

Best experienced payoff dynamics and cooperation in the centipede game

WILLIAM H. SANDHOLM

Department of Economics, University of Wisconsin

SEGISMUNDO S. IZQUIERDO

BioEcoUva, Department of Industrial Organization, Universidad de Valladolid

LUIS R. IZQUIERDO

Department of Civil Engineering, Universidad de Burgos

We study population game dynamics under which each revising agent tests each of his strategies a fixed number of times, with each play of each strategy being against a newly drawn opponent, and chooses the strategy whose total payoff was highest. In the centipede game, these *best experienced payoff dynamics* lead to cooperative play. When strategies are tested once, play at the almost globally stable state is concentrated on the last few nodes of the game, with the proportions of agents playing each strategy being largely independent of the length of the game. Testing strategies many times leads to cyclical play.

KEYWORDS. Evolutionary game theory, backward induction, centipede game, computational algebra.

JEL CLASSIFICATION. C72, C73.

1. INTRODUCTION

The discrepancy between the conclusions of backward induction reasoning and observed behavior in certain canonical extensive form games is a basic puzzle of game theory. The centipede game (Rosenthal (1981)), the finitely repeated prisoner's dilemma, and related examples can be viewed as models of relationships in which each participant has repeated opportunities to take costly actions that benefit his partner and in which there is a commonly known date at which the interaction will end. Experimental and anecdotal evidence suggests that cooperative behavior may persist until close to

William H. Sandholm: whs@ssc.wisc.edu

Segismundo S. Izquierdo: segis@eii.uva.es

Luis R. Izquierdo: lrizquierdo@ubu.es

We thank Dan Friedman, Drew Fudenberg, Ken Judd, Panayotis Mertikopoulos, Erik Mohlin, Ignacio Monzón, Arthur Robson, Marzena Rostek, Ariel Rubinstein, Larry Samuelson, Ryoji Sawa, Andy Skrzypacz, Lones Smith, Mark Voorneveld, Marek Weretka, and, especially, Antonio Penta for helpful discussions and comments. Financial support from the U.S. National Science Foundation (Grants SES-1458992 and SES-1728853), the U.S. Army Research Office (Grants W911NF-17-1-0134 MSN201957), Project ECO2017-83147-C2-2-P (MINECO/AEI/FEDER, UE), and the Spanish Ministerio de Educación, Cultura, y Deporte (Grants PRX15/00362 and PRX16/00048) is gratefully acknowledged.

the exogenous terminal date (McKelvey and Palfrey (1992)). But the logic of backward induction leads to the conclusion that there will be no cooperation at all.

Work on epistemic foundations provides room for wariness about unflinching appeals to backward induction. To support this prediction, one must assume that there is always common belief that all players will act as payoff maximizers at all points in the future, even when many rounds of previous choices argue against such beliefs.¹ Thus, the simplicity of backward induction belies the strength of the assumptions needed to justify it, and this strength may help explain why backward induction does not yield descriptively accurate predictions in some classes of games.²

This paper studies a dynamic model of behavior in games that maintains the assumption that agents respond optimally to the information they possess. But rather than imposing strong assumptions about agents' knowledge of opponents' intentions, we suppose instead that agents' information comes from direct but incomplete experience with playing the strategies available to them. As with the earlier work of Osborne and Rubinstein (1998) and Sethi (2000), our model is best viewed not as one that incorporates irrational choices, but rather as one of rational choice under particular restrictions on what agents know.

Following the standard approach of evolutionary game theory, we suppose that two populations of agents are recurrently randomly matched to play a two-player game. This framework accords with some experimental protocols, and can be understood more broadly as a model of the formation of social norms (Young (1998)). At random times, each agent receives opportunities to switch strategies. At these moments the agent plays each of his strategies against κ opponents drawn at random from the opposing population, with each play of each strategy being against a newly drawn opponent. He then switches to the strategy that achieved the highest total payoff, breaking ties in favor of the lowest-numbered strategy. Standard results imply that when the populations are large, the agents' aggregate behavior evolves in an essentially deterministic fashion, obeying a differential equation that describes the expected motion of the stochastic process described above (Benaïm and Weibull (2003)). We study the properties of this differential equation when agents play the centipede game.

Our model builds on earlier work on games played by "procedurally rational players." If we replace our tie-breaking rule with uniform tie-breaking, then the rest points of the process (with $\kappa = k$) would correspond to the $S(k)$ equilibria of Osborne and Rubinstein (1998). The corresponding dynamics were studied by Sethi (2000). These and other dynamics are instances of the broader family of *best experienced payoff dynamics* (BEP dynamics for short; Sandholm et al. (2019)), which allow for variation in how ties are resolved and in the selection of sets of candidate strategies considered by revising agents. The results we present here are robust to many different model specifications within the family of BEP dynamics.

¹For formal analyses, see Binmore (1987), Reny (1992), Stalnaker (1996), Ben-Porath (1997), Halpern (2001), and Perea (2014).

²As an alternative, one could apply Nash equilibrium, which also predicts noncooperative behavior in the games mentioned above, but doing so replaces assumptions about future rationality with the assumption of equilibrium knowledge, which may not be particularly more appealing; see Dekel and Gul (1997).

Our analysis of best experienced payoff dynamics in the centipede game uses techniques from dynamical systems theory. What is more novel is our reliance on algorithms from computational algebra and perturbation bounds from linear algebra, which allow us to solve exactly for the rest points of our differential equations and to perform rigorous stability analyses in centipede games with up to six decision nodes. We complement this approach with numerical analyses of cases in which analytical results cannot be obtained.

Our initial results focus on dynamics under which each tested strategy is tested exactly once ($\kappa = 1$), so that agents' choices depend only on ordinal properties of payoffs. In centipede games, under the BEP dynamics studied here, the backward induction state—the state at which all agents in both populations stop at their first opportunity—is a rest point. However, we prove that this rest point is always *repelling*: the appearance of agents in either population who cooperate to any degree is self-reinforcing and eventually causes the backward induction solution to break down completely.

We next obtain strong lower bounds on the total weight placed on cooperative strategies at any other rest points of the BEP dynamic. At any such rest point, the probability that play during a random match leads to one of the last five terminal nodes is above 0.96, and the probability that play leads to one of the last seven terminal nodes is virtually 1. We then use tools from computational algebra to perform an exact analysis of games with up to six decision nodes, and we perform numerical analyses of longer games. In all cases, we find that besides the unstable backward induction state, the dynamics have exactly one other rest point.³ The form of this rest point is essentially independent of the length of the game. The rest point has virtually all players choosing to continue until the last few nodes of the game. Moreover, this rest point is dynamically stable, attracting solutions from all initial conditions other than the backward induction state. Thus if agents make choices based on experienced payoffs, testing each strategy once and choosing the one that performed best, then play converges to a stable rest point that exhibits high levels of cooperation.

To explain why, we first observe that cooperative strategies are most disadvantaged when they are most rare—specifically, in the vicinity of the backward induction state. Near this state, the most cooperative agents would obtain higher *expected payoffs* by stopping earlier. However, when an agent considers switching strategies, he tests each of his strategies against new, independently drawn opponents. He may thus test a cooperative strategy against a cooperative opponent, and test less cooperative strategies against less cooperative opponents, in which case his best *experienced payoff* will come from the cooperative strategy. Our analysis confirms that this possibility indeed leads to instability.⁴ After this initial entry, the high payoffs generated by cooperative strategies

³While traditional equilibrium notions in economics require stasis of choice, interior rest points of population dynamics represent situations in which individuals' choices fluctuate even as the expected change in *aggregate* behavior is null; see Section 2.2.

⁴Specifically, linearizing any given specification of the dynamics at the backward induction state identifies a single eigenvector with a positive eigenvalue (Appendix A). This eigenvector describes the mixture of strategies in the two populations whose entry is self-reinforcing and identifies the direction toward which all other disturbances of the backward induction state are drawn. Direct examination of the dynamics provides a straightforward explanation of why the given mixture of entrants is successful (Example 2).

when matched against one another spurs their continued growth. This growth is only abated when virtually all agents are choosing among the most cooperative strategies.

Our final results consider the effects of the number of trials κ of each strategy during testing on predictions of play. It seems clear that if the number of trials is made sufficiently large, so that the agents' information about opponents' behavior is quite accurate, then the population's behavior should come to resemble a Nash equilibrium. Indeed, when agents possess exact information, so that aggregate behavior evolves according to the best response dynamic (Gilboa and Matsui (1991), Hofbauer (1995)), the results of Xu (2016) imply that every solution trajectory converges to the set of Nash equilibria, all of which entail stopping at the initial node.

Our analysis shows, however, that stable cooperative behavior can persist even for substantial numbers of trials. To start, we prove that the backward induction state is unstable as long as the number of trials κ is less than the length of the game. For a larger number of trials, the backward induction state becomes locally stable, but numerical evidence suggests that its basin of attraction is very small. Examining centipede games of length $d = 4$ in detail, we find that a unique, attracting interior rest point with substantial cooperation persists for moderate numbers of trials. With many trials, numerical analysis suggests the attractor is always a single cycle that includes significant amounts of cooperation for numbers of trials as large as 200. We discuss in Section 4 how the robustness of cooperation to fairly large numbers of trials can be explained using simple central limit theorem arguments.

Our main technical contribution lies in the use of methods from computational algebra and perturbation theorems from linear algebra to prove results about the properties of our dynamics. The starting point for this analysis—one that suggests a broader scope for our approach—is that decision procedures based on sampling from a population are described by multivariate polynomials with rational coefficients. In particular, BEP dynamics are described by systems of such equations, so finding their rest points amounts to finding the zeros of these polynomial systems. To accomplish this, we compute a *Gröbner basis* for the set of polynomials that defines each instance of our dynamics; this new set of polynomials has the same zeros as the original set, but its zeros can be computed by finding the roots of a single (possibly high-degree) univariate polynomial.⁵ Exact representations of these roots, known as *algebraic numbers*, can then be obtained by factoring the polynomial into irreducible components and then using algorithms based on classical results to isolate each component's real roots.⁶ With these exact solutions in hand, we can rigorously assess the rest points' local stability through a linearization analysis. So as to obviate certain intractable exact calculations, this analysis takes advantage of both an eigenvalue perturbation theorem and a bound on the condition number of a matrix that does not require the computation of its inverse.

The code used to obtain the exact and numerical results is available as a Mathematica notebook posted on GitHub and on the authors' websites. The Supplemental Material provides background and details about both the exact and the numerical analyses, and reports certain numerical results in full detail.

⁵See Buchberger (1965) and Cox et al. (2015). For applications of Gröbner bases in economics, see Kubler et al. (2014).

⁶See von zur Gathen and Gerhard (2013), McNamee (2007), and Akritas (2010).

Related literature

Previous work relating backward induction and deterministic evolutionary dynamics has focused on the replicator dynamic of [Taylor and Jonker \(1978\)](#) and the best response dynamic of [Gilboa and Matsui \(1991\)](#) and [Hofbauer \(1995\)](#). [Cressman and Schlag \(1998\)](#) (see also [Cressman \(1996, 2003\)](#)) show that in generic perfect information games, every interior solution trajectory of the replicator dynamic converges to a Nash equilibrium. Likewise, [Xu \(2016\)](#) (see also [Cressman \(2003\)](#)) shows that in such games, every solution trajectory of the best response dynamic converges to a component of Nash equilibria. In both cases, the Nash equilibria approached need not be subgame perfect and the Nash equilibrium components generally are not locally stable. Focusing on the centipede game with three decision nodes, [Ponti \(2000\)](#) shows numerically that perturbed versions of the replicator dynamic exhibit cyclical behavior, with trajectories approaching and then moving away from the Nash component. In contrast, we show that for small and moderate numbers of tests, best experienced payoff dynamics lead to a stable distribution of cooperative strategies far from the Nash component.

[Osborne and Rubinstein's \(1998\)](#) notion of $S(k)$ equilibrium corresponds to the rest points of the BEP dynamic under which agents test all strategies, subject each to k trials, and break ties via uniform randomization.⁷ While most of their analysis focuses on simultaneous move games, they show that in centipede games, the probability with which player 1 stops immediately in any $S(1)$ equilibrium must vanish as the length of the game grows large. As we soon observe ([Observation 1](#)), this conclusion may fail if uniform tie-breaking is not assumed, with the backward induction state being an equilibrium. Nevertheless, more detailed analyses below show that this equilibrium state is unstable under BEP dynamics.

Building on [Osborne and Rubinstein \(1998\)](#), [Sethi \(2000\)](#) introduces BEP dynamics under which all strategies are tested and ties are broken uniformly.⁸ He shows that both dominant strategy equilibria and strict equilibria can be unstable under these dynamics, while dominated strategies can be played in stable equilibria. The latter fact is a basic component of our analysis of cooperative behavior. [Berkemer \(2008\)](#) considers the local stability of the unique rationalizable strategy profile in the traveler's dilemma of [Basu \(1994\)](#) under [Sethi's \(2000\)](#) dynamics, obtaining a sufficient condition for the instability of the rationalizable state. He shows numerically that the stable $S(1)$ equilibrium becomes independent of the number of strategies in the game and he provides evidence from agent-based simulations that larger numbers of trials during testing can lead to cyclical behavior.⁹

Earlier efforts to explain cooperative behavior in centipede and related games have followed a different approach, applying equilibrium analyses to augmented versions of the game. The best known example of this approach is the work of [Kreps et al. \(1982\)](#).

⁷For extensions of $S(k)$ equilibrium to more complex testing procedures, see [Rustichini \(2003\)](#).

⁸[Cárdenas et al. \(2015\)](#) and [Mantilla et al. \(2019\)](#) use these dynamics to explain stable non-Nash behavior in public goods games.

⁹For complementary models of dynamics based on a single sample, see [Sandholm \(2001\)](#), [Kosfeld et al. \(2002\)](#), [Droste et al. \(2003\)](#), and [Oyama et al. \(2015\)](#).

