilibrium, since the stage game Nash equilibrium is unique.

Proposition 1. *The coordinated attacker defender game has a unique Nash equilibrium social state under which $x_1^A = x_2^A = x_1^D = x_2^D = \frac{1}{2}$.*

Proof. See appendix. \square

3.3. Evolutionary stability

A Nash equilibrium is said to be *evolutionarily stable* if any sufficiently small proportion of agents who deviate from equilibrium earn less than agents who maintain equilibrium behavior. More formally, a Nash equilibrium x is said to be *evolutionarily stable* if there exists some $\varepsilon > 0$ such that for any other social state $y \neq x$ and any proportion $\alpha \in (0, \varepsilon)$ we have $x \cdot \pi(z) > y \cdot \pi(z)$ where $z = \alpha y + (1 - \alpha)x$. The less restrictive condition of *neutral stability* is obtained by exchanging this strict inequality for the weak inequality $x \cdot \pi(z) \geq y \cdot \pi(z)$. These evolutionary stability criteria were originally developed for games played by a monomorphic population of agents by Smith and Price (1973) and were extended to the case of multiple polymorphic populations by Taylor (1979).

If a Nash equilibrium social state is not evolutionarily stable, then it is said to be evolutionarily unstable. If a Nash equilibrium is not even neutrally stable, then it is said to be strictly unstable. If an equilibrium is strictly unstable then an arbitrarily small proportion of agents deviating from equilibrium can earn more than agents who maintain equilibrium behavior. More formally, a Nash equilibrium social state x is said to be *strictly unstable* if there exists some other social state $y \neq x$ such that for all $\varepsilon > 0$ there exists some proportion $\alpha \in (0, \varepsilon)$ such that $x \cdot \pi(z) < y \cdot \pi(z)$ where $z = \alpha y + (1 - \alpha)x$.

A fully mixed Nash equilibrium is said to be *positive definite* if any sufficiently small deviation from equilibrium creates incentives that push behavior farther away from equilibrium. More formally, a fully mixed Nash equilibrium x is said to be positive definite if the linear approximation of the payoff function $\pi'(x)$ at the equilibrium point x is positive definite on $Z = \{z \in \mathbb{R}^m : \sum_{i=1}^{n_p} z_i^p = 0\}$. Conversely, a fully mixed Nash equilibrium is said to be *negative definite* if any sufficiently small deviation from equilibrium creates incentives that push behavior back towards equilibrium. More formally, a fully mixed Nash equilibrium x is said to be negative definite if the linear approximation of the payoff function $\pi'(x)$ at the equilibrium point x is negative definite on Z . A Nash equilibrium that is neither positive definite nor negative definite is said to be *indefinite*. These dynamic instability criteria were originally described by Hopkins (1999) and were extended to more general payoff functions by Hopkins and Seymour (2002). Hopkins and Seymour (2002) show that these criteria characterize the stability properties of positive definite adaptive dynamics, while Hopkins (1999) shows that perturbed best response dynamics share stability properties with positive definite adaptive dynamics when the perturbation is sufficiently small.

Proposition 2. *Under zero coordination incentives, the unique Nash equilibrium of the coordinated attacker defender game is neutrally stable and indefinite. Under positive coordination incentives, it is strictly unstable and positive definite.*

Proof. See appendix. \square

3.4. Evolutionary dynamics

Dynamic evolutionary models provide a system of differential equations that describe an adaptive process from which Nash equilibrium can emerge as a stationary point. A *revision protocol* $\rho_{ij}^p(x)$ gives the rate at which agents in population p switch from strategy i to strategy j as a function of the social state x . This method of describing evolutionary dynamics in population games is widely employed by evolutionary game theorists including Björnerstedt and Weibull (1994), Sandholm (2009), and Bulò and Bomze (2011). The *evolutionary dynamics* induced by a revision protocol ρ are given by the system of differential equations:

$$\dot{x}_i^p = \sum_{j \in S^p} x_j^p \rho_{ji}^p(x) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(x) \quad (2)$$

The first summation in this expression describes the inflow of agents into strategy i from other strategies, while the second summation describes the outflow of agents from strategy i onto other strategies. The difference between these two summations yields the net rate of change \dot{x}_i^p in the proportion of population p employing the pure strategy i . An evolutionary dynamic is said to be *sign-preserving* if agents exclusively switch from lower performing strategies to higher performing strategies. More formally, an evolutionary dynamic is said to be *sign-preserving* if, for every social state x and every population p , the conditional switch rate $\rho_{ij}^p(x)$ is positive if and only if $\pi_j^p(x)$ is strictly greater than $\pi_i^p(x)$.

Evolutionary dynamics explicitly describe the adaptive process of behavioral adjustment from which equilibrium can emerge. While classical game theorists often describe Nash equilibrium as the outcome of introspective reasoning prior to play, dynamic evolutionary models explain Nash equilibrium as the long run outcome of an adaptive process. Table 1 provides the revision protocols for three widely employed evolutionary dynamics: the best response dynamic, the Smith dynamic, and the logit dynamic. Fig. 1 illustrates these revision protocols for populations with two pure strategies. The

Table 1
Revision protocols for three evolutionary dynamics.

Name of dynamic	Revision protocol
Best response	$\rho_{ij}^p(x) = 1 \left(j \in \operatorname{argmax}_{k \in SP} \pi_k^p(x) \right)$
Smith	$\rho_{ij}^p(x) = \max \left\{ 0, \pi_j^p(x) - \pi_i^p(x) \right\}$
Logit	$\rho_{ij}^p(x) = \frac{\exp \left(\eta^{-1} \pi_j^p(x) \right)}{\sum_{k \in SP} \exp \left(\eta^{-1} \pi_k^p(x) \right)}$




Fig. 1. Conditional switch rates in populations with two pure strategies.

remainder of this section describes the underlying behavioral assumptions of these dynamics and highlights the key differences between their theoretical predictions.

In their analysis of Nash equilibrium social stability, Gilboa and Matsui (1991) consider an evolutionary dynamic under which agents myopically switch to their best response under the current social state. Since agents switch to inferior strategies, the best response dynamic is sign-preserving. Since agents switch exclusively to best responses, the fixed points of the best response dynamic always coincides with the set of Nash equilibria.