These authors modify the finitely repeated prisoner's dilemma by assuming that one player attaches some probability to his opponent having a fixed preference for cooperative play. They show that in all sequential equilibria of long enough versions of the resulting Bayesian game, both players act cooperatively for a large number of initial rounds.¹⁰ To justify this approach, one must assume that the augmentation of the original game is commonly understood by the players, that the players act in accordance with a rather complicated equilibrium construction, and that the equilibrium knowledge assumptions required to justify sequential equilibrium apply. In contrast, our model makes no changes to the original game other than placing it in a population setting, and it is built upon the assumption that agents' choices are optimal given their experiences during play.

2. BEST EXPERIENCED PAYOFF DYNAMICS IN THE CENTIPEDE GAME

2.1 Normal form games and population games

A two-player normal form game $G = \{(S^1, S^2), (A, B)\}$ is defined by pairs of strategy sets $S^p = \{1, \dots, s^p\}$ and payoff matrices $A, B \in \mathbb{R}^{s^p \times s^q}$, $p, q \in \{1, 2\}$, $p \neq q$. Entries A_{ij} and B_{ij} represent the two players' payoffs when strategy profile $(i, j) \in S^1 \times S^2$ is played. When considering extensive form games, our analysis focuses on the reduced normal form, whose strategies specify an agent's "plan of action" for the game, but not his choices at decision nodes that are ruled out by his own previous choices.

In our population model, members of two unit-mass populations are matched to play a two-player game. A *population state* for population 1 is an element of $X = \{x \in \mathbb{R}_+^{s^1} : \sum_{i \in S^1} x_i = 1\}$, where x_i is the fraction of population 1 players choosing strategy i . Likewise $Y = \{y \in \mathbb{R}_+^{s^2} : \sum_{i \in S^2} y_i = 1\}$ is the set of population states for population 2. Thus, x and y are formally equivalent to mixed strategies for players 1 and 2, and elements of the set $\Xi = X \times Y$ are formally equivalent to mixed strategy profiles. In a slight abuse of terminology, we also refer to elements of Ξ as population states.

2.2 Revision protocols and evolutionary dynamics

To define evolutionary game dynamics, we follow the standard approach of specifying microfoundations in terms of revision protocols.¹¹ We suppose that at all times $t \in [0, \infty)$, each agent has a strategy he uses when matched to play game G . The empirical distributions of these strategies are described by the population state $\xi(t) = (x(t), y(t))$.

¹⁰McKelvey and Palfrey (1992) show that this analysis extends to the centipede game. A different augmentation is considered by Jehiel (2005), who assumes that agents bundle decision nodes from contiguous stages into analogy classes and view the choices at all nodes in a class interchangeably. Alternatively, following Radner (1980), one can consider versions of centipede in which the stakes of each move are small and analyze these games using ε equilibrium; see Friedman and Oprea (2012) for a discussion. But as Binmore (1998) observes, the existence of a non-Nash ε equilibrium depends on the relative sizes of the stakes and of ε , and the backward induction solution always persists as a Nash equilibrium and, hence, as an ε equilibrium.

¹¹See Björnerstedt and Weibull (1996), Weibull (1995), Sandholm (2010a,b, 2015), and Izquierdo et al. (2018).

Agents occasionally receive opportunities to switch strategies according to independent rate 1 Poisson processes. An agent who receives an opportunity considers switching to a new strategy, making his decision by applying a revision protocol. Formally, a *revision protocol* for population 1 is described by a map $(A, y) \mapsto \sigma^1(A, y) \in X^{S^1}$ that assigns own payoff matrices and opposing population states to matrices of *conditional switch probabilities*, where $\sigma_{ij}^1(A, y)$ is the probability that an agent playing strategy $i \in S^1$ who receives a revision opportunity switches to strategy $j \in S^1$. Likewise, a revision protocol for population 2 is described by a map $(B, x) \mapsto \sigma^2(B, x) \in Y^{S^2}$ with an analogous interpretation.¹²

It is well known that if the population sizes are large, the Markov process implicitly defined by the above procedure is well approximated by solutions to a differential equation defined by the expected motion of the process (Benaïm and Weibull (2003)). Here this differential equation takes the form

$$\begin{aligned} \dot{x}_i &= \sum_{j \in S^1} x_j \sigma_{ji}^1(A, y) - x_i \quad \text{for all } i \in S^1, \\ \dot{y}_i &= \sum_{j \in S^2} y_j \sigma_{ji}^2(B, x) - y_i \quad \text{for all } i \in S^2. \end{aligned} \tag{1}$$

Equation (1) is easy to interpret. Since revision opportunities are assigned to agents randomly, there is an outflow from each strategy i proportional to its current level of use. To generate inflow to i , an agent playing some strategy j must receive a revision opportunity, and applying his revision protocol must lead him to play strategy i .

Outside of monomorphic (i.e., pure) cases, the rest points of the dynamic (1) should not be understood as equilibria in the traditional game-theoretic sense. Rather, they represent situations in which agents perpetually switch among strategies, but with the *expected* change in the use of each strategy equaling zero.¹³ At states that are locally stable under the dynamic (1), fluctuations in any direction are generally undone by the action of (1) itself. Contrariwise, fluctuations away from unstable equilibria are reinforced, so we should not expect such states to be observed.

2.3 Best experienced payoff protocols and dynamics

We now introduce the class of revision protocols and dynamics that we study in this paper. A *best experienced payoff protocol* is defined by a triple (τ, κ, β) consisting of a *test set rule* τ , a *number of trials* κ , and a *tie-breaking rule* β . The triple (τ, κ, β) defines a revision protocol in the following way. When an agent currently using strategy $i \in S^p$ receives an opportunity to switch strategies, he draws a set of strategies $R^p \subseteq S^p$ to test according to the distribution τ^p on the collection of subsets of S^p with at least two elements. He then plays each strategy in R^p in κ random matches against members of

¹²When σ_{ij}^1 and σ_{ij}^2 are independent of the current strategy i , as is true for the dynamic (3) we focus on here, it is equivalent to interpret the process as one in which agents play a fixed strategy until leaving the population, when they are replaced by new agents whose strategies are determined by applying σ^1 and σ^2 .

¹³Thus in the finite-population version of the model, variations in the use of each strategy would be observed. For a formal analysis, see Sandholm (2003).

the opposing population. He thus engages in $\#R^p \times \kappa$ random matches in total, facing distinct sets of opponents when testing different strategies. The agent then selects the strategy in R^p that earned him the highest total payoff, breaking ties according to rule β . The triple (τ, κ, β) thus defines a revision protocol σ^p for each population p . Inserting these revision protocols into (1) defines a *best experienced payoff dynamic*.

Our analysis here focuses on the test-set rule *test-all*, τ^{all} , under which a revising agent tests all of his strategies, and on the tie-breaking rule *min-if-tie*, β^{min} , which chooses the lowest-numbered optimal strategy. We refer to the resulting dynamics (1) as *BEP*($\tau^{\text{all}}, \kappa, \beta^{\text{min}}$) *dynamics*.

BEP dynamics based on other specifications of test-set and tie-breaking rules are studied in a companion paper (Sandholm et al. (2019)); they are also discussed briefly in Section 3.4. Importantly, if we retain τ^{all} , but replace β^{min} with uniform tie-breaking, then the rest points of the dynamic (1) are the $S(k)$ equilibria of Osborne and Rubinstein (1998) (with $k = \kappa$), and the dynamic itself is the one studied by Sethi (2000). In extensive form games like centipede, different strategies often earn the same payoffs, so the choice of tie-breaking rule matters. Because of our convention for numbering strategies in the centipede game (see below), the min-if-tie rule will be the one that is least conducive to cooperative play.

Choice probabilities under best experienced payoff dynamics depend only on the payoffs strategies earn during testing; they do not require agents to track the choices made by their opponents. This property makes the dynamics appealing as a simple model of play for extensive form games. Typically, a single play of an extensive form game does not reveal the strategy chosen by one's opponent, but only the portion of that strategy required to determine the path of play. Consequently, it is not straightforward to specify how agents should use their experience of play to assess opponents' choices of strategies. Because they focus on the performances of own strategies, best experienced payoff dynamics avoid such ambiguities.

2.4 The centipede game

Centipede (Rosenthal (1981)) is a two-player extensive form game with $d \geq 2$ decision nodes (Figure 1). Each node presents two actions: stop and continue. The nodes are arranged linearly, with the first one assigned to player 1 and subsequent nodes assigned in an alternating fashion. A player who stops ends the game. A player who continues suffers a cost of 1 but benefits his opponent 3, and sends the game to the next decision node if one exists.




FIGURE 1. The centipede game of length $d = 8$.

In a centipede game of length d , players 1 and 2 have $d^1 = \lfloor \frac{d+1}{2} \rfloor$ and $d^2 = \lfloor \frac{d}{2} \rfloor$ decision nodes, respectively. Thus player p has $s^p = d^p + 1$ strategies, where strategy $i < s^p$ is the plan to continue at his first $i - 1$ decision nodes and to stop at his i th decision node, and strategy s^p is the plan to continue at all d^p of his decision nodes. Of course, the portion of a player's plan that is actually carried out depends on the plan of his opponent. The payoff matrices (A, B) of centipede's reduced normal form can be expressed concisely as

$$(A_{ij}, B_{ij}) = \begin{cases} (2i - 2, 2i - 2) & \text{if } i \leq j \\ (2j - 3, 2j + 1) & \text{if } j < i. \end{cases} \tag{2}$$

It will sometimes be convenient to number strategies starting from the end of the game. To do so, we write $[k] \equiv s^p - k$ for $k \in \{0, \dots, d^p\}$, so that $[0]$ denotes continuing at all nodes, and $[k]$ with $k \geq 1$ denotes stopping at player p 's k th-to-last node.

We noted above that best experienced payoff dynamics with $\kappa = 1$ depend only on ordinal properties of payoffs. In this case, what matters in (2) is that a player is better off continuing at a given decision node if and only if his opponent will continue at the subsequent decision node. If the cost of continuing is 1, this property holds as long as the benefit obtained when one's opponent continues exceeds 2. This ordering of payoffs also holds for typical specifications in which total payoffs grow exponentially over time. When there are multiple trials of each tested strategy ($\kappa = 1$), then cardinal properties of payoffs matter; in this case, Rosenthal's (1981) specification (2) keeps the potential benefits from continuing relatively modest.

The backward induction solution to centipede has both players stop at each of their decision nodes. We thus call the population state $\xi^\dagger = (x^\dagger, y^\dagger) \in \Xi$ with $x_1^\dagger = y_1^\dagger = 1$ the (*reduced*) *backward induction state*. It is well known that all Nash equilibria of centipede have player 1 stop at his initial node. This makes player 2 indifferent among all of her strategies, so Nash equilibrium requires that she choose a mixed strategy that makes stopping immediately optimal for player 1.

Of course, these predictions require assumptions about what the players know. In the traditional justification of Nash equilibrium, players are assumed to anticipate opponents' play correctly. Likewise, traditional justifications of the backward induction solution require agents to maintain common belief in rational future play, even if behavior contradicting this belief has been observed in the past.

2.5 Best experienced payoff dynamics for the centipede game

We can now introduce the explicit formula for the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic in the centipede game:¹⁴

$$\dot{x}_i = \left(\sum_{k=i}^{s^2} y_k \right) \left(\sum_{m=1}^i y_m \right)^{s^1-i} + \sum_{k=2}^{i-1} y_k \left(\sum_{\ell=1}^{k-1} y_\ell \right)^{i-k} \left(\sum_{m=1}^k y_m \right)^{s^1-i} - x_i \tag{3a}$$

¹⁴We follow the convention that a sum whose lower limit exceeds its upper limit evaluates to 0.

$$\dot{y}_j = \begin{cases} \left(\sum_{k=2}^{s^1} x_k \right) (x_1 + x_2)^{s^2-1} + (x_1)^{s^2} - y_1 & \text{if } j = 1 \\ \left(\sum_{k=j+1}^{s^1} x_k \right) \left(\sum_{m=1}^{j+1} x_m \right)^{s^2-j} & \\ + \sum_{k=2}^j x_k \left(\sum_{\ell=1}^{k-1} x_\ell \right)^{j-k+1} \left(\sum_{m=1}^k x_m \right)^{s^2-j} - y_j & \text{otherwise.} \end{cases} \tag{3b}$$

Under test-all with min-if-tie, the choice made by a revising agent does not depend on his original strategy. The first two terms of (3a) describe the two types of matchings that lead a revising agent in the role of player 1 to choose strategy i . First, it could be that when the agent tests i , his opponent plays i or higher (so that the agent is the one to stop the game), and that when the agent tests higher strategies, his opponents play strategies i or lower. In this case, only strategy i yields the agent his highest payoff. Second, it could be that when the agent tests i , his opponent plays strategy $k < i$; when he tests strategies between k and $i - 1$, his opponents play strategies less than k ; and when he tests strategies above i , his opponents play strategies less than or equal to k . In this case, strategy i is the lowest strategy that achieves the optimal payoff, and so it is chosen by the revising agent under the min-if-tie rule. Similar logic and accounting for the fact that player 2's j th node is followed by player 1's $(j + 1)$ st node lead to (3b).

We conclude this section with a simple observation about the backward induction solution of centipede under best experienced payoff dynamics.

OBSERVATION 1. *Under the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic, the backward induction state ξ^\dagger is a rest point.*

Osborne and Rubinstein (1998) show that if all strategies are tested once and ties are broken uniformly, then in a long centipede game, stationarity requires that play is almost never stopped at the initial node. Observation 1 shows that this conclusion depends on the assumption that ties are broken uniformly. If instead ties are broken in favor of the lowest-numbered strategy (or, alternatively, an agent's current strategy), then the backward induction state is a rest point. Even so, the analyses to follow will explain why the backward induction state is not a compelling prediction of play even under these tie-breaking rules.

3. ANALYSIS OF DYNAMICS WITH ONE TRIAL OF EACH STRATEGY

In this section, we analyze the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic in centipede. Since tie-breaking rule β^{min} selects the optimal strategy that stops soonest, it is the tie-breaking rule that is most favorable toward backward induction. Before proceeding, we review some standard definitions and results from dynamical systems theory, and follow this with a simple example.

Consider a C^1 differential equation $\dot{\xi} = V(\xi)$ defined on Ξ whose forward solutions $\{x(t)\}_{t \geq 0}$ do not leave Ξ . State ξ^* is a rest point if $V(\xi^*) = 0$, so that the unique solution

starting from ξ^* is stationary. Rest point ξ^* is *Lyapunov stable* if for every neighborhood $O \subset \Xi$ of ξ^* , there exists a neighborhood $O' \subset \Xi$ of ξ^* such that every forward solution that starts in O' is contained in O . If ξ^* is not Lyapunov stable, it is *unstable*, and it is *repelling* if there is a neighborhood $O \subset \Xi$ of ξ^* such that solutions from all initial conditions in $O \setminus \{\xi^*\}$ leave O .

Rest point ξ^* is *attracting* if there is a neighborhood $O \subset \Xi$ of ξ^* such that all solutions that start in O converge to ξ^* . A state that is Lyapunov stable and attracting is *asymptotically stable*. In this case, the maximal (relatively) open set of states from which solutions converge to ξ^* is called the *basin* of ξ^* . If the basin of ξ^* contains $\text{int}(\Xi)$, we call ξ^* *almost globally asymptotically stable*; if it is Ξ itself, we call ξ^* *globally asymptotically stable*.

The C^1 function $L: O \rightarrow \mathbb{R}_+$ is a *strict Lyapunov function* for rest point $\xi^* \in O$ if $L^{-1}(0) = \{\xi^*\}$ and if its time derivative $\dot{L}(\xi) \equiv \nabla L(\xi)'V(\xi)$ is negative on $O \setminus \{\xi^*\}$. Standard results imply that if such a function exists, then ξ^* is asymptotically stable.¹⁵ If L is a strict Lyapunov function for ξ^* with domain $O = \Xi \setminus \{\xi^\dagger\}$ and ξ^\dagger is repelling, then ξ^* is almost globally asymptotically stable; if the domain is Ξ , then ξ^* is globally asymptotically stable.

EXAMPLE 1. As a preliminary, we consider $\text{BEP}(\tau, 1, \beta^{\min})$ dynamics for the centipede game of length 2. Since each player has two strategies, all test-set rules τ have revising agents test both of them. Focusing on the fractions of agents choosing to continue, we can express the dynamics as

$$\begin{aligned} \dot{x}_2 &= y_2 - x_2 \\ \dot{y}_2 &= x_2x_1 - y_2. \end{aligned} \tag{4}$$

By way of interpretation, a revising agent in population 1 chooses to continue if his opponent when he tests continue also continues. A revising agent in population 2 chooses to continue if her opponent continues when she tests continue, and her opponent stops when she tests stop.¹⁶

Writing $1 - x_2$ for x_1 in (4) and then solving for the zeros, we find that the unique rest point of (4) is the backward induction state: $x_2^\dagger = y_2^\dagger = 0$. Moreover, defining the function $L: [0, 1]^2 \rightarrow \mathbb{R}_+$ by $L(x_2, y_2) = \frac{1}{2}((x_2)^2 + (y_2)^2)$, we see that $L^{-1}(0) = \{\xi^\dagger\}$ and that

$$\dot{L}(x_2, y_2) = x_2\dot{x}_2 + y_2\dot{y}_2 = x_2y_2 - (x_2)^2 + y_2x_2 - y_2(x_2)^2 - (y_2)^2 = -(x_2 - y_2)^2 - y_2(x_2)^2,$$

which is nonpositive on $[0, 1]^2$ and equals zero only at the backward induction state. Since L is a strict Lyapunov function for ξ^\dagger on Ξ , state ξ^\dagger is globally asymptotically stable. ◇

In light of this example, our analyses to come will focus on centipede games of lengths $d \geq 3$.

¹⁵See, e.g., Sandholm (2010b, Appendix 7.B).

¹⁶Compare the discussion after (3) and Example 2 below.

3.1 Analytical results

As we know from [Observation 1](#), the backward induction state ξ^\dagger of the centipede game is a rest point of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic. Our first result shows that this rest point is always repelling.

PROPOSITION 1. *In centipede games of lengths $d \geq 3$, the backward induction state ξ^\dagger is repelling under the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic.*

The proof of [Proposition 1](#), which is presented in [Appendix A](#), is based on a somewhat nonstandard linearization argument. While we are directly concerned with the behavior of the BEP dynamics on the state space Ξ , it is useful to view (1) as defining dynamics throughout the affine hull $\text{aff}(\Xi) = \{(x, y) \in \mathbb{R}^{s^1+s^2} : \sum_{i \in S^1} x_i = \sum_{j \in S^2} y_j = 1\}$, which is then invariant under (1). Vectors of motion through $\text{aff}(\Xi)$ are elements of the tangent space $T\Xi = \{(z^1, z^2) \in \mathbb{R}^{s^1+s^2} : \sum_{i \in S^1} z_i^1 = \sum_{j \in S^2} z_j^2 = 0\}$. Note that $T\Xi$ is a subspace of $\mathbb{R}^{s^1+s^2}$ and that $\text{aff}(\Xi)$ is obtained from $T\Xi$ via translation: $\text{aff}(\Xi) = T\Xi + \xi^\dagger$.

A standard linearization argument is enough to prove that ξ^\dagger is unstable. Let the vector field $V : \text{aff}(\Xi) \rightarrow T\Xi$ be defined by the right-hand side of (1). To start the proof, we obtain an expression for the derivative matrix $DV(\xi^\dagger)$ that holds for any game length d . We then derive formulas for the d linearly independent eigenvectors of $DV(\xi^\dagger)$ in the subspace $T\Xi$ and for their corresponding eigenvalues. We find that $d - 1$ of the eigenvalues are negative and one is positive. The existence of the latter implies that ξ^\dagger is unstable.

To prove that ξ^\dagger is repelling, we show that the hyperplane through ξ^\dagger defined by the span of the set of $d - 1$ eigenvectors with negative eigenvalues supports the convex state space Ξ at state ξ^\dagger . Results from dynamical systems theory—specifically, the Hartman–Grobman and stable manifold theorems ([Perko \(2013, Sections 2.7 and 2.8\)](#))—then imply that in some neighborhood $O \subset \text{aff}(\Xi)$ of ξ^\dagger , the set of initial conditions from which solutions converge to ξ^\dagger is disjoint from $\Xi \setminus \{\xi^\dagger\}$, and that solutions from the remaining initial conditions eventually move away from ξ^\dagger .

The following example provides intuition for the instability of the backward induction state; the logic is similar in longer games and for other specifications of BEP dynamics.

EXAMPLE 2. In a centipede game of length $d = 4$, writing out display (3) shows that the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic is described by

$$\begin{aligned} \dot{x}_1 &= (y_1)^2 - x_1, & \dot{y}_1 &= (x_2 + x_3)(x_1 + x_2)^2 + (x_1)^3 - y_1 \\ \dot{x}_2 &= (y_2 + y_3)(y_1 + y_2) - x_2, & \dot{y}_2 &= x_3 + x_2x_1(x_1 + x_2) - y_2 \\ \dot{x}_3 &= y_3 + y_2y_1 - x_3, & \dot{y}_3 &= x_2(x_1)^2 + x_3(x_1 + x_2) - y_3. \end{aligned} \tag{5}$$

The linearization of this system at (x^\dagger, y^\dagger) has the positive eigenvalue 1 corresponding to eigenvector $(z^1, z^2) = ((-2, 1, 1), (-2, 1, 1))$ (this is (14) with $m \equiv d^1 = 2$ and $n \equiv d^2 = 2$). Thus at state $(x, y) = ((1 - 2\varepsilon, \varepsilon, \varepsilon), (1 - 2\varepsilon, \varepsilon, \varepsilon))$ with $\varepsilon > 0$ small, we have $(\dot{x}, \dot{y}) \approx ((-2\varepsilon, \varepsilon, \varepsilon), (-2\varepsilon, \varepsilon, \varepsilon))$.

To understand why the addition of agents in both populations playing the cooperative strategies 2 and 3 is self-reinforcing, we build on the discussion following (3). Consider, for instance, component y_3 , which represents the fraction of agents in population 2 who continue at both decision nodes. The last expression in (5) says that a revising population 2 agent switches to strategy 3 if (i) when testing strategy 3 she meets an opponent playing strategy 2, and when testing strategies 1 and 2 she meets opponents playing strategy 1, or (ii) when testing strategy 3 she meets an opponent playing strategy 3, and when testing strategy 2 she meets an opponent playing strategy 1 or 2. These events have total probability $\varepsilon(1 - \varepsilon) + \varepsilon(1 - 2\varepsilon)^2 \approx 2\varepsilon$. Since there are $y_3 = \varepsilon$ agents currently playing strategy 3, outflow from this strategy occurs at rate ε . Combining the inflow and outflow terms shows that $\dot{y}_3 \approx 2\varepsilon - \varepsilon = \varepsilon$. Analogous arguments explain the changes in the values of the other components of the state. \diamond

It may seem surprising that the play of a weakly dominated strategy—continuing by the last mover at the last decision node—is positively reinforced at an interior population state. This is possible because revising agents test each of their strategies against newly drawn opponents: as just described, a revising population 2 agent will choose to continue at both of her decision nodes if her opponent’s strategy when she tests strategy 3 is more cooperative than her opponents’ strategies when she tests her own less cooperative strategies.

Since the backward induction state is unstable, we next try to determine where in the state space the dynamics may converge. As a start, we prove that except at the rest point ξ^\dagger , motion from states on the boundary of the state space proceeds immediately into the interior of the state space.

PROPOSITION 2. *In centipede games of all lengths $d \geq 3$, solutions to the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic from every initial condition $\xi \in \text{bd}(\Xi) \setminus \{\xi^\dagger\}$ immediately enter $\text{int}(\Xi)$.*

The proof of Proposition 2, which is presented in Appendix B, starts with a simple differential inequality (Lemma 1) that lets us obtain explicit positive lower bounds on the use of any initially unused strategy i at times $t \in (0, T]$. The bounds are given in terms of the probabilities of test results that lead i to be chosen, and, thus, backing up one step, in terms of the usage levels of the opponents’ strategies occurring in those tests (16). With this preliminary result in hand, we prove inward motion from $\xi \neq \xi^\dagger$ by constructing a sequence that contains all unused strategies, and whose k th strategy could be chosen by a revising agent after a test result that includes only strategies that were initially in use or that appeared earlier in the sequence.

Together, Propositions 1 and 2 imply that in centipede games of length $d \geq 3$, the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic has an interior rest point. Our next result places strong bounds on the weights on the most cooperative strategies in any such rest point. For any population state $\xi = (x, y)$ and for $k \in \{0, \dots, d^1\}$, let $\bar{x}_{[k]} = x_{[0]} + \dots + x_{[k]}$ be the mass of population 1 agents who stop at their k th-to-last decision node or later. Likewise, for $\ell \in \{0, \dots, d^2\}$, let $\bar{y}_{[\ell]} = y_{[0]} + \dots + y_{[\ell]}$.

PROPOSITION 3. *Let $\xi = (x, y)$ be a rest point of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic in centipede other than the backward induction state ξ^\dagger .*

(i) *If d is even, so that player 2 moves last, then $\bar{y}_{[2]} > 0.9657$, $\bar{x}_{[2]} > 1 - 10^{-4}$, and $\bar{y}_{[3]} > 1 - 10^{-16}$.*

(ii) *If d is odd, so that player 1 moves last, then $\bar{x}_{[2]} > 0.9657$, $\bar{y}_{[2]} > 1 - 10^{-4}$, and $\bar{x}_{[3]} > 1 - 10^{-16}$.*

The proposition implies that the probability that play in a random match leads to the last five terminal nodes, namely $\bar{y}_{[2]}\bar{x}_{[2]}$, is at least 0.9656, and play is virtually guaranteed to reach the last eight terminal nodes.

PROOF OF PROPOSITION 3. We consider the case in which d is even. We start by introducing three inequalities that a rest point (x, y) must satisfy:

$$1 - \bar{x}_{[k]} \leq (1 - \bar{y}_{[k]})^{k+1} \tag{6}$$

$$1 - \bar{y}_{[k+1]} \leq (1 - \bar{x}_{[k]})^{k+2} \tag{7}$$

$$\bar{x}_{[k]} \geq (1 - (1 - \bar{y}_{[k+1]})^{k+1})(1 - \bar{y}_{[k+1]}). \tag{8}$$

A necessary condition for a population 1 agent to choose a strategy outside of $\{[k], \dots, [0]\}$ is that when testing strategies in this set, he is never matched with a population 2 agent playing a strategy in $\{[k], \dots, [0]\}$. This fact gives us inequality (6). Likewise, for a population 2 agent not to choose a strategy in $\{[k + 1], \dots, [0]\}$, it is necessary that when testing strategies in this set, she is never matched with a population 1 opponent playing a strategy in $\{[k], \dots, [0]\}$; this gives us inequality (7). Finally, inequality (8) follows from the observation that for a revising population 1 agent to choose a strategy in $\{[k], \dots, [0]\}$, it is enough that both (i) when playing at least one such strategy, her match opponent chooses a strategy in $\{[k + 1], \dots, [0]\}$, and (ii) when playing strategy $[k + 1]$, her opponent does not choose a strategy in $\{[k + 1], \dots, [0]\}$. Inequality (8) can also be written as

$$1 - \bar{x}_{[k]} \leq (1 - \bar{y}_{[k+1]})^{k+2} + \bar{y}_{[k+1]}. \tag{9}$$

Substituting inequality (9) into inequality (7) and rearranging yields

$$((1 - \bar{y}_{[k+1]})^{k+2} + \bar{y}_{[k+1]})^{k+2} + \bar{y}_{[k+1]} - 1 \geq 0. \tag{10}$$

When $k = 0$, the left-hand side of (10) is a polynomial whose real roots are 0 and r , with r close to but greater than .4301, and that is negative on $(0, r)$ and positive on (r, ∞) . (This r is the lone real root of the polynomial $z^3 - 2z^2 + 3z - 1$.) Thus, any rest point (x, y) with $\bar{y}_{[1]} \neq 0$ satisfies $\bar{y}_{[1]} \geq r > 0.4301$. By Proposition 2, this requirement holds for every rest point besides ξ^\dagger .