In his analysis of highway traffic congestion dynamics, Smith (1984) considered an evolutionary dynamic under which agents switch from lower performing strategies to higher performing strategies at a rate proportional to the difference in payoffs. The Smith dynamic is a pairwise comparison dynamic because the conditional switch rate from strategy i to strategy j in population p depends exclusively on the pairwise comparison $\pi_j^p(x) - \pi_i^p(x)$. Sandholm (2009) proves that the fixed points of sign-preserving pairwise comparison dynamics, such as the Smith dynamic, exactly coincide with the set Nash equilibria.

Fudenberg and Levine (1998) describe an evolutionary dynamic under which agents switch towards their perceived best response under a noisy perturbation of their payoffs. This logit dynamic is closely related to the logit quantal response equilibrium described by McKelvey and Palfrey (1995). Under the logit dynamic, agents are more likely to switch to strategies that yield higher payoffs. Here η denotes the noise level in an agent's perception of payoffs. As η becomes large, agents become increasingly insensitive to payoff differences and the conditional switch rate approaches uniformly random behavior. Conversely, as η becomes small, agents perceive their payoffs more precisely, and the conditional switch rate approaches that of the best response dynamic. Unlike the Smith dynamic and the best response dynamic, the logit dynamic is not sign-preserving, so its fixed points frequently fail to coincide with Nash equilibria.

Alos-Ferrer and Netzer (2010) characterize the stochastically stable states of the logit dynamic. They show that the logit dynamic converges to Nash equilibria in best-response potential games, but fails to select potential maximizers outside this class. In general, the fixed points of the logit dynamic correspond to the set of logit quantal response equilibria (QRE). In coordinated attacker defender games, the unique logit QRE is identical to the unique Nash equilibrium since equilibrium mixing probabilities are equal across strategies. This feature of coordinated attacker defender games provides a clean separation between the logit dynamic and the logit QRE. Different coordination incentives induce different logit dynamics but identical logit QRE.

3.5. Coordinated attacker–defender dynamics

Fig. 2 depicts evolutionary dynamics in coordinated attacker defender games. The horizontal axis of each graph indicates the proportion of attackers attacking target 2 and the vertical axis of each graph indicates the proportion of defenders defending target 2. Since defenders have an incentive to defend the target that is most likely to be attacked and attackers have an incentive to attack the target that is least likely to be defended, evolutionary dynamics consistently predict cyclic behavior in the coordinated attacker defender game. Under zero coordination incentives, the social state gradually spirals inwards towards equilibrium. Under sufficiently strong coordination incentives, small perturbations from equilibrium cause the social state to diverge from equilibrium to a stable limit cycle.

Fig. 3 depicts stable limit cycles in coordinated attacker defender games with positive coordination incentives. All three dynamics predict greater deviation from equilibrium under stronger coordination incentives. In contrast, conventional evolutionary stability criteria classify all coordinated attacker defender games with positive coordination incentives identically.

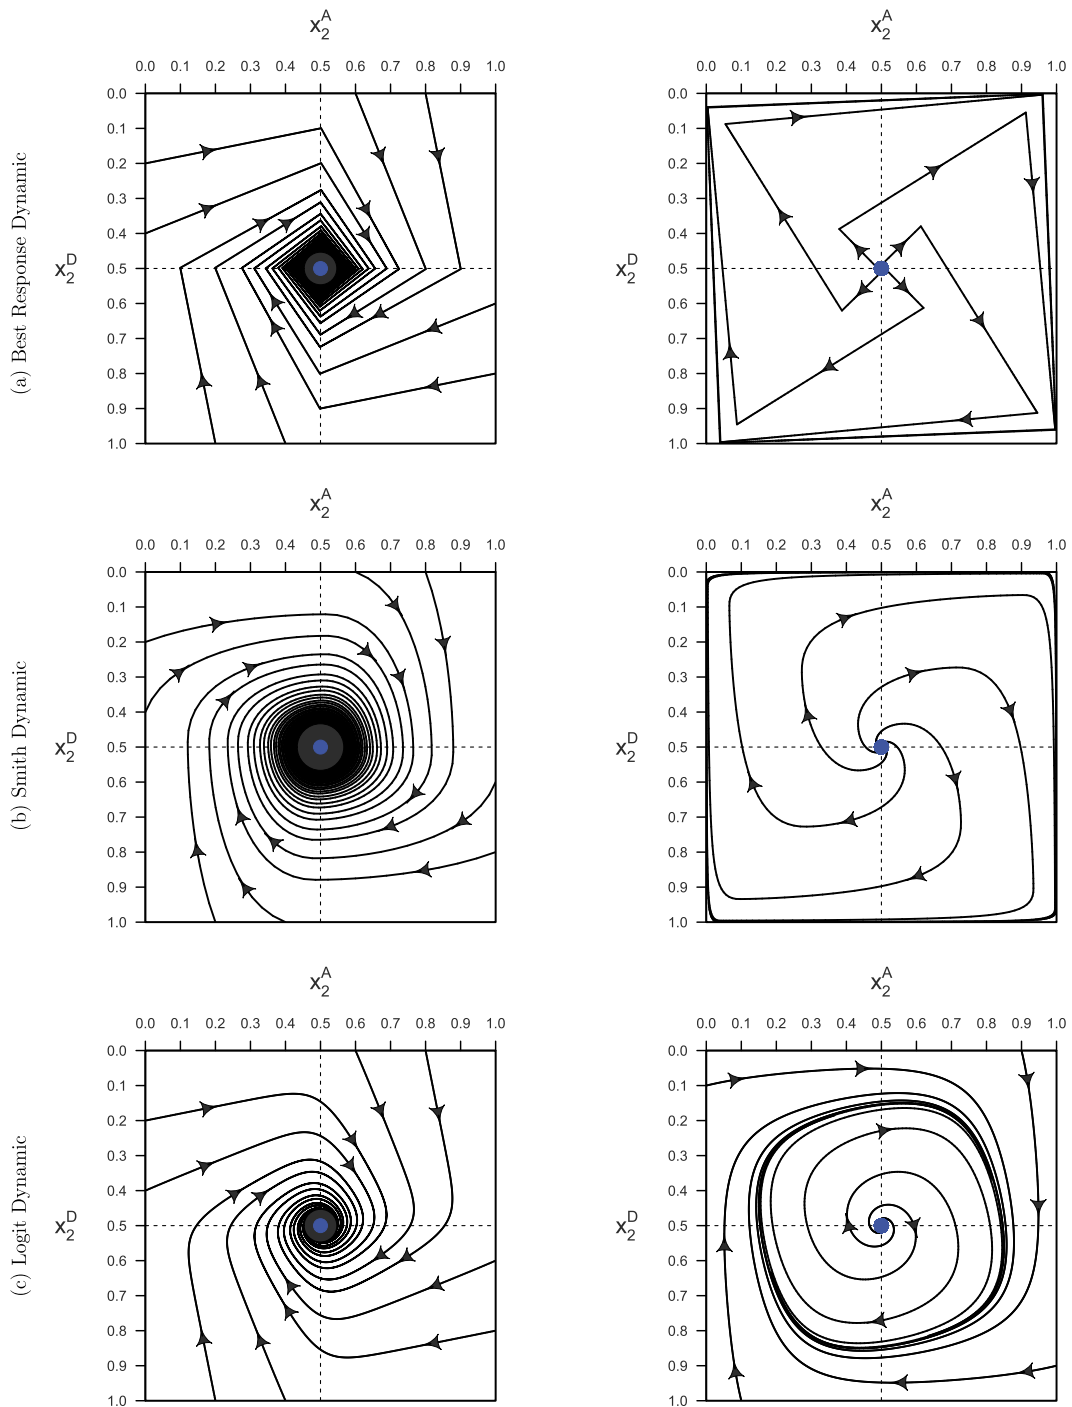


Fig. 2. Evolutionary dynamics in the coordinated attacker defender game. Theoretical predictions for the zero coordination case where $M = 5$ and $C = 0$ are shown on the left. Theoretical predictions for the nonzero coordination case where $M = 2.6$ and $C = 2.4$ are shown on the right. These parameters are selected to reflect those implemented in the experiment. The first row depicts the best response dynamics, the second row depicts the Smith dynamics, and the third row depicts the logit dynamics with $\eta = 0.6$.

Fig. 4 depicts the amplitude of the logit limit cycle as a function of the noise parameter η . As η becomes large, agents become increasingly insensitive to payoff differences, so agents switch to each strategy with closer to equal probability, pushing the social state towards the center of the state space. Accordingly, Proposition 3 states that the Nash equilibrium is globally asymptotically stable under the logit dynamic if the noise parameter η is sufficiently large. The critical thresh-