Now, applying (6) and (7) sequentially yields the inequalities

$$\begin{aligned}
 \bar{x}_{[1]} &\geq 1 - (1 - \bar{y}_{[1]})^2 \geq 1 - (1 - r)^2 > 0.6752 \\
 \bar{y}_{[2]} &\geq 1 - (1 - \bar{x}_{[1]})^3 \geq 1 - (1 - r)^6 > 0.9657 \\
 \bar{x}_{[2]} &\geq 1 - (1 - \bar{y}_{[2]})^3 \geq 1 - (1 - r)^{18} > 1 - 10^{-4} \\
 \bar{y}_{[3]} &\geq 1 - (1 - \bar{x}_{[2]})^4 \geq 1 - (1 - r)^{72} > 1 - 10^{-16}.
 \end{aligned}
 \tag{11}$$

□

Part of the intuition behind the proof of Proposition 3 is straightforward. Consider an interior rest point $\xi = (x, y)$ of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic when d is even. Inequality (6) says that if $1 - \bar{y}_{[k]}$ is small (i.e., if few population 2 players stop before $[k]$), then $1 - \bar{x}_{[k]}$ is even smaller (i.e., few population 1 players stop before $[k]$), since some test of a more cooperative strategy will very likely lead to a high payoff. Likewise, by (7), if $1 - \bar{x}_{[k]}$ is small, then $1 - \bar{y}_{[k+1]}$ is smaller still. These inequalities imply that any lower bound on $\bar{y}_{[k]}$ will quickly propagate into much stronger lower bounds on $\bar{x}_{[k]}$, $\bar{y}_{[k+1]}$, $\bar{x}_{[k+1]}$, and so on (see (11)). Initiating this chain of reasoning requires a less obvious step: we combine inequality (7) with a weak bound (9) on $1 - \bar{x}_{[k]}$ that depends on $\bar{y}_{[k+1]}$ alone. This combination gives us the initial inequality $\bar{y}_{[1]} \geq 0.4301$, which in turn leads to the strong bounds stated in the proposition.

3.2 Results based on exact computations

Proposition 3 places strong lower bounds on the degree of cooperation arising at any rest point other than the unstable backward induction state. To gain a more precise understanding of the form and stability of such rest points, we turn to exact computations. Because the dynamic (3) is a system of polynomials with rational coefficients, its zeros can in principle be found by computing a Gröbner basis for the system. The Gröbner basis is a new system of equations that has the same zeros as the original system, but can be solved by backward substitution. Once the Gröbner basis has been obtained, polynomial factoring and root finding algorithms can be used to identify its zeros and, hence, the zeros of the original system. Applying these techniques, which are described in detail in Appendix C, leads to part (i) of the following result.

PROPOSITION 4. *In centipede games of lengths $3 \leq d \leq 6$, the following statements hold:*

- (i) *The $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic has exactly two rest points: ξ^\dagger , and $\xi^* = \xi^*(d) \in \text{int}(\Xi)$.*
- (ii) *The rest point ξ^* is asymptotically stable.*

Exact solutions can only be obtained for games of length at most 6 because of the computational demands of computing the Gröbner bases. Two indications of these demands are that when $d = 6$, the leading (univariate) polynomial from the Gröbner basis is of degree 221 and a coefficient of one of the polynomials in the basis has 13,278 digits.

Table 1 reports the approximate values of the interior rest points $\xi^* = \xi^*(d)$, referring to strategies using the last-to-first notation $[k]$ introduced in Section 2.4. Evidently, the

	Population p				Population q			
	[3]	[2]	[1]	[0]	[3]	[2]	[1]	[0]
$d = 3$			0.618034	0.381966		0.381966	0.381966	0.236068
$d = 4$		0.113625	0.501712	0.384663		0.337084	0.419741	0.243175
$d = 5$		0.113493	0.501849	0.384658	0.001462	0.335672	0.419706	0.243160
$d = 6$	3.12×10^{-9}	0.113493	0.501849	0.384658	0.001462	0.335672	0.419706	0.243160

Note: The p denotes the owner of the penultimate decision node; the q denotes the owner of the last decision node.

TABLE 1. “Exact” interior rest points $\xi^* = \xi^*(d)$ of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic.

$d = 3$	-1 ± 0.3820	-1	
$d = 4$	$-1.1411 \pm 0.3277i$	$-0.8589 \pm 0.3277i$	
$d = 5$	$-1.1355 \pm 0.3284i$	$-0.8645 \pm 0.3284i$	-1
$d = 6$	$-1.1355 \pm 0.3284i$	$-0.8645 \pm 0.3284i$	$-1 \pm 9.74 \times 10^{-5}i$

TABLE 2. Eigenvalues of the derivative matrices $DV(\xi^*)$ of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic.

masses on each strategy are nearly identical for games of lengths 4, 5, and 6, with nearly all of the weight in both populations being placed on continuing to the end, stopping at the last node, or stopping at the penultimate node.

In principle, it is possible to prove the local stability of the rest points $\xi^* = \xi^*(d)$ using linearization. But since the components of ξ^* are algebraic numbers, computing the eigenvalues of $DV(\xi^*)$ requires finding the exact roots of a polynomial with algebraic coefficients, a computationally intensive problem. Fortunately, we can prove local stability without doing so. Instead, we compute the eigenvalues of the matrix $DV(\xi)$, where ξ is a rational point that is very close to ξ^* , showing that these eigenvalues all have negative real part. Proposition 6 in Appendix D establishes an upper bound on the distances between the eigenvalues of $DV(\xi)$ and $DV(\xi^*)$. Importantly, this bound can be evaluated without having to compute the roots of a polynomial with algebraic coefficients or to invert a matrix with algebraic components, as both of these operations quickly become computationally infeasible. Combining these steps allows us to conclude that the eigenvalues of $DV(\xi^*)$ also have negative real part. For a detailed presentation of this argument, see Appendix D.

The approximate eigenvalues of $DV(\xi^*)$ are reported in Table 2. Note that the eigenvalues for games of length 5 and 6 are nearly identical, with the replacement of an eigenvalue of -1 by a pair of complex eigenvalues that are very close to -1 .

3.3 Numerical results

Because exact methods allow us to determine only the rest points of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic in centipede games of lengths $d \leq 6$, we use numerical methods to study games of lengths 7–20. We know from Proposition 3 that at any rest point besides the backward induction state ξ^\dagger , the weight on strategies that stop before either player’s




FIGURE 2. The stable rest point $\xi^* = \xi^*(d)$ of centipede under the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic for game lengths $d = 3, \dots, 10$ and $d = 20$. Stacked bars, from the bottom to the top, represent weights on strategy [0] (continue at all decision nodes), [1] (stop at the last node), [2] (stop at the second-to-last node), etc. The dashed line separates exact ($d \leq 6$) and numerical ($d \geq 7$) results.

third-to-last node is very small. This suggests that the presence of earlier nodes should have little bearing on how the game is played.

Our numerical analysis suggests that for game lengths $7 \leq d \leq 20$, there are exactly two rest points: the backward induction state ξ^\dagger and an interior rest point $\xi^* = \xi^*(d)$. As Figure 2 illustrates, the form of the interior rest point follows the pattern from Table 1: regardless of the length of the game, nearly all of the mass is placed on each population’s three most cooperative strategies, and the weights on these strategies are essentially independent of the length of the game. Precise numerical estimates of these rest points are provided in Appendix III available in a supplementary file on the journal website, <http://econtheory.org/supp3565/supplement.pdf>, as are numerical estimates of the eigenvalues of the derivative matrices $DV(\xi^*)$. The latter are essentially identical to those presented in Table 2 for $d = 6$, with the addition of an eigenvalue of ≈ -1 for each additional decision node.

These numerical results suggest that the conclusions about rest points established analytically for games of lengths $d \leq 6$ continue to hold for longer games: there are always exactly two rest points: the backward induction state ξ^\dagger and a stable interior rest point ξ^* whose form barely varies with the length of the game.

The facts that the vertex ξ^\dagger is repelling, the interior rest point $\xi^* = \xi^*(d)$ is attracting, and these are the only two rest points give us a strong reason to suspect that state ξ^* attracts all solutions of the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic other than the stationary solution at ξ^\dagger .¹⁷ To argue that ξ^* is almost globally stable, we introduce the candidate Lyapunov

¹⁷For there to be other solutions that did not converge to ξ^* without the dynamics having another rest point, the flow of the dynamics would need to have very special topological properties. For instance, in a two-dimensional setting, this could occur if ξ^* were contained in a pair of concentric closed orbits, where the inner orbit is repelling and the outer orbit is attracting.

function

$$L(x, y) = \sum_{i=2}^{s_1} (x_i - x_i^*)^2 + \sum_{j=2}^{s_2} (y_j - y_j^*)^2.$$

In words, $L(x, y)$ is the squared Euclidean distance of (x, y) from (x^*, y^*) if the points in the state space Ξ are represented in \mathbb{R}^d by omitting the first components of x and y .

The Gröbner basis techniques used in Section 3.2 are not suitable for establishing that L is a Lyapunov function. For the centipede game of length $d = 3$, we are able to verify that L is a Lyapunov function using an algorithm from real algebraic geometry called *cylindrical algebraic decomposition* (Collins (1975)). However, exact implementations of this algorithm fail to terminate in longer games.

We therefore verify numerically that L is a Lyapunov function. For games of lengths 4–20, we chose one billion (10^9) points from the state space Ξ uniformly at random, and evaluated a floating-point approximation of \dot{L} at each point. In all instances, the approximate version of \dot{L} evaluated to a negative number. This numerical procedure covers the state space fairly thoroughly for the game lengths we consider,¹⁸ and so provides strong numerical evidence that the interior rest point ξ^* is an almost global attractor.

3.4 Other specifications of the dynamics

To test the robustness of the preceding results, we repeat the analyses for other specifications of $\text{BEP}(\tau, 1, \beta)$ dynamics. In addition to the test-all rule τ^{all} , we also studied a test-set rule under which the revising agent considers only his current strategy and one other strategy (τ^{two}), as well as a rule under which the revising agent considers only his current strategy and one adjacent strategy (τ^{adj}). The qualitative behavior under these test-set rules is similar to that under τ^{all} . The differences worth mentioning are that stable play is concentrated on a larger number of strategies (e.g., nine strategies in total have mass of at least 0.01 under $\text{BEP}(\tau^{\text{two}}, 1, \beta^{\text{min}})$) and that the rate of decay of the weights on strategies that stop earlier is not as severe as under τ^{all} . The intuition here is simple. Under test-all, a revising agent will try out all of his most cooperative strategies, providing many opportunities for some such strategy to perform best; if instead only two strategies are tested at a time, the selective pressure against less cooperative strategies is weaker.

In addition, we also considered alternative tie-breaking rules: *stick/min-if-tie*, which chooses the agent's current strategy if it is optimal and chooses the lowest-numbered strategy otherwise, and *uniform-if-tie*, which randomizes uniformly among the optimal strategies (as in Osborne and Rubinstein (1998) and Sethi (2000)). As we noted in Section 2.5, uniform tie-breaking implies that the backward induction state ξ^\dagger is not a rest point, rendering a stability analysis of this rest point unnecessary. In other respects, alternate choices of tie-breaking rules have little qualitative impact on behavior.

¹⁸By a standard combinatoric formula, the number of states in a grid in $\Xi = X \times Y$ with mesh $\frac{1}{m}$ is $\binom{m+s^1-1}{m} \binom{m+s^2-1}{m}$. Applying this formula shows for a game of length 10 that 10^9 is between the numbers of states in grids in Ξ of meshes $\frac{1}{17}$ (since $\binom{22}{17}^2 = 693,479,556$) and $\frac{1}{18}$ (since $\binom{23}{18}^2 = 1,132,255,201$). For a game of length 15, the comparable meshes are $\frac{1}{10}$ and $\frac{1}{11}$, and for length 20 are $\frac{1}{7}$ and $\frac{1}{8}$.

In summary, the results presented in previous sections are highly robust to alternative specifications of the dynamics.

4. LARGER NUMBERS OF TRIALS

The analysis thus far has focused on cases in which agents test each strategy in their test sets exactly once. We now examine aggregate behavior when each strategy is subject to larger numbers of trials κ , focusing on $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamics.

4.1 *Instability and stability of the backward induction state*

Proposition 1 shows that the backward induction state ξ^\dagger is a repellor under the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic with $\kappa = 1$. The following proposition shows that ξ^\dagger remains unstable as long as the number of trials κ is less than the length of the game d and then becomes stable for larger numbers of trials. The statement is complicated slightly by the dependence of the crossover point on whether d is even or odd.