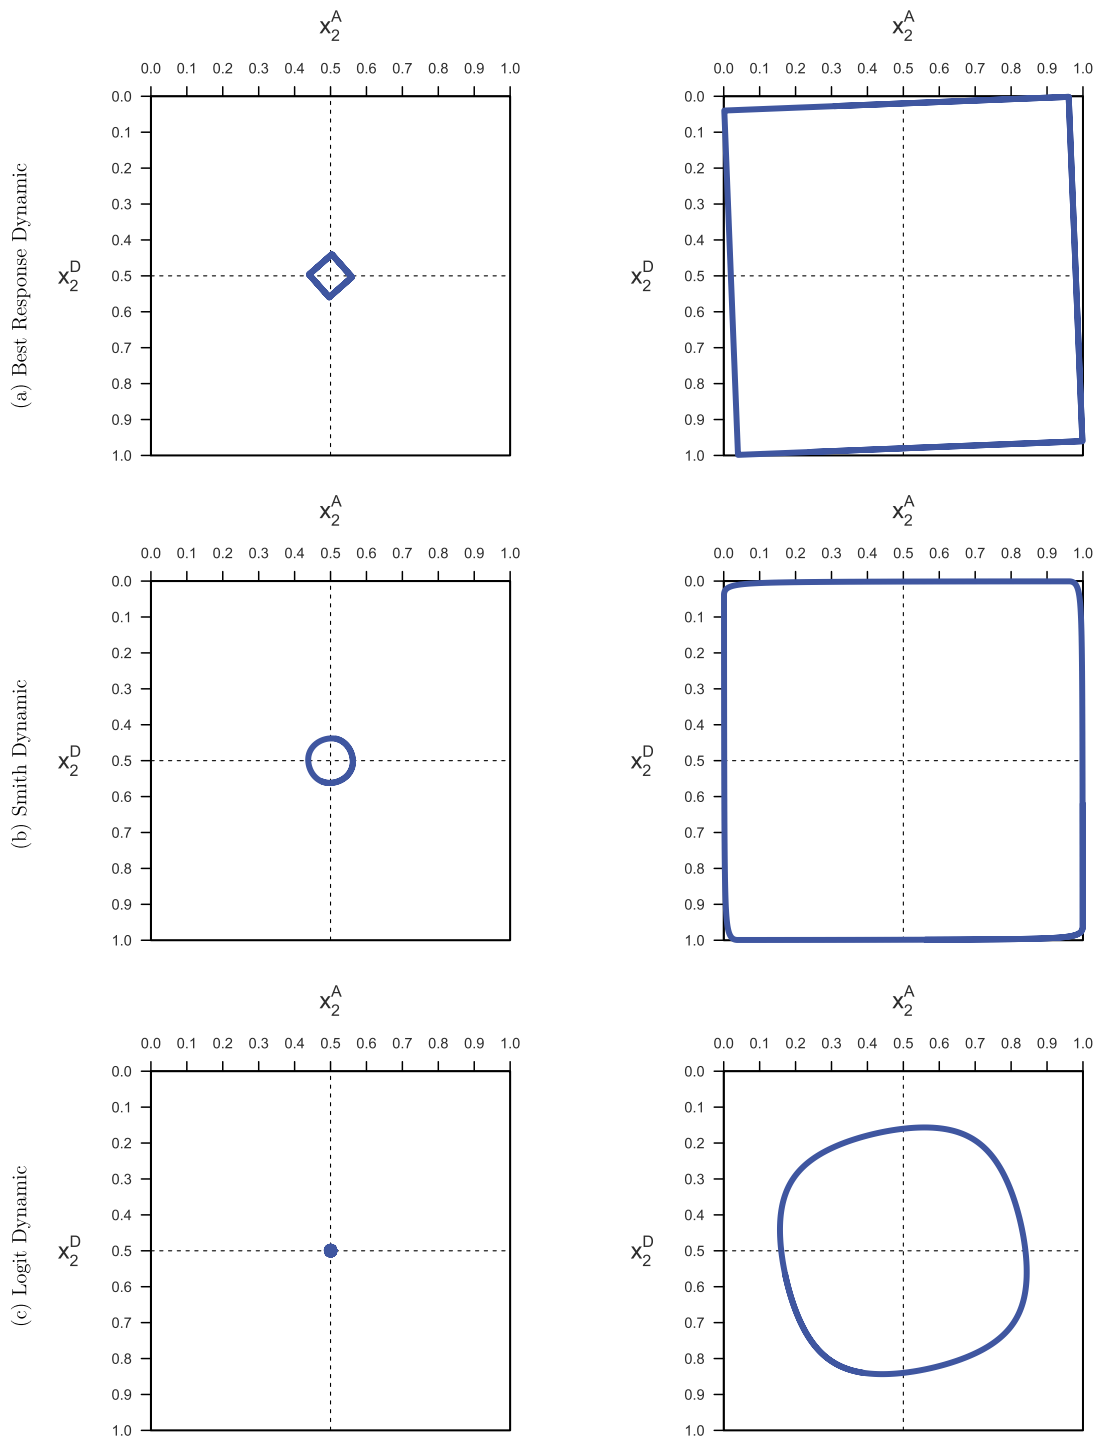


Fig. 3. Stable limit cycles in the coordinated attacker defender game. Theoretical predictions for the weak coordination case where $M = 4.8$ and $C = 0.2$ are shown on the left. Theoretical predictions for the strong coordination case where $M = 2.6$ and $C = 2.4$ are shown on the right. These parameters are selected to reflect those implemented in the experiment. The first row depicts the best response dynamics, the second row depicts the Smith dynamics, and the third row depicts the logit dynamics with $\eta = 0.6$.




Fig. 4. Mean distance from equilibrium in the stable limit cycle of the logit dynamic as a function of the noise parameter η . The blue line illustrates the predicted distance from equilibrium in the strong coordination case where $M = 2.6$ and $C = 2.4$. The green line illustrates the predicted distance from equilibrium in the weak coordination case where $M = 4.8$ and $C = 0.2$. These parameters are selected to reflect those implemented in the experiment.

Table 2
Parameter values for each of the three treatment conditions.

Treatment	M	C
No-coordination	\$5.00	\$0.00
Weak coordination	\$4.80	\$0.20
Strong coordination	\$2.60	\$2.40

old is increasing in the coordination parameter C and is equal to zero when $C = 0$, so the logit dynamic exhibits global convergence to equilibrium if coordination incentives are sufficiently weak.

Proposition 3. For every coordinated attacker–defender game there is $\eta^*(M, C) > 0$ such that the Nash equilibrium is globally asymptotically stable under the logit dynamic if $\eta > \eta^*(M, C)$. Further, $\eta^*(M, C)$ is strictly increasing in C with $\eta^*(M, 0) = 0$.

Proof. See appendix. □

4. Experimental design and procedures

The experimental design has three treatment conditions, each of which implements the coordinated attacker defender game with a different level of coordination incentives. The first treatment has zero coordination incentives, the second treatment has weak coordination incentives, and the third treatment has strong coordination incentives. All three treatments exhibit identical Nash equilibrium predictions. Evolutionary stability criteria classify the zero coordination treatment as neutrally stable and the other two as strictly unstable. Evolutionary dynamics predict convergence to Nash equilibrium under sufficiently weak coordination incentives and convergence to stable limit cycles under sufficiently strong coordination incentives. Evolutionary dynamics consistently predict that limit cycles will exhibit larger deviations from equilibrium under stronger coordination incentives. Earnings per minute in each treatment were given by

$$\begin{aligned}
 \pi_1^A(x) &= Mx_2^D + Cx_1^A & \pi_2^A(x) &= Mx_1^D + Cx_2^A \\
 \pi_1^D(x) &= Mx_1^A + Cx_1^D & \pi_2^D(x) &= Mx_2^A + Cx_2^D
 \end{aligned}
 \tag{3}$$

where π_i^g denotes the payoff to a subject in group g who employs the pure strategy i and x_j^p denotes the proportion of other subjects in group p employing the pure strategy j . Table 2 provides the values of the coordination parameter C and the conflict parameter M under each of the three treatment conditions. These treatment conditions implement the parameter values illustrated by Fig. 2 and Fig. 3 in the previous section. In equilibrium, these parameter values equalize the equilibrium earnings rate at \$2.50 per minute across all three treatment conditions.

Each experimental session was conducted with twenty subjects and lasted for approximately thirty minutes. On average, each subject earned a total of \$18.63, including a five dollar show-up payment. We employ a between-subjects design, so each subject participated in only one experimental session and each session implemented only one of the three experimental treatments. Two of the six sessions implemented each of the three experimental treatments, for a total of 120




Fig. 5. Experimental decision interface. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

distinct experimental subjects. At the beginning of each session, subjects were randomly divided into two equally sized population groups. Each group of attackers consisted of exactly 10 subjects and each group of defenders consisted of exactly 10 subjects. Subjects stayed in the same population group for the duration of the session. Each session consisted of eight identical periods during which subjects played a coordinated attacker–defender game. Each period lasted for exactly forty seconds.

A screenshot of the full experimental display is provided by Fig. 5. Subjects could continuously adjust their mixed strategy throughout each period. They were not restricted to exclusively selecting pure strategies. The horizontal position of the green bar represented a subject's probability of selecting the pure strategy *B*. When the green bar was all the way on the right side of the graph, the subject employed the pure strategy *B* with certainty. Conversely, when the green bar was all the way on the left side of the graph, the subject employed the pure strategy *A* with certainty. When the green bar was in the interior of the graph, the subject employed a non-trivial mixed strategy. Throughout each period, subjects earned continuous flow payoffs at a rate proportional to the expected payoff from their mixed strategy.

The diagonal black line illustrates the current earnings rate for each strategy. The height of the line at the left side of the graph indicates the current earnings rate for the pure strategy *A*. Similarly, the height of the line at the right side of the graph shows the current earnings rate for the pure strategy *B*. The height of the green bar indicates the subject's current earnings rate. The subject's current earnings rate and the subject's accumulated earnings were also listed at the bottom of the decision interface. Subjects could directly observe their current strategy and their current payoff but not the payoffs or strategies of others. At the end of each session, subjects received their total accumulated earnings plus a five dollar show-up payment.