PROPOSITION 5. *Under the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic in the centipede game of length d , the backward induction state ξ^\dagger is unstable if $\kappa \leq 2\lfloor \frac{d}{2} \rfloor$ and is asymptotically stable otherwise.*

Like those of the earlier stability analyses, the proof of **Proposition 5**, which is presented in **Appendix E**, is based on linearization. The key observation is that linearizations of BEP dynamics around pure rest points are driven by match results in which exactly one out of all κs^p match partners plays a strategy different from the equilibrium strategy. We show that if $\kappa \leq d - 1$ (for d even) or $\kappa \leq d$ (for d odd), there is enough sensitivity to perturbations of the state to ensure the existence of an unstable manifold through ξ^\dagger ; however, unlike in the $\kappa = 1$ case, ξ^\dagger need not be a repellor. Conversely, if these inequalities are violated, strategy 1 $\in S^1$ earns the highest total payoff after any matching with exactly one discrepant opponent. This insensitivity of population 1's behavior to small changes in population 2's behavior ensures local stability.

So as to assess the practical relevance of the stability of the backward induction state for larger numbers of trials, we use numerical analysis to estimate the basin of attraction of ξ^\dagger and to determine the position of the interior saddle point of the dynamics, which lies on the manifold separating the basin of ξ^\dagger from the basin of the main attractor. We focus for tractability on games of length $d = 4$ and numbers of trials $\kappa \leq 100$. Details of these analyses are presented in Appendices IV and V in the Supplemental Material.

We have two main findings from this numerical analysis. First, the basin of attraction is always minuscule, with volumes always smaller than 0.01% of the total volume of the state space Ξ . Second, ξ^\dagger is almost completely nonrobust to changes in behavior in population 1. Evidence for this lies in the position of the saddle points, which have more than 99.8% of population 1 agents choosing strategy 1, indicating that changes in the behavior of 0.2% of population 1 agents are enough to disrupt the stability of ξ^\dagger . Thus, the exact stability analysis of the backward induction state for larger numbers of trials is undercut by a thorough numerical analysis of the dynamics in the vicinity of that state.

4.2 Persistence of the stable interior rest point

When agents test their strategies thoroughly, the distributions of opponents' choices they face when testing each strategy will come to resemble the current distribution of play in the opposing population. Since agents choose the strategy whose total payoff during testing was highest, this suggests that the rest points of the resulting dynamics should approximate Nash equilibria. Indeed, when agents possess exact information, so that play adjusts according to the exact best response dynamic (Gilboa and Matsui (1991), Hofbauer (1995)), the results of Xu (2016) imply that every solution trajectory converges to the set of Nash equilibria; in centipede, all Nash equilibria entail all population 1 agents stopping immediately.

While the intuition suggested above is correct for large enough numbers of trials, it is nevertheless the case that stable cooperative behavior can persist when the number of trials of each strategy is substantial. To illustrate this, we consider play in the centipede game of length $d = 4$ under the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic. Figures 3 and 4 present the stable rest points of this dynamic for numbers of trials κ up to 50, which we computed using numerical methods. While increasing the number of trials shifts mass toward uncooperative strategies, it is clear from the figures that this shifting takes place gradually: even with rather thorough testing, significant levels of cooperation are still maintained. We note as well that the fraction of population 2 agents who play the weakly dominated strategy [0] (always continue) becomes fixed between 7% and 6.5% once $\kappa \geq 15$, even as the fraction of population 1 agents who play strategy [0] remains far from 0 (specifically, between 28% and 18%).

While surprising at first glance, these facts can be explained by considering both the expectations and the *dispersions* in the payoffs obtained through repeated trials of each strategy. As an illustration, consider the stable rest point when $\kappa = 32$, namely $\xi^* = (x^*, y^*) \approx ((0.2140, 0.5738, 0.2122), (0.6333, 0.3010, 0.0657))$. Let Π_j be a random variable that represents the payoff obtained by a population 2 agent who plays strategy j in a single random match at this state. By (2) (or Figure 1), the expected payoffs to this agent's three strategies are

$$\mathbb{E}(\Pi_1) = (0, 3, 3) \cdot x^* = 2.3580,$$

$$\mathbb{E}(\Pi_2) = (0, 2, 5) \cdot x^* = 2.2086,$$

$$\mathbb{E}(\Pi_3) = (0, 2, 4) \cdot x^* = 1.9964.$$


From this we anticipate that the strategy weights in population 2 satisfy $y_1^* > y_2^* > y_3^*$.

To explain why these weights take the values they do, we also need to know how dispersed are the payoffs from testing each strategy. We thus compute the variances of the single-test payoffs Π_j :


$$\text{Var}(\Pi_1) = 1.5138,$$

$$\text{Var}(\Pi_2) = 2.7223,$$

$$\text{Var}(\Pi_3) = 1.7048.$$



(i) population 1



(ii) population 2

FIGURE 3. The stable interior rest point of the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic in the centipede game of length $d = 4$, $\kappa = 1, \dots, 50$. Stacked bars, from the bottom to the top, represent weights on strategies [0], [1], and [2].

Using these calculations and the central limit theorem, we find that the difference between the average payoffs from 32 tests of strategy 3 and 32 tests of strategy 2 is approximately normally distributed with mean $\mathbb{E}(\Pi_3) - \mathbb{E}(\Pi_2) = -0.2122$ and standard deviation $\sqrt{(\text{Var}(\Pi_3) + \text{Var}(\Pi_2))/32} \approx 0.3720$. The latter statistic is commensurate with the former. Thus the weakly dominated strategy 3 yields a higher total payoff than the dominating strategy 2 with approximate probability $\mathbb{P}(Z \geq 0.57) \approx 0.28$ and so is not a rare event. Likewise, evaluating the appropriate multivariate normal integrals shows that the probabilities of strategies 1, 2, and 3 yielding the highest total payoff are approximately 0.61, 0.32, and 0.07, figures which accord fairly well with the components of y^* .

As the number of trials κ becomes larger, greater averaging reduces the variation in each strategy's payoffs per trial. At the same time, increasing κ increases the weight x_1^* on stopping immediately at the expense of population 1's other two strategies, hence reducing the differences in the expected payoffs of population 2's strategies. This explains




FIGURE 4. The stable interior rest point in centipede of length $d = 4$ under $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamics for $\kappa = 1, \dots, 34$ trials of each tested strategy. Lighter shading corresponds to larger numbers of trials. Dashed lines represent boundaries of best response regions.

why the strategy weights in population 2 do not vary very much as κ increases and why the weight on the weakly dominated strategy hardly varies at all.

4.3 Convergence to cycles

Figure 3 does not record rest points for certain numbers of trials above 34. For these values of κ , the population state does not converge to a rest point. Instead, our numerical analyses indicate that for all κ with empty entries in Figure 3 and all κ between 51 and 100, the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamic converges to a periodic orbit. Figure 5 presents the cycles under the $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamics for $\kappa = 50, 100$, and 200 . In all three cases, we observe substantial levels of cooperative play in population 1 over the course of the cycle, with the fraction of the population choosing to continue at the initial node varying between 0.50 and 0.83 for $\kappa = 50$, between 0.28 and 0.70 for $\kappa = 100$, and between 0.16 and 0.45 for $\kappa = 200$. These examples illustrate that cooperative behavior can persist even when agents have substantial amounts of information about opponents' play.

From a methodological point of view, the existence of attracting limit cycles under BEP dynamics suggests that solution concepts like $S(k)$ equilibrium and logit equilibrium that are motivated as steady states of dynamic disequilibrium processes should be applied with some caution. Existence results for such solution concepts can generally be proved by appeals to suitable fixed point theorems. But the fact that static solutions exist need not imply that any are stable, and it may happen that no static solution provides a good prediction of the behavior of the underlying dynamic process.

5. CONCLUSION

In this paper, we introduce a class of game dynamics built on natural assumptions about the information agents obtain when revising, and we show that these dynamics lead to cooperative behavior in the centipede game. One key feature of the agents' revision process is that conditional on the current population state, the experienced payoffs to each




FIGURE 5. Stable cycles in centipede of length $d = 4$ under $\text{BEP}(\tau^{\text{all}}, \kappa, \beta^{\text{min}})$ dynamics for $\kappa = 50, 100,$ and 200 . Lighter shading represents faster motion. The small circles represent the unstable interior rest points. For $\kappa = 50$ and 100 , shapes synchronize positions along the cycle.

strategy are independent of one another. This allows cooperative strategies with sub-optimal expected payoffs to be played with nonnegligible probabilities, even when the testing of each strategy involves substantial numbers of trials. The use of any such strat-

egy increases the expected payoffs of other cooperative strategies, creating a virtuous circle that sustains cooperative play.

APPENDIX A: PROOF OF PROPOSITION 1

A.1 Generalities

Letting $s = s^1 + s^2$, we denote the tangent space of the state space $\Xi = X \times Y$ by $T\Xi = TX \times TY = \{(z^1, z^2)' \in \mathbb{R}^s : \sum_{i \in S^1} z_i^1 = 0 \text{ and } \sum_{j \in S^2} z_j^2 = 0\}$ and we denote the affine hull of Ξ by $\text{aff}(\Xi) = T\Xi + \xi^\dagger$. Writing our dynamics as

$$\dot{\xi} = V(\xi), \tag{D}$$

we have $V: \text{aff}(\Xi) \rightarrow T\Xi$, and so $DV(\xi)z \in T\Xi$ for all $\xi \in \Xi$ and $z \in T\Xi$. We can thus view $DV(\xi)$ as a linear map from $T\Xi$ to itself, and the behavior of the dynamics in the neighborhood of a rest point is determined by the eigenvalues and eigenvectors of this linear map. The latter are obtained by computing the eigenvalues and eigenvectors of the product matrix $\Phi DV(\xi)\Phi$, where $\mathcal{V}: \mathbb{R}^s \rightarrow \mathbb{R}^s$ is the natural extension of V to \mathbb{R}^s , and Φ is the orthogonal projection of \mathbb{R}^s onto $T\Xi$, i.e., the block diagonal matrix with diagonal blocks $I - \frac{1}{s^1}\mathbf{1}\mathbf{1}' \in \mathbb{R}^{s^1 \times s^1}$ and $I - \frac{1}{s^2}\mathbf{1}\mathbf{1}' \in \mathbb{R}^{s^2 \times s^2}$, where $\mathbf{1} = (1, \dots, 1)'$. Since V maps Ξ into $T\Xi$, the projection is only needed when there are eigenspaces of $DV(\xi)$ that intersect both the set $T\Xi$ and its complement.

We prove that the backward induction state ξ^\dagger is a repellor using the following argument. Computing the eigenvalues and eigenvectors of $DV(\xi^\dagger)$ as described above, we find that ξ^\dagger is a *hyperbolic* rest point, meaning that all of the eigenvalues have a nonzero real part.

The linearization of the dynamic (D) at rest point ξ^\dagger is the linear differential equation

$$\dot{z} = DV(\xi^\dagger)z \tag{L}$$

on $T\Xi$. The *stable subspace* $E^s \subseteq T\Xi$ of (L) is the span of the real and imaginary parts of the eigenvectors and generalized eigenvectors of $DV(\xi^\dagger)$ corresponding to eigenvalues with a negative real part. The *unstable subspace* $E^u \subseteq T\Xi$ of (L) is defined analogously. The basic theory of linear differential equations implies that solutions to (L) on E^s converge to the origin at an exponential rate, that solutions to (L) on E^u diverge from the origin at an exponential rate, and that the remaining solutions approach E^u and then diverge from the origin at an exponential rate.

Let $A^s = E^s + \xi^\dagger$ and $A^u = E^u + \xi^\dagger$ denote the affine spaces that are parallel to E^s and E^u and that pass through ξ^\dagger . In Appendix A.2, we prove that under the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic, the dimensions of E^s and E^u are $d - 1$ and 1 , and that A^s is a supporting hyperplane to Ξ at ξ^\dagger .

Combining these facts with fundamental results from dynamical systems theory lets us complete the proof that ξ^\dagger is a repellor. By the Hartman–Grobman theorem (Perko (2013, Section 2.8)), there is a homeomorphism h between a neighborhood of ξ^\dagger in $\text{aff}(\Xi)$ and a neighborhood of $\mathbf{0}$ in $T\Xi$ that maps solutions of (D) to solutions of (L). By

the stable manifold theorem (Perko (2013, Section 2.7)), there is an invariant *stable manifold* $M^s \subset \text{aff}(\Xi)$ of dimension $\dim(E^s) = d - 1$ that is tangent to A^s at ξ^\dagger such that solutions to (D) in M^s converge to ξ^\dagger at an exponential rate. Combining these results shows that there is a neighborhood $O \subset \text{aff}(\Xi)$ of ξ^\dagger with these properties: $O \cap \Xi \cap M^s = \{\xi^\dagger\}$; the initial conditions in O from which solutions converge exponentially quickly to ξ^\dagger are those in $O \cap M^s$; and solutions from initial conditions in $(O \cap \Xi) \setminus \{\xi^\dagger\}$ eventually move away from ξ^\dagger . Thus the properties stated in the previous paragraph imply that state ξ^\dagger is a repeller of the dynamic (D) on Ξ .

A.2 Computation of eigenvalues and eigenvectors

Starting from formula (3) and using the notations $m \equiv d^1 = s^1 - 1$ and $n \equiv d^2 = s^2 - 1$, it is easy to verify that under the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic,

$$D^{\mathcal{V}}(\xi^\dagger) = \begin{pmatrix} -1 & 0 & \cdots & 0 & m+1 & 1 & \cdots & 1 \\ 0 & -1 & \ddots & \vdots & 0 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & 0 & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & -1 & 0 & 1 & \cdots & 1 \\ \hline n+1 & 1 & \cdots & 1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 1 & 0 & -1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 1 & \cdots & 1 & 0 & \cdots & 0 & -1 \end{pmatrix}.$$

Write $\delta^i \in \mathbb{R}^s$ and $\varepsilon^j \in \mathbb{R}^s$ for the standard basis vectors corresponding to strategies $i \in S^1$ and $j \in S^2$, respectively. For $d \geq 3$, the eigenvalues of $DV(\xi^\dagger)$ with respect to $T\Xi$ and the bases for their eigenspaces are

$$-1, \quad \{\delta^2 - \delta^i : i \in \{3, \dots, m+1\}\} \cup \{\varepsilon^2 - \varepsilon^j : j \in \{3, \dots, n+1\}\} \tag{12}$$

$$-1 - \sqrt{mn}, \quad \{(\sqrt{mn}, -\sqrt{n/m}, \dots, -\sqrt{n/m} | -n, 1, \dots, 1)'\} \tag{13}$$

$$-1 + \sqrt{mn}, \quad \{(-\sqrt{mn}, \sqrt{n/m}, \dots, \sqrt{n/m} | -n, 1, \dots, 1)'\}. \tag{14}$$

The eigenvectors in (12) and (13) span the stable subspace E^s of the linear equation (L). The normal vector to E^s is

$$z^\perp = \left(-\frac{n+1}{m+1} \sqrt{\frac{m}{n}}, \frac{n+1}{(m+1)\sqrt{mn}}, \dots, \frac{n+1}{(m+1)\sqrt{mn}} \mid -1, \frac{1}{n}, \dots, \frac{1}{n} \right)'.$$

This vector satisfies

$$(z^\perp)'(\delta^i - \delta^1) = \frac{n+1}{\sqrt{mn}} > 0 \quad \text{for } i \in S^1 \setminus \{1\}$$

$$(z^\perp)'(\varepsilon^j - \varepsilon^1) = \frac{n+1}{n} > 0 \quad \text{for } j \in S^2 \setminus \{1\}.$$

The collection of vectors $\{\delta^i - \delta^1 : i \in S^1\} \cup \{\varepsilon^j - \varepsilon^1 : j \in S^2\}$ describes the motions along all edges of the convex set Ξ emanating from state ξ^\dagger . Thus, the fact that their inner products with z^\perp are all positive implies that the translation of E^s to ξ^\dagger is a hyperplane that supports Ξ at ξ^\dagger . Since the remaining eigenvalue, from (14), is positive, the arguments from the start of the section allow us to conclude that ξ^\dagger is a repellor.

APPENDIX B: PROOF OF PROPOSITION 2

The following differential inequality allows us to obtain simple lower bounds on the use of initially unused strategies. In all cases in which we apply the lemma, $v(0) = 0$.

LEMMA 1. *Let $v : [0, T] \rightarrow \mathbb{R}_+$ satisfy $\dot{v}(t) \geq a(t) - v(t)$ for some $a : [0, T] \rightarrow \mathbb{R}_+$. Then*

$$v(t) \geq e^{-t} \left(v(0) + \int_0^t e^s a(s) ds \right) \quad \text{for all } t \in [0, T]. \tag{15}$$

PROOF. Clearly $v(t) = v(0) + \int_0^t \dot{v}(s) ds \geq v(0) + \int_0^t (a(s) - v(s)) ds$. The final expression is the time t value of the solution to the differential equation $\dot{v}(s) + v(s) = a(s)$ with initial condition $v(0)$. Using the integrating factor e^s to solve this equation yields the right-hand side of (15). □

For the analysis to come, it will be convenient to work with the set $S = S^1 \cup S^2$ of all strategies from both populations and to drop population superscripts from notation related to the state; for instance, writing ξ_i rather than ξ_i^p .

We use Lemma 1 to prove inward motion from the boundary under BEP dynamics in the following way. Write $\dot{\xi}_i = r_i(\xi) - \xi_i$, where $r_i(\xi)$ is the polynomial appearing in the formula (1) for the BEP($\tau, 1, \beta$) dynamic. Let $\{\xi(t)\}_{t \geq 0}$ be the solution to the dynamic with initial condition $\xi(0)$. Let $S_0 = \text{supp}(\xi(0))$ and $Q = \min\{\xi_h(0) : h \in S_0\}$, and, finally, let $S_1 = \{i \in S \setminus S_0 : r_i(\xi(0)) > 0\}$ and $R = \frac{1}{2} \min\{r_k(\xi(0)) : r_k(\xi(0)) > 0\}$.

By the continuity of (1), there is a neighborhood $O \subset \Xi$ of $\xi(0)$ such that every $\chi \in O$ satisfies $\chi_h > Q$ for all $h \in S_0$ and $\dot{\chi}_i \geq R$ for all $i \in S_1$, and since (1) is smooth, there is a time $T > 0$ such that $\xi(t) \in O$ for all $t \in [0, T]$. Thus, applying Lemma 1 shows that

$$\xi_i(t) \geq R(1 - e^{-t}) \quad \text{for all } t \in [0, T] \text{ and } i \in S_1.$$

Now let S_2 be the set of $j \notin S_0 \cup S_1$ for which there is a term of polynomial r_j whose factors all correspond to elements of S_0 or S_1 . If this term has a factors in S_0 , b factors in S_1 , and coefficient c , then the foregoing claims and Lemma 1 imply that

$$\xi_j(t) \geq cQ^a e^{-t} \int_0^t e^s (R(1 - e^{-s}))^b ds \quad \text{for all } t \in [0, T]. \tag{16}$$

Proceeding sequentially, we can obtain positive lower bounds on the use of any strategy for times $t \in (0, T]$ by considering as-yet-unconsidered strategies k whose polynomials r_k have a term whose factors all correspond to strategies for which lower bounds

have already been obtained. Below, we prove that solutions to the $\text{BEP}(\tau^{\text{all}}, 1, \beta^{\text{min}})$ dynamic from states $\xi(0) \neq \xi^\dagger$ immediately enter $\text{int}(\Xi)$ by showing that the strategies in $\mathcal{S} \setminus \mathcal{S}_0$ can be considered in a sequence that satisfies the property just stated.

To proceed, we use the notations $i^{[1]}$ and $i^{[2]}$ to denote the i th strategies of players 1 and 2. We also introduce the linear order $<$ on \mathcal{S} defined by $1^{[1]} < 1^{[2]} < 2^{[1]} < 2^{[2]} < 3^{[1]} < \dots$, which arranges the strategies according to how early they stop play in centipede.

PROOF OF PROPOSITION 2. Fix an initial condition $\xi(0) \neq \xi^\dagger$. We can sequentially add all strategies in $\mathcal{S} \setminus \mathcal{S}_0$ in accordance with the property above as follows:

(i) We add the strategies $\{i \in \mathcal{S} \setminus \mathcal{S}_0 : i < \max \mathcal{S}_0\}$ in decreasing order. At the point that i has been added, i 's successor h has already been added, and strategy i is the unique best response when the revising agent tests all strategies against opponents playing h . Let \mathcal{S}_1 denote the set of strategies added during this stage. The assumption that $\xi(0) \neq \xi^\dagger$ implies that $\mathcal{S}_0 \cup \mathcal{S}_1$ contains $1^{[1]}$, $1^{[2]}$, and $2^{[1]}$.

(ii) We add the strategies $j \in \mathcal{S}^2 \setminus (\mathcal{S}_0 \cup \mathcal{S}_1)$. We can do so because j is the unique best response when it is tested against $2^{[1]}$ and all other strategies are tested against $1^{[1]}$.

(iii) We add the strategies $k \in \mathcal{S}^1 \setminus (\mathcal{S}_0 \cup \mathcal{S}_1)$. We can do so because k is the unique best response when it is tested against $2^{[2]}$ and other strategies are tested against $1^{[2]}$. \square

APPENDIX C: EXACT SOLUTIONS OF SYSTEMS OF POLYNOMIAL EQUATIONS

In this section, we describe the algebraic tools that we use to compute the exact rest points of BEP dynamics in centipede games (Proposition 4(i)).

C.1 Gröbner bases

Let $\mathbb{Q}[z_1, \dots, z_n]$ ($\mathbb{Q}[z]$ for short) denote the collection (more formally, the *ring*) of polynomials in the variables z_1, \dots, z_n with rational coefficients. Let $F = \{f_1, \dots, f_m\} \subset \mathbb{Q}[z]$ be a set of such polynomials. Let Z be a subset of \mathbb{R}^n and consider the problem of finding the set of points $z^* \in Z$ that are zeros of all polynomials in F .

To do so, it is convenient first to consider finding all zeros in \mathbb{C}^n of the polynomials in F . In this case, the set of interest,

$$\mathbf{V}(f_1, \dots, f_m) = \{z^* \in \mathbb{C}^n : f_j(z^*) = 0 \text{ for all } 1 \leq j \leq m\}, \tag{17}$$

is called the *variety* (or *algebraic set*) *generated by* f_1, \dots, f_m . To characterize (17), it is useful to introduce the *ideal generated by* f_1, \dots, f_m :

$$\langle f_1, \dots, f_m \rangle = \left\{ \sum_{j=1}^m h_j f_j : h_j \in \mathbb{C}[z] \text{ for all } 1 \leq j \leq m \right\}. \tag{18}$$

Thus, the ideal (18) is the set of linear combinations of the polynomials f_1, \dots, f_m , where the coefficients on each are themselves polynomials in $\mathbb{C}[z]$. It is easy to verify that any other collection of polynomials in $\mathbb{C}[z]$ whose linear combinations generate the ideal (18)—that is, any other *basis* for the ideal—also generates the variety (17).

For our purposes, the most useful basis for the ideal (18) is the *reduced lex-order Gröbner basis*, which we denote by $G \subset \mathbb{Q}[z]$. This basis, which contains no superfluous polynomials and is uniquely determined by its ideal and the ordering of the variables, has this convenient property: it consists of polynomials in z_n only, polynomials in z_n and z_{n-1} only, polynomials in z_n, z_{n-1} , and z_{n-2} only, and so forth. Thus, if the variety (17) has cardinality $|\mathbf{V}| < \infty$, then it can be computed sequentially by solving univariate polynomials and substituting backward.¹⁹

In many cases, including all that arise in this paper, the basis G is of the simple form

$$G = \{g_n(z_n), z_{n-1} - g_{n-1}(z_n), \dots, z_1 - g_1(z_n)\} \quad (19)$$

for some univariate polynomials g_n, \dots, g_1 , where g_n has degree $\deg(g_n) = |\mathbf{V}|$ and where $\deg(g_k) < |\mathbf{V}|$ for $k < n$.²⁰ In such cases, one computes the variety (17) by finding the $|\mathbf{V}|$ complex roots of g_n and then substituting each into the other $n - 1$ polynomials to obtain the $|\mathbf{V}|$ elements of (17).²¹

C.2 Algebraic numbers

The first step in finding the zeros of the polynomials in $G \subset \mathbb{Q}[z]$ is to find the roots of the univariate polynomial g_n . There are well known limits to what can be accomplished here: Abel's theorem states that there is no solution in radicals to general univariate polynomial equations of degree 5 or higher. Nevertheless, tools from computational algebra allow us to represent such solutions exactly.

Let $\bar{\mathbb{Q}} \subset \mathbb{C}$ denote the set of *algebraic numbers*: the complex numbers that are roots of nonzero polynomials with rational coefficients. The set $\bar{\mathbb{Q}}$ is a subfield of \mathbb{C} , and this fact and the definition of algebraic numbers are summarized by saying that $\bar{\mathbb{Q}}$ is the *algebraic closure* of \mathbb{Q} .²²

Every univariate polynomial $g \in \mathbb{Q}[x]$ can be factored as a product of *irreducible polynomials* in $\mathbb{Q}[x]$, which cannot themselves be further factored into products of nonconstant elements of $\mathbb{Q}[x]$.²³ If an irreducible polynomial $h \in \mathbb{Q}[x]$ is of degree k , it has k distinct roots $a_1, \dots, a_k \in \bar{\mathbb{Q}}$. The multiple of h whose leading term has coefficient 1 is

¹⁹The notion of Gröbner bases and the basic algorithm for computing them are due to Buchberger (1965). Cox et al. (2015) provide an excellent current account of Gröbner basis algorithms as well as a thorough introduction to the ideas summarized above.

²⁰According to the *shape lemma*, a sufficient condition for the reduced lex-order basis to be of form (19) is that each point in (17) has a distinct z_n component and that (18) be a *radical ideal*, meaning that if it includes some integer power of a polynomial, then it includes the polynomial itself. See Becker et al. (1994) and Kubler et al. (2014).

²¹Although we are only interested in elements of the variety (17) that lie in the state space Ξ , the solution methods described above work only if (17) has a finite number of solutions in \mathbb{C}^n .

²²Like \mathbb{C} , $\bar{\mathbb{Q}}$ is *algebraically closed*, in that every univariate polynomial with coefficients in $\bar{\mathbb{Q}}$ has a root in $\bar{\mathbb{Q}}$. It follows from this and the existence of lex-order Gröbner bases that when the variety (17) has a finite number of elements, the components of its elements are algebraic numbers.

²³"Typical" polynomials in $\mathbb{Q}[x]$ are irreducible: for instance, the quadratic $ax^2 + bx + c$ with $a, b, c \in \mathbb{Q}$ is reducible only if $\sqrt{b^2 - 4ac} \in \mathbb{Q}$. By *Gauss's lemma*, polynomial factorization in $\mathbb{Q}[x]$ is effectively equivalent to polynomial factorization in $\mathbb{Z}[x]$. For an excellent presentation of polynomial factorization algorithms, see von zur Gathen and Gerhard (2013, Chapters 14–16).

called the *minimal polynomial* of these roots. One often works instead with the multiple of h that is *primitive* in $\mathbb{Z}[x]$, meaning that its coefficients are integers with greatest common divisor 1.