5. Hypotheses

Proposition 1 states that all three experimental treatment conditions have an identical unique Nash equilibrium. This feature of the coordinated attacker–defender game is part of what allows the experimental design to disentangle theoretical predictions of evolutionary models from those of other solution concepts. In contrast, evolutionary stability criteria classify the zero coordination treatment differently from the other two treatments. Proposition 2 states that the unique Nash equilibrium is neutrally stable and indefinite under zero coordination incentives, but strictly unstable and positive definite under positive coordination incentives. These stability criteria motivate the first hypothesis.

Hypothesis 1. Subjects in the zero coordination treatment will exhibit smaller deviations from equilibrium than subjects in either of the other two treatments.

Dynamic evolutionary models consistently predict that the social state will exhibit convergence to equilibrium under zero coordination incentives, gradually spiraling inwards towards the Nash equilibrium as shown in Fig. 2. Under sufficiently strong coordination incentives, evolutionary dynamics predict convergence to a stable limit cycle. As detailed in Subsection 3.5, different evolutionary dynamics yield different predictions regarding the convergence process. In particular, the widely employed class of sign-preserving dynamics predict that subjects will exclusively switch from lower earning strategies to higher earning strategies. These models motivate the second hypothesis.

Hypothesis 2. Subjects will exclusively switch from lower earning strategies to higher earning strategies.




Fig. 6. The observed path of the social state over each period of each session. The left column depicts the first half of each period. The right column depicts the second half of each period. Two of the six sessions implemented each of the three experimental treatment conditions. One is shown in green and the other is shown in blue. Grey lines depict the stable limit cycles predicted by the logit dynamics with $\eta = 0.6$.

Evolutionary stability criteria classify all coordinated attacker defender games with positive coordination incentives identically. In contrast, dynamic evolutionary models consistently predict that stable limit cycles will exhibit greater deviation from equilibrium under stronger coordination incentives. This theoretical prediction motivates the third hypothesis.

Hypothesis 3. Subjects will exhibit greater deviation from equilibrium in the strong coordination treatment than the weak coordination treatment.

Table 3

Deviation from equilibrium as measured by the time average of the Euclidean distance between the Nash equilibrium social state and the observed social state over the entirety of each period.

Treatment	Dev. from Eq.	Null hypothesis	p-value
Zero coordination	0.067 (0.0030)	Zero = weak	0.539
Weak coordination	0.071 (0.0033)	Zero = strong	<0.001
Strong coordination	0.226 (0.0065)	Weak = strong	<0.001

(a) Mean deviation from equilibrium under all three treatments. Standard errors are given in parentheses.

(b) Nonparametric Mann–Whitney tests for treatment effects on deviation from equilibrium.

Coefficient	Estimate	(s.e.)	p-value
Constant	0.067	(0.005)	<0.001
Weak coordination	0.003	(0.007)	0.644
Strong coordination	0.159	(0.007)	<0.001

(c) A random effects regression where the dependent variable is the mean deviation from equilibrium. The default treatment is the zero coordination treatment.

6. Results

Fig. 6 illustrates the observed path of the social state over every period of every experimental treatment. The path of the social state over each period of each session is depicted by a distinct line. The first row depicts the observed behavior under the no coordination treatment, the second row depicts the observed behavior under the weak coordination treatment, and the third row depicts the observed behavior under the strong coordination treatment. The left column depicts observed behavior during the first half of each period. The right column depicts observed behavior under second half of each period. The horizontal axis of each graph illustrates the proportion of attackers that chose to attack target 2 and the vertical axis in each graph illustrates the proportion of defenders that chose to defend target 2.

6.1. Deviation from equilibrium

The strong coordination treatment exhibited considerably greater deviation from equilibrium than either the weak coordination treatment or the no coordination treatment. Table 3 provides the average deviation from equilibrium across each of the three treatment conditions and presents hypothesis tests for treatment effects. Deviation from equilibrium is measured as the time-average of the Euclidean distance between the Nash equilibrium social state and the observed social state over the entirety of each period. The unit of observation is one period. Both nonparametric Mann–Whitney tests and a parametric random-effects regression find that strong coordination incentives significantly increase deviation from equilibrium at the one percent level, while weak coordination incentives have no significant effect.

Result 1. Subjects exhibited significantly greater deviation from equilibrium under the strong coordination treatment than under the weak coordination treatment or the no coordination treatment.

Result 1 is consistent with dynamic evolutionary models but inconsistent with Nash equilibrium, logit QRE, and conventional evolutionary stability criteria. Proposition 1 states that the unique Nash equilibrium is identical under all three experimental treatments. Since equilibrium mixing probabilities are equal across strategies, the logit QRE coincides with the Nash equilibrium in all three treatments. Proposition 2 states that the weak coordination treatment and the strong coordination treatment are classified identically by evolutionary stability criteria. In line with the experimental data, Fig. 3 shows that dynamic evolutionary models predict larger deviations from equilibrium under the strong coordination treatment than under the weak coordination treatment, so Result 1 provides empirical evidence for the ability of explicitly dynamic evolutionary models to predict relative stability more accurately than conventional evolutionary stability criteria. Further, Proposition 3 indicates that the logit dynamic predicts global convergence to equilibrium when coordination incentives are sufficiently weak, consistent with the lack of a significant difference in deviation from equilibrium under the no coordination treatment and the weak coordination treatment.

6.2. Cyclical behavior

Under all three treatment conditions, subject behavior exhibited significant behavioral cycles orbiting the unique Nash equilibrium. Table 4 provides the average value of the cycle rotation index (Cason et al., 2013) under each treatment condition and presents hypothesis tests for the presence of significant cyclical behavior. The unit of observation is one period. Both nonparametric Wilcoxon sign-rank tests and a parametric random-effects regression find that subjects in all three treatments exhibited significant behavioral cycles orbiting the unique Nash equilibrium, consistent with theoretical prediction from dynamic evolutionary models.

Table 4
Cyclical behavior as measured by the cycle rotation index (CRI) under all three treatments.

Treatment	CRI	Treatment	p-value
Zero coordination	0.839 (0.043)	Zero coordination	<0.001
Weak coordination	0.857 (0.036)	Weak coordination	<0.001
Strong coordination	0.936 (0.035)	Strong coordination	<0.001

(a) Mean cycle rotation index under all three treatments. Standard errors are given in parentheses.

Treatment	p-value
Zero coordination	<0.001
Weak coordination	<0.001
Strong coordination	<0.001

(b) Nonparametric Wilcoxon signed-rank tests for significant behavioral cycling.