Each algebraic number is uniquely identified by its minimal polynomial h and a label that distinguishes the roots of h from one another. For instance, one can label each root $a_j \in \bar{\mathbb{Q}}$ with a numerical approximation that is sufficiently accurate to distinguish a_j from the other roots. In computer algebra systems, the algebraic numbers with minimal polynomial h are represented by pairs consisting of h and an integer in $\{1, \dots, k\}$ that ranks the roots of h with respect to some ordering; for instance, the lowest integers are commonly assigned to the real roots of h in increasing order. Just as the symbol $\sqrt{2}$ is a label for the positive solution to $x^2 - 2 = 0$, the approach above provides labels for every algebraic number.²⁴

If the Gröbner basis G is of form (19), then we need only look for the roots of the irreducible factors h of the polynomial g_n , which are the possible values of $x_n \in \bar{\mathbb{Q}}$; then substitution into the univariate polynomials g_{n-1}, \dots, g_1 determines the corresponding values of the other variables. The fact that these latter values are generated from a fixed algebraic number allows us to work in subfields of $\bar{\mathbb{Q}}$ in which arithmetic operations are easy to perform. If the minimal polynomial h of $\alpha \in \bar{\mathbb{Q}}$ has degree $\deg(h)$, then for any polynomial f , one can find a polynomial f^* of degree $\deg(f^*) < \deg(h)$ such that $f(\alpha) = f^*(\alpha)$. It follows that the values of $g_{n-1}(\alpha), \dots, g_1(\alpha)$ are all elements of

$$\mathbb{Q}(\alpha) = \left\{ \sum_{k=0}^{\deg(h)-1} a_k \alpha^k : a_0, \dots, a_{\deg(h)-1} \in \mathbb{Q} \right\} \subset \bar{\mathbb{Q}},$$

called the *field extension* of \mathbb{Q} generated by α . Straightforward arguments show that the representation of elements of $\mathbb{Q}(\alpha)$ by sequences of coefficients (a_0, \dots, a_d) makes addition and multiplication in $\mathbb{Q}(\alpha)$ simple to perform. For further details on algebraic numbers and field extensions, we refer the reader to Dummit and Foote (2004, Chapter 13) and Cohen (1993, Chapter 4).

C.3 Examples

To illustrate the techniques above, we use them to compute the rest points of the $\text{BEP}(\tau^{\text{two}}, 1, \beta^{\text{min}})$ dynamic in the centipede games with $d = 3$ and $d = 4$ decision nodes. Since $x \in X$ and $y \in Y$, we need not explicitly write the laws of motion for the final components of x and y , as those components can be deduced from others and the simplex constraints.²⁵

²⁴There are exact methods based on classic theorems of Sturm and Vincent for isolating the real roots of a polynomial with rational coefficients; see McNamee (2007, Chapters 2 and 3) and Akritas (2010).

²⁵The Gröbner basis algorithm sometimes runs faster if all components are retained and the left-hand sides of the constraints $\sum_{i=1}^{s_1} x_i - 1 = 0$ and $\sum_{j=1}^{s_2} y_j - 1 = 0$ are included in the initial set of polynomials. Our Mathematica notebook includes both implementations.

EXAMPLE 3. The $\text{BEP}(\tau^{\text{two}}, 1, \beta^{\min})$ dynamic in centipede of length $d = 3$ is

$$\begin{aligned}\dot{x}_1 &= \frac{1}{2}(y_1(x_1 + x_2) + y_1(x_1 + x_3)) - x_1 \\ \dot{x}_2 &= \frac{1}{2}(y_2(x_1 + x_2) + (y_2 + (y_1)^2)(x_2 + x_3)) - x_2 \\ \dot{y}_1 &= ((x_2 + x_3)(x_1 + x_2) + (x_1)^2)(y_1 + y_2) - y_1.\end{aligned}\tag{20}$$

To find the rest points of this system, we substitute $x_3 = 1 - x_1 - x_2$ and $y_2 = 1 - y_1$ in the right-hand sides of (20) to obtain a system of three equations and three unknowns. We then compute a Gröbner basis of form (19) for the right-hand sides of (20):

$$\begin{aligned}\{ & 3(y_1)^4 - 8(y_1)^3 + 13(y_1)^2 - 12y_1 + 4, 4x_2 + 3(y_1)^3 - 5(y_1)^2 + 6y_1 - 4, \\ & 8x_1 - 3(y_1)^3 + 2(y_1)^2 - 9y_1 + 2\}.\end{aligned}\tag{21}$$

The initial quartic in (21) has roots $1, \frac{2}{3}$, and $(1 \pm \sqrt{7}i)/2$. Of course, only the first two roots could be components of states in Ξ . Substituting $y_1 = 1$ in the remaining polynomials in (21) and equating them to 0 yields $x_1 = 1$ and $x_2 = 0$, which with the simplex constraints gives us the backward induction state ξ^\dagger . Substituting $y_1 = \frac{2}{3}$ instead yields the interior state $\xi^* = (x^*, y^*) = ((\frac{1}{2}, \frac{1}{3}, \frac{1}{6}), (\frac{2}{3}, \frac{1}{3}))$. This is the complete set of rest points of the dynamic (20). \diamond

EXAMPLE 4. The $\text{BEP}(\tau^{\text{two}}, 1, \beta^{\min})$ dynamic in centipede of length $d = 4$ is

$$\begin{aligned}\dot{x}_1 &= \frac{1}{2}(y_1(x_1 + x_2) + y_1(x_1 + x_3)) - x_1 \\ \dot{x}_2 &= \frac{1}{2}((y_2 + y_3)(x_1 + x_2) + (y_2 + y_1(y_1 + y_3))(x_2 + x_3)) - x_2 \\ \dot{y}_1 &= \frac{1}{2}((x_2 + x_3)(x_1 + x_2) + (x_1)^2)((y_1 + y_2) + (y_1 + y_3)) - y_1 \\ \dot{y}_2 &= \frac{1}{2}((x_2 + x_3)(x_1 + x_3)(y_1 + y_2) + (x_1 + x_2(x_2 + x_3) + (x_3)^2)(y_2 + y_3)) - y_2.\end{aligned}$$

We can again compute a Gröbner basis of form (19). Its univariate polynomial is

$$\begin{aligned} & 4096(y_2)^8 - 28,608(y_2)^7 + 79,812(y_2)^6 - 64,332(y_2)^5 + 39,744(y_2)^4 \\ & - 9180(y_2)^3 + 648(y_2)^2 - 243y_2.\end{aligned}\tag{22}$$

This polynomial has root 0, which again generates the backward induction state ξ^\dagger . Dividing (22) by y_2 yields an irreducible seventh degree polynomial. Using the algorithms mentioned above, one can show that this polynomial has one real root, which we designate by $y_2^* = \text{Root}[4096\alpha^7 - 28,608\alpha^6 + 79,812\alpha^5 - 64,332\alpha^4 + 39,744\alpha^3 - 9180\alpha^2 + 648\alpha - 243, 1] \approx 0.3607$, and six complex roots. Substituting y_2^* into the remaining polynomials from the Gröbner basis and using the simplex constraints, we obtain an exact expression for the interior rest point ξ^* , whose components are elements of the field extension $\mathbb{Q}(y_2^*)$; their approximate values are $\xi^* = (x^*, y^*) \approx ((0.2575, 0.4358, 0.3068), (0.4095, 0.3607, 0.2298))$. \diamond

APPENDIX D: PROOF OF PROPOSITION 4(II)

The interior rest point ξ^* of the dynamic $\dot{x} = V(x)$ is locally stable if all eigenvalues of the derivative matrix $DV(\xi^*)$ have negative real part. Since each entry of the derivative matrix $D\mathcal{V}(\xi^*)$ is a polynomial with many terms that is evaluated at a state whose components are algebraic numbers, it is not feasible to compute its eigenvalues exactly. We circumvent this problem by computing the eigenvalues of the derivative matrix at a nearby rational state ξ and making use of a bound on the distances between the eigenvalues of the two matrices. This bound is established in Proposition 6.

As in Appendix A, let $s = s^1 + s^2 = d + 2$, let $\dot{\xi} = V(\xi)$, $V: \text{aff}(\Xi) \rightarrow T\Xi$ denote an instance of the BEP dynamics, and let $\mathcal{V}: \mathbb{R}^s \rightarrow \mathbb{R}^s$ denote the natural extension of V to \mathbb{R}^s . Observe that if $D\mathcal{V}(\xi)$ is diagonalizable, then so is $DV(\xi)$, and all eigenvalues of the latter are eigenvalues of the former. To state the proposition, we write $S = S^1 \cup S^2$ and omit population superscripts to define

$$\Delta = \max_{i \in S} \max_{k \in S} \sum_{j \in S} \frac{\partial^2 \mathcal{V}_i}{\partial \xi_j \partial \xi_k}(1, \dots, 1 | 1, \dots, 1). \tag{23}$$

PROPOSITION 6. *Suppose that $D\mathcal{V}(\xi)$ is (complex) diagonalizable with $D\mathcal{V}(\xi) = Q \times \text{diag}(\lambda)Q^{-1}$, and let λ^* be an eigenvalue of $D\mathcal{V}(\xi^*)$. Then there is an eigenvalue λ_i of $D\mathcal{V}(\xi)$ such that*

$$|\lambda^* - \lambda_i| < \frac{2\Delta}{s^{s/2-1}} \frac{\text{tr}(Q^*Q)^{s/2}}{|\det(Q)|} \sum_{k \in S} |\xi_k - \xi_k^*|. \tag{24}$$

The eigenvalue perturbation theorem (26) that begins the proof of the proposition bounds the distances between the eigenvalues of $D\mathcal{V}(\xi^*)$ and $D\mathcal{V}(\xi)$, but neither term on its right-hand side is feasible to compute. The second paragraph of the proof provides a bound on the condition number $\kappa_\infty(Q)$ that does not require the computation of the inverse of the (algebraic-valued) eigenvector matrix Q . The third paragraph provides a bound on the norm of $D\mathcal{V}(\xi) - D\mathcal{V}(\xi^*)$, which is needed because numerically evaluating of the entries of $D\mathcal{V}(\xi^*)$ with guaranteed precision is computationally infeasible. Two further devices that we employ to improve the bound and speed its computation are described after the proof of the proposition.

PROOF OF PROPOSITION 6. For $M \in \mathbb{R}^{s \times s}$, let

$$\|M\|_\infty = \max_{1 \leq i \leq s} \sum_{j=1}^s |M_{ij}| \tag{25}$$

denote the maximum row sum norm of M . Let $\kappa_\infty(Q) = \|Q\|_\infty \|Q^{-1}\|_\infty$ be the condition number of Q with respect to norm (25). The following eigenvalue perturbation theorem (Horn and Johnson (2013, Observation 6.3.1)) follows from the Geršgorin disk theorem and the submultiplicativity of matrix norms:

$$|\lambda^* - \lambda_i| \leq \kappa_\infty(Q) \|D\mathcal{V}(\xi) - D\mathcal{V}(\xi^*)\|_\infty. \tag{26}$$

To bound $\kappa_\infty(Q)$, let $\|M\|_2$ denote the spectral norm of M (i.e., the largest singular value of M) and let $\kappa_2(Q) = \|Q\|_2 \|Q^{-1}\|_2$ be the condition number of Q with respect to this norm. Since the maximum row sum and spectral norms differ by a factor of at most \sqrt{s} (Horn and Johnson (2013, Problem 5.6.P23)), it follows that

$$\kappa_\infty(Q) \leq s \kappa_2(Q). \quad (27)$$

Also, Guggenheimer et al. (1995) (see also Merikoski et al. (1997)) show that

$$\kappa_2(Q) < \frac{2}{|\det(Q)|} \left(\frac{\text{tr}(Q^*Q)}{s} \right)^{s/2}. \quad (28)$$

To bound the final expression in (26), note that by construction, each component of the BEP dynamics $\dot{x} = \mathcal{V}(x)$ is the difference between a sum of monomials in the components of ξ with positive coefficients and a linear term. Thus, the second derivatives of $\mathcal{V}_i(\xi)$ are sums of monomials with positive coefficients. Since every component of every state $\xi \in \Xi$ is at most 1, we therefore have

$$\max_{\xi \in \Xi} \left| \frac{\partial^2 \mathcal{V}_i}{\partial \xi_j \partial \xi_k}(\xi) \right| \leq \frac{\partial^2 \mathcal{V}_i}{\partial \xi_j \partial \xi_k}(1, \dots, 1 | 1, \dots, 1). \quad (29)$$

Thus, the fundamental theorem of calculus, (29), and (23) imply that

$$\begin{aligned} \|D\mathcal{V}(\xi) - D\mathcal{V}(\xi^*)\|_\infty &\leq \max_{i \in S} \sum_{j \in S} \sum_{k \in S} \frac{\partial^2 \mathcal{V}_i}{\partial \xi_j \partial \xi_k}(1, \dots, 1 | 1, \dots, 1) \times |\xi_k - \xi_k^*| \\ &\leq \Delta \sum_{k \in S} |\xi_k - \xi_k^*|. \end{aligned} \quad (30)$$

Combining inequalities (26), (27), (28), and (30) yields inequality (24). \square

When applying Proposition 6, one can choose Q to be any matrix of eigenvectors of $D\mathcal{V}(\xi)$. Guggenheimer et al. (1995) suggest that choosing the eigenvectors to have Euclidean norm 1 (which if done exactly makes the expression in parentheses in (28) equal 1) leads to the lowest bounds. We apply this normalization in the final step of our analysis.

To use this bound to establish the stability of the interior rest point ξ^* , we choose a rational point ξ close to ξ^* , compute the eigenvalues of the derivative matrix $D\mathcal{V}(\xi)$, and evaluate the bound from Proposition 6. The eigenvalues of $D\mathcal{V}(\xi)$ all have negative real part as long as ξ is reasonably close to ξ^* .

If ξ is close enough to ξ^* that the bound is smaller than the magnitude of the real part of any eigenvalue of $D\mathcal{V}(\xi)$, we can conclude that the eigenvalues of $D\mathcal{V}(\xi^*)$ all have negative real part and, hence, that ξ^* is asymptotically stable.