Coefficient	Estimate	(s.e.)	p-value
Constant	0.839	(0.052)	<0.001
Weak coordination	0.018	(0.073)	0.801
Strong coordination	0.097	(0.073)	0.181

(c) A random effects regression where the dependent variable is the cycle rotation index. The default treatment is the zero coordination treatment.




Fig. 7. The solid blue line depicts the empirical switch rate. It illustrates the percentage of subjects who chose to switch from strategy i to strategy j when the difference in payoffs was given by $\pi_j^p - \pi_i^p$. The dotted black line depicts the switch rate predicted by the logit dynamic under the estimated parameters.

Result 2. The social state exhibited significant clockwise cyclical dynamics under all three treatment conditions.

6.3. Sign preservation

Fig. 7 illustrates the observed rate at which subjects switch from one strategy to another, conditional on the difference in payoffs. Subjects switched from lower performing strategies to higher performing strategies more often than they switched from higher performing strategies to lower performing strategies. However, in contrast to Hypothesis 2, subjects also frequently switched from higher performing strategies to lower performing strategies. This result is inconsistent with the assumption of sign-preservation which is maintained by both the Smith dynamic and the best response dynamic, but is consistent with theoretical predictions from the logit dynamic.

Result 3. The empirical switch rates did not satisfy the assumption of sign-preservation as subjects frequently switched from higher performing strategies to lower performing strategies.

Result 3 suggests that the widely maintained assumption of sign preservation may not be appropriate for modeling the adaptive process giving rise to Nash equilibrium in human behavior. Accordingly, dynamic evolutionary models that do not satisfy sign preservation may provide a more reliable characterization of human behavior than sign-preserving dynamics.

Table 5

Maximum likelihood estimates for the logit dynamic and the implied deviations from equilibrium. Standard errors are obtained via bootstrap and are clustered at the subject level.

Parameter		Estimate	Standard error
Noise	η	0.599	(0.034)
Speed	β	0.674	(0.021)

Deviation from equilibrium					
Strong coordination		Weak coordination		No coordination	
Predicted	Observed	Predicted	Observed	Predicted	Observed
0.411	0.226	0	0.070	0	0.067

Table 5 presents maximum likelihood estimates for the logit dynamic and compares the predicted deviation from equilibrium to the empirically observed deviation from equilibrium under each treatment. Under the logit dynamic, the rate ρ_{ij}^p at which agents in population p switch from strategy i to strategy j per unit time depends on the payoff to each strategy and is given by

$$\rho_{ij}^p = \frac{\exp(\eta^{-1}\pi_i^p)}{\sum_{k \in S^p} \exp(\eta^{-1}\pi_k^p)} \quad (4)$$

The noise parameter η denotes the level of noise in the perception of payoffs. Accordingly, the probability that an agent in population p switches from strategy i to strategy j is given by $P_{ij}^p = \beta \rho_{ij}^p$ where β denotes the adjustment rate. Standard errors are obtained via bootstrap and are clustered at the subject level. As illustrated by Fig. 4, the logit dynamic's predicted deviation from equilibrium under the estimated parameters is 0.411 in the strong coordination treatment, 0 in the weak coordination treatment, and 0 in the no coordination treatment. Consistent with these theoretical predictions, significantly larger deviations from equilibrium were observed in the strong coordination treatment than in the other two treatments.

7. Conclusion

This study experimentally tests the predictive validity of evolutionary models in coordinated attacker–defender games. The experiment implements three treatment conditions, one with strong coordination incentives, one with weak coordination incentives, and one with zero coordination incentives. Evolutionary models yield distinct predictions for each treatment, but Nash equilibrium and several other behavioral models yield identical predictions under all three experimental treatments. Evolutionary stability criteria classify the unique Nash equilibrium of the coordinated attacker–defender game as neutrally stable under zero coordination incentives but strictly unstable under all positive coordination incentives.

This paper investigates three widely employed evolutionary dynamics: the best response dynamic, the Smith dynamic, and the logit dynamic. The Smith dynamic and the best response dynamic satisfy the widely maintained assumption of sign-preservation. They predict global convergence to equilibrium under zero coordination incentives but convergence to a stable limit cycle under nonzero coordination incentives. The logit dynamic does not impose sign preservation. It predicts global convergence to equilibrium under sufficiently weak coordination incentives but convergence to a stable limit cycle under sufficiently strong coordination incentives. All three dynamics predict that stable limit cycles will exhibit greater deviation from equilibrium under stronger coordination incentives.

Under the zero coordination treatment and the weak coordination treatment, observed behavior was tightly clustered around the Nash equilibrium. Under the strong coordination treatment, behavior was widely dispersed from the Nash equilibrium. This result was anticipated by evolutionary dynamics but not by evolutionary stability criteria, which classify the strong and weak treatments identically. This result suggests that evolutionary dynamics can characterize behavior with greater precision than evolutionary stability criteria. In contrast to the conventional class of sign-preserving dynamics, subjects frequently switched to lower performing strategies, suggesting that non-sign-preserving dynamics may provide a more accurate characterization of human behavior.

Sign preservation is frequently assumed by dynamic evolutionary models because it helps ensure that the fixed points of an evolutionary dynamic coincide with the set of Nash equilibria. The frequent violation of sign-preservation in the experimental data suggests that long run behavioral outcomes may often fail to correspond with Nash equilibria. Further research is needed to determine the extent to which non-sign-preserving evolutionary dynamics can predict and explain deviations from Nash equilibrium in a wider class of strategic environments. These results have important policy implications for strategic environments that are commonly modeled by attacker–defender games such as network security (Ibidunmoye et al., 2013), property crime (Bruni et al., 2013), and counter-terrorism (Powell, 2007). They suggest that, in such environments, policy makers ought not to rely exclusively on equilibrium models and stability criteria, but should also employ explicitly adaptive models.

Appendix

Proof of Proposition 1. Let x denote the social state under which both populations are evenly divided between their two pure strategies, so $\pi_1^A(x) = \pi_2^A(x) = \pi_1^D(x) = \pi_2^D(x) = \frac{1}{2}M + \frac{1}{2}C$ and x is a Nash equilibrium social state. Now let $y \neq x$ denote some alternate social state under which at least one population is unevenly divided between their two pure strategies. Without loss of generality, suppose that a larger fraction of the attacker population chooses to attack target 1, so $y_1^A > y_2^A$. Assume for contradiction that y is a Nash equilibrium social state. If $y_1^A = 1$ then $\pi_1^D(y) = M + Cy_1^D > Cy_2^D = \pi_2^D(y)$ since $M > C$ and $y^D \in [0, 1]^2$. Otherwise we have $y_1^A > y_2^A > 0$ so the payoff to attacking target 1 must equal the payoff to attacking target 2. In this case $My_2^D + Cy_1^A = My_1^D + Cy_2^A$ and we have $y_1^D > y_2^D$ since $y_1^A > y_2^A$. Then $My_1^A + Cy_2^D > My_2^A + Cy_1^D$ and we still have $\pi_1^D(y) > \pi_2^D(y)$. Thus, in either case, we must have $y_1^D = 1$ since defenders earn a higher payoff by defending target 1. Thus $\pi_2^A(y) = M + Cy_2^A > Cy_1^A = \pi_1^A(y)$ since $M > C$ and $y^A \in [0, 1]^2$. But then attackers would earn a higher payoff by attacking target 2, contradicting the assumption that y is a Nash equilibrium social state. □

Proof of Proposition 2.

Neutral stability: Under zero coordination incentives we have $C = 0$. Let x denote the equilibrium social state under which both populations are evenly divided between their two pure strategies. For each population $p \in \{A, D\}$ and strategy $i \in \{1, 2\}$ we have $\pi_i^p(x) = \frac{1}{2}M$ since $C = 0$. Let y denote some other social state. Now since $C = 0$ we have $\pi_1^A(y) = My_2^D, \pi_2^A(y) = My_1^D, \pi_1^D(y) = My_1^A$, and $\pi_2^D(y) = My_2^A$. Thus $x \cdot \pi(y) = y \cdot \pi(x) = y \cdot \pi(y) = M$ since $y_2^D + y_1^D = 1 = y_1^A + y_2^A$. Now let $\alpha \in (0, 1)$ and $z = \alpha y + (1 - \alpha)x$. Then we have $x \cdot \pi(z) = y \cdot \pi(z) = M$, so if a fraction of the population deviates away from equilibrium to some alternate strategy they never earn a higher average payoff than agents who maintain equilibrium behavior.

Strict instability: Under positive coordination incentives we have $C > 0$. Let x denote the equilibrium social state under which both populations are evenly divided between their two pure strategies, so for each population $p \in \{A, D\}$ and strategy $i \in \{1, 2\}$ we have $x_i^p = \frac{1}{2}$ and $\pi_i^p(x) = \frac{1}{2}M + \frac{1}{2}C$. Let y denote the social state under which attackers always attack target 1 and defenders always defend target 1, so $y_1^A = y_1^D = 1, y_2^A = y_2^D = 0, \pi_1^A(y) = C$, and $\pi_1^D(y) = M + C$. Hence $x \cdot \pi(y) = y \cdot \pi(x) = y \cdot \pi(y) = M + C$, but $y \cdot \pi(y) = M + 2C$. Now let $\alpha \in (0, 1)$ and $z = \alpha y + (1 - \alpha)x$. Then $x \cdot \pi(z) = M + C$ but $y \cdot \pi(z) = M + C + \alpha C$, so if any fraction of the population deviates to a strategy under which all agents select the same target they never earn a lower average payoff than agents who maintain equilibrium behavior. Further, under positive coordination incentives we have $C > 0$ so $x \cdot \pi(z) = M + C < M + C + \alpha C = y \cdot \pi(z)$, so the deviating agents earn a strictly higher payoff than the agents who maintain equilibrium behavior.

Positive definiteness: Let $x = (x_1^D, x_2^D, x_1^A, x_2^A) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}) \in \mathbb{R}^4$ denote the Nash equilibrium social state. So the linear approximation $\pi'(x) \in \mathbb{R}^{4 \times 4}$ at the equilibrium point x is given by

$$\pi'(x) = \begin{bmatrix} \nabla \pi_1^D(x)^T \\ \nabla \pi_2^D(x)^T \\ \nabla \pi_1^A(x)^T \\ \nabla \pi_2^A(x)^T \end{bmatrix} = \begin{bmatrix} C & 0 & M & 0 \\ 0 & C & 0 & M \\ 0 & M & C & 0 \\ M & 0 & 0 & C \end{bmatrix} \tag{5}$$

Let $z = (z_1^D, z_2^D, z_1^A, z_2^A) \in \mathbb{R}^4$ such that $z_1^D + z_2^D = 0$ and $z_1^A + z_2^A = 0$. Then the corresponding quadratic form is given by $z^T \pi'(x) z = M(z_1^D z_1^A + z_2^D z_2^A + z_1^A z_2^D + z_2^A z_1^D) + C \sum_{p=1}^2 \sum_{i=1}^2 (z_i^p)^2$. Now since $z_2^D = -z_1^D$ and $z_2^A = -z_1^A$ we have $z_1^D z_1^A + z_2^D z_2^A + z_1^A z_2^D + z_2^A z_1^D = 0$ so $z^T \pi'(x) z = C \sum_{p=1}^2 \sum_{i=1}^2 (z_i^p)^2$. Thus $\pi'(x)$ is positive definite on Z when $C > 0$ and indefinite on Z when $C = 0$. □

Proof of Proposition 3. Let $\psi(y) = (1 + e^{-y/\eta})^{-1} - \frac{1}{2}$ so $\psi'(y) = \eta e^{-y/\eta} (1 + e^{-y/\eta})^{-2}$, $\psi(0) = 0$, $\psi'(0) = \frac{1}{4\eta}$, and $\psi'(y) < \frac{1}{4\eta}$ for all $y > 0$. Hence $\psi(y) < \frac{y}{4\eta}$ for all $y > 0$. In the coordinated attacker defender game, the conditional switch rates are given by $\rho_{ij}^p = \psi(\pi_j^p - \pi_i^p) + \frac{1}{2}$ under the logit dynamic. Since ψ is an odd function, the difference between the conditional switch rates is given by $\rho_{ij}^p - \rho_{ji}^p = 2\psi(\pi_j^p - \pi_i^p) = 2\rho_{ij}^p + 1$. Let $z_A = x_1^A - \frac{1}{2}$ and $z_D = x_1^D - \frac{1}{2}$ so

$$\dot{z}_p = x_2^p \rho_{21}^p - x_1^p \rho_{12}^p = (1 - x_1^p) \rho_{21}^p - x_1^p \rho_{12}^p \tag{6}$$

$$= (1 - 2x_1^p) \rho_{21}^p + x_1^p (\rho_{21}^p - \rho_{12}^p) \tag{7}$$

$$= (1 - 2x_1^p) \rho_{21}^p + x_1^p (2\rho_{21}^p - 1) \tag{8}$$

$$= \rho_{21}^p - x_1^p = \psi(\pi_1^p - \pi_2^p) - z_p \tag{9}$$

In the coordinated attacker defender game $\pi_1^A - \pi_2^A = 2Cz_A - 2Mz_D$ and $\pi_1^D - \pi_2^D = 2Cz_A + 2Mz_D$. Hence $\dot{z}_A = \psi(2Cz_A - 2Mz_D) - z_A$ and $\dot{z}_D = \psi(2Cz_D + 2Mz_A) - z_D$. Let $V(z_A, z_D) = z_A^2 + z_D^2$ so

$$\dot{V}(z_A, z_D) = \frac{\partial V}{\partial z_A} \dot{z}_A + \frac{\partial V}{\partial z_D} \dot{z}_D \quad (10)$$

$$= 2z_A[\psi(2Cz_A - 2Mz_D) - z_A] + 2z_D[\psi(2Cz_D + 2Mz_A) - z_D] \quad (11)$$

$$= 2[z_A\psi(2Cz_A - 2Mz_D) + z_D\psi(2Cz_D + 2Mz_A)] - 2(z_A^2 + z_D^2) \quad (12)$$

Let $y = \max(|z_A|, |z_D|)$. If $C > 0$, $y > 0$, and $\eta > C + M$ then $\dot{V}(z_A, z_D) \leq 4y\psi(2[C + M]y) - 2y^2 < 0$ since $\psi(y) < \frac{y}{4\eta}$ for all $y > 0$. If $C = 0$, $y > 0$, and $\eta > 0$ then $\dot{V}(z_A, z_D) \leq 4y\psi(2My) - 2y^2 < 0$. \square

References

- Alos-Ferrer, C., Netzer, N., 2010. The logit-response dynamics. *Games Econ. Behav.* 68 (2), 413–427.
- Aumann, R., Brandenburger, A., 1995. Epistemic conditions for Nash equilibrium. *Econometrica* 63 (5), 1161–1180.
- Benaïm, M., Hofbauer, J., Hopkins, E., 2009. Learning in games with unstable equilibria. *J. Econ. Theory* 144 (4), 1694–1709.
- Benndorf, V., Martinez-Martinez, I., 2017. Perturbed best response dynamics in a hawk–dove game. *Econ. Lett.* 153, 61–64.
- Benndorf, V., Martinez-Martinez, I., Normann, H.-T., 2016. Equilibrium selection with coupled populations in hawk–dove games: theory and experiment in continuous time. *J. Econ. Theory* 165, 472–486.
- Björnerstedt, J., Weibull, J., 1994. Nash equilibrium and evolution by imitation (No. 407). IUI working paper.
- Bruni, C., Nuño, J.C., Primicerio, M., 2013. What if criminals optimize their choice? Optimal strategies to defend public goods. *Physica A* 392 (4), 840–850.
- Bulò, S.R., Bomze, I.M., 2011. Infection and immunization: a new class of evolutionary game dynamics. *Games Econ. Behav.* 71 (1), 193–211.
- Cason, T., Friedman, D., Hopkins, E., 2013. Cycles and instability in a rock–paper–scissors population game: a continuous time experiment. *Rev. Econ. Stud.* 81 (1), 112–136.
- Cournot, A.-A., 1838. *Researches into the Mathematical Principles of the Theory of Wealth*. chez L. Hachette.
- Crawford, V.P., 1991. An evolutionary interpretation of van Huyck, Battalio, and Beil's experimental results on coordination. *Games Econ. Behav.* 3 (1), 25–59.
- Friedman, D., Huck, S., Oprea, R., Weidenholzer, S., 2015. From imitation to collusion: long-run learning in a low-information environment. *J. Econ. Theory* 155, 185–205.
- Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning In Games*, vol. 2. MIT Press.
- Gilboa, I., Matsui, A., 1991. Social stability and equilibrium. *Econometrica*, 859–867.
- Hofbauer, J., Sigmund, K., 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- Hopkins, E., 1999. A note on best response dynamics. *Games Econ. Behav.* 29 (1), 138–150.
- Hopkins, E., Seymour, R.M., 2002. The stability of price dispersion under seller and consumer learning. *Int. Econ. Rev.* 43 (4), 1157–1190.
- Ibidunmoye, E., Alese, B., Ogundele, O., 2013. Modeling attacker–defender interaction as a zero-sum stochastic game. *J. Comput. Sci. Appl.* 1 (2), 27–32.
- Lim, W., Neary, P.R., 2016. An experimental investigation of stochastic adjustment dynamics. *Games Econ. Behav.* 100, 208–219.
- McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. *Games Econ. Behav.* 10 (1), 6–38.
- Nash, J., 1950. *Equilibrium Points in n-Person Games*. PhD thesis. Princeton University.
- Oprea, R., Henwood, K., Friedman, D., 2011. Separating the hawks from the doves: evidence from continuous time laboratory games. *J. Econ. Theory* 146 (6), 2206–2225.
- Powell, R., 2007. Defending against terrorist attacks with limited resources. *Am. Polit. Sci. Rev.* 101 (03), 527–541.
- Sandholm, W.H., 2009. Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. *Games* 1 (1), 3–17.
- Smith, M.J., 1984. The stability of a dynamic model of traffic assignment – an application of a method of Lyapunov. *Transp. Sci.* 18 (3), 245–252.
- Smith, J.M., Price, G.R., 1973. The logic of animal conflict. *Nature* 246, 15.
- Stahl, D.O., Wilson, P.W., 1994. Experimental evidence on players' models of other players. *J. Econ. Behav. Organ.* 25 (3), 309–327.
- Taylor, P.D., 1979. Evolutionarily stable strategies with two types of player. *J. Appl. Probab.* 16 (1), 76–83.
- Van Huyck, J.B., Battalio, R.C., Beil, R.O., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. *Am. Econ. Rev.* 80 (1), 234–248.