Selecting state ξ involves a trade-off: choosing ξ closer to ξ^* reduces the bound, but doing so also leads the components of ξ to have larger numerators and denominators, which slows the computation of the bound significantly. In all cases, we are able to choose ξ satisfactorily and to conclude that ξ^* is asymptotically stable. For further details about how the computations are implemented, see the Supplemental Material.

APPENDIX E: PROOF OF PROPOSITION 5

Letting $K = \{1, \dots, \kappa\}$, we can write the population 1 equations of the BEP($\tau^{\text{all}}, \kappa, \beta^{\text{min}}$) dynamic as

$$\dot{x}_i = \sum_{r: S^1 \times K \rightarrow S^2} \left(\prod_{\ell \in S^1, \lambda \in K} y_{r_{\ell\lambda}} \right) \mathbf{1}[i = \min(\text{argmax}_{k \in S^1} \pi_k^1(r))] - x_i,$$

where $\pi_k^1(r) = \sum_{m=1}^{\kappa} A_{kr_{km}}$. (31)

The result function $(\ell, \lambda) \mapsto r_{\ell\lambda}$ specifies the strategy in S^2 played by an agent's match partner during the λ th test of strategy ℓ for all $\ell \in S^1$ and $\lambda \in K$. The second piece of (31) specifies the probability of a given result, and the third piece indicates whether strategy i is the minimal optimal strategy for this result.

If there are two or more occurrences of strategies from S^2 other than 1, then all partial derivatives of the product in (31) equal 0. Thus, for the purposes of computing the Jacobian, we need only consider results in which there are 0 or 1 match partners playing strategies other than strategy $1 \in S^2$. These results comprise the following possibilities:

(i) If all match partners play strategy $1 \in S^2$, then strategy $1 \in S^1$ earns total payoff $\kappa \cdot 0$ and all other strategies earn total payoff $\kappa \cdot (-1)$, so strategy 1 has the best experienced payoff.

(ii) If the lone match against another strategy $j \in S^2 \setminus \{1\}$ occurs when the revising agent plays strategy $1 \in S^1$, then total payoffs are as above and strategy 1 has the best experienced payoff.

(iii) If the lone match against another strategy occurs when the revising agent plays strategy $i \in S^1 \setminus \{1\}$ and if this match occurs against an opponent playing strategy $j \in S^2 \setminus \{1\}$, then (using the payoffs A_{ij} defined in (2)) strategy 1 is the minimal strategy earning the best experienced payoff if

$$\kappa \cdot 0 \geq (\kappa - 1) \cdot (-1) + \begin{cases} 2i - 2 & \text{if } i \leq j \\ 2j - 3 & \text{if } i > j; \end{cases}$$

otherwise, strategy i uniquely obtains the best experienced payoff.

Accounting for all of these possibilities, including the fact that the matches in cases (ii) and (iii) can occur during any of the κ tests of the strategy in question, we have

$$\dot{x}_1 = (y_1)^{\kappa s^1} + \kappa (y_1)^{\kappa s^1 - 1} \left(\sum_{j=2}^{s^2} y_j + \sum_{i=2}^{s^1} \left(\sum_{j=2}^{i-1} y_j \mathbf{1}_{2j-2 \leq \kappa} + \sum_{j=i}^{s^2} y_j \mathbf{1}_{2i-1 \leq \kappa} \right) \right) - x_1 + O((y_{-1})^2)$$
(32a)

$$\dot{x}_i = \kappa (y_1)^{\kappa s^1 - 1} \left(\sum_{j=2}^{i-1} y_j \mathbf{1}_{2j-3 \geq \kappa} + \sum_{j=i}^{s^2} y_j \mathbf{1}_{2i-2 \geq \kappa} \right) - x_i + O((y_{-1})^2),$$

where $y_{-1} = \sum_{j=2}^{s^2} y_j$ and $i \in S^1 \setminus \{1\}$.

Turning to population 2, the test results with 0 or 1 match against opponents playing strategies other than $1 \in S^1$ comprise the following possibilities:

- (i) If all match partners play strategy $1 \in S^1$, then all strategies earn total payoff $\kappa \cdot 0$, so strategy 1 is the minimal strategy earning the best experienced payoff.
- (ii) If the lone match against another strategy occurs when the revising agent plays strategy $j \in S^2$, then strategy j earns a positive total payoff and other strategies earn total payoff 0, so strategy j has the best experienced payoff.

Accounting for both possibilities, we obtain

$$\begin{aligned} \dot{y}_1 &= (x_1)^{\kappa s^2} + \kappa(x_1)^{\kappa s^2-1} \sum_{i=2}^{s^1} x_i - y_1 + O((x_{-1})^2) \\ \dot{y}_j &= \kappa(x_1)^{\kappa s^2-1} \sum_{i=2}^{s^1} x_i - y_j + O((x_{-1})^2), \end{aligned} \tag{32b}$$

where $x_{-1} = \sum_{i=2}^{s^1} x_i$ and $j \in S^1 \setminus \{1\}$.

Taking the derivative of (32a) and (32b) at state ξ^\dagger , we obtain the matrix (we write this matrix for the case of d even, so that $s^1 = s^2$; roughly speaking, the case of d odd corresponds to removing the final column)

$$D\mathcal{V}(\xi^\dagger) = \begin{pmatrix} -1 & 0 & \cdots & \cdots & 0 & \kappa s^1 & + & \cdots & \cdots & + \\ 0 & -1 & \ddots & \ddots & \vdots & 0 & \kappa \mathbf{1}_{2i-2 \geq \kappa} & \cdots & \cdots & \kappa \mathbf{1}_{2i-2 \geq \kappa} \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \kappa \mathbf{1}_{2j-3 \geq \kappa} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & 0 & -1 & 0 & \kappa \mathbf{1}_{2j-3 \geq \kappa} & \cdots & \kappa \mathbf{1}_{2j-3 \geq \kappa} & \kappa \mathbf{1}_{2i-2 \geq \kappa} \\ \hline \kappa s^2 & \kappa & \cdots & \cdots & \kappa & -1 & 0 & \cdots & \cdots & 0 \\ 0 & \kappa & \cdots & \cdots & \kappa & 0 & -1 & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \kappa & \cdots & \cdots & \kappa & 0 & \cdots & 0 & 0 & -1 \end{pmatrix}. \tag{33}$$

Each + above represents the number that makes the column sum in the block equal κs^1 . If all indicator functions in the upper-right block of $D\mathcal{V}(\xi^\dagger)$ equal 0, then each + in (33) equals κs^1 , implying that this block is the zero operator on TY . In this case, (33) acts as a block triangular matrix on $T\Xi$, and so its lone eigenvalue with respect to $T\Xi$ is -1 , implying that ξ^\dagger is stable.

To check that the indicators are all 0 when d is even, it is enough to consider the indicator for entry $j = i = s^2 \equiv \frac{1}{2}d + 1$, which is 0 if and only if $\kappa \geq d + 1$. When d is odd, it is enough to check the indicator for entry $j = i = s^1 - 1 \equiv \frac{1}{2}(d + 1)$, which is 0 if and only if $\kappa \geq d$. We conclude that ξ^\dagger is asymptotically stable in these cases.

To show that ξ^\dagger is unstable in the remaining cases (when $\kappa \geq 2$ and $d \geq 3$), write

$$\chi_{ij}^\kappa = \begin{cases} \mathbf{1}_{2i-2 \geq \kappa} & \text{if } i \leq j \\ \mathbf{1}_{2j-3 \geq \kappa} & \text{if } i > j \end{cases}, \quad \chi_{i\Sigma}^\kappa = \sum_{j=2}^{s^2} \chi_{ij}^\kappa, \quad \text{and} \quad \chi_{\Sigma\Sigma}^\kappa = \sum_{i=2}^{s^1} \sum_{j=2}^{s^2} \chi_{ij}^\kappa$$

for $i, j \geq 2$. A straightforward calculation shows that $\lambda = \kappa \sqrt{\chi_{\Sigma\Sigma}^\kappa} - 1$ is an eigenvalue of $D\mathcal{V}(\xi^\dagger)$ that corresponds to eigenvector

$$z = \left(-\chi_{\Sigma\Sigma}^\kappa, \chi_{2\Sigma}^\kappa, \dots, \chi_{s^1\Sigma}^\kappa \mid - (s^2 - 1) \sqrt{\chi_{\Sigma\Sigma}^\kappa}, \sqrt{\chi_{\Sigma\Sigma}^\kappa}, \dots, \sqrt{\chi_{\Sigma\Sigma}^\kappa} \right).$$

Since $\kappa \geq 2$, λ is positive whenever at least one of the indicators in $D\mathcal{V}(\xi^\dagger)$ equals 1. Combining this with the previous argument, we conclude that ξ^\dagger is unstable whenever it is not asymptotically stable.

REFERENCES

- Akritis, A. G. (2010), “Vincent’s theorem of 1836: Overview and future research.” *Journal of Mathematical Sciences*, 168, 309–325. [1350, 1375]
- Basu, Kaushik (1994), “The traveler’s dilemma: Paradoxes of rationality in game theory.” *American Economic Review Papers and Proceedings*, 84, 391–395. [1351]
- Becker, Eberhard, Maria Grazia Marinari, Teo Mora, and Carlo Traverso (1994), “The shape of the Shape Lemma.” In *ISSAC ’94: Proceedings of the International Symposium on Symbolic and Algebraic Computation* (J. von zur Gathen and M. Giesbrecht, eds.), 129–133, ACM. [1374]
- Ben-Porath, Elchanan (1997), “Rationality, Nash equilibrium and backwards induction in perfect-information games.” *Review of Economic Studies*, 64, 23–46. [1348]
- Benaïm, Michel and Jörgen W. Weibull (2003), “Deterministic approximation of stochastic evolution in games.” *Econometrica*, 71, 873–903. [1348, 1353]
- Berkemer, Rainer (2008), “Disputable advantage of experience in the travelers’ dilemma.” Unpublished paper, Technical University of Denmark, Department of Mathematics. [1351]
- Binmore, Ken (1987), “Modeling rational players: Part I.” *Economics and Philosophy*, 3, 179–214. [1348]
- Binmore, Ken (1998), *Game Theory and the Social Contract, Volume 2: Just Playing*. MIT Press, Cambridge. [1352]
- Björnerstedt, Jonas and Jörgen W. Weibull (1996), “Nash equilibrium and evolution by imitation.” In *The Rational Foundation of Economic Behavior* (Kenneth J. Arrow, Enrico Colombatto, Mark Perlman, and Christian Schmidt, eds.), 155–171, London: MacMillan. [1352]

- Buchberger, Bruno (1965), "Ein Algorithmus zum Auffinden der Basiselemente des Restklassenrings nach einem nulldimensionalen Polynomideal." Translated by M. P. Abramson as "An algorithm for finding the basis elements of the residue class ring of a zero-dimensional polynomial ideal" in *Journal of Symbolic Computation* 41 (2006), 475–511. [1350, 1374]
- Cárdenas, Juan Camilo, César Mantilla, and Rajiv Sethi (2015), "Stable sampling equilibrium in common pool resource games." *Games*, 6, 299–317. [1351]
- Cohen, Henri (1993), *A Course in Computational Algebraic Number Theory*. Springer, Berlin. [1375]
- Collins, George E. (1975), "Quantifier elimination for the theory of real closed fields by cylindrical algebraic decomposition." In *Automata Theory and Formal Languages 2nd GI Conference Kaiserslautern*, volume 33 of Lecture Notes in Computer Science, 134–183, Springer, Berlin. [1364]
- Cox, David, John Little, and Donal O'Shea (2015), *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*, Fourth edition. Springer International, Cham, Switzerland. [1350, 1374]
- Cressman, Ross (1996), "Evolutionary stability in the finitely repeated prisoner's dilemma game." *Journal of Economic Theory*, 68, 234–248. [1351]
- Cressman, Ross (2003), *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, Massachusetts. [1351]
- Cressman, Ross and Karl H. Schlag (1998), "On the dynamic (in)stability of backwards induction." *Journal of Economic Theory*, 83, 260–285. [1351]
- Dekel, Eddie and Faruk Gul (1997), "Rationality and knowledge in game theory." In *Advances in Economics and Econometrics: Theory and Applications*, volume 1 (David M. Kreps and K. F. Wallis, eds.), 87–172, Cambridge University Press, Cambridge. [1348]
- Droste, Edward, Michael Kosfeld, and Mark Voorneveld (2003), "Best-reply matching in games." *Mathematical Social Sciences*, 46, 291–309. [1351]
- Dummit, David S. and Richard M. Foote (2004), *Abstract Algebra*, Third edition. Wiley, Hoboken, New Jersey. [1375]
- Friedman, Daniel and Ryan Oprea (2012), "A continuous dilemma." *American Economic Review*, 102, 337–363. [1352]
- Gilboa, Itzhak and Akihiko Matsui (1991), "Social stability and equilibrium." *Econometrica*, 59, 859–867. [1350, 1351, 1366]
- Guggenheimer, Heinrich W., Alan S. Edelman, and Charles R. Johnson (1995), "A simple estimate of the condition number of a linear system." *College Mathematics Journal*, 26, 2–5. [1378]
- Halpern, Joseph Y. (2001), "Substantive rationality and backward induction." *Games and Economic Behavior*, 37, 425–435. [1348]

- Hofbauer, Josef (1995), “Stability for the best response dynamics.” Unpublished manuscript, University of Vienna. [1350, 1351, 1366]
- Horn, Roger A. and Charles R. Johnson (2013), *Matrix Analysis*, Second edition. Cambridge University Press, New York. [1377, 1378]
- Izquierdo, Luis R., Segismundo S. Izquierdo, and William H. Sandholm (2018), “An introduction to ABED: Agent-based simulation of evolutionary game dynamics.” Unpublished manuscript, Universidad de Burgos, Universidad de Valladolid, and University of Wisconsin. [1352]
- Jehiel, Philippe (2005), “Analogy-based expectation equilibrium.” *Journal of Economic Theory*, 123, 81–104. [1352]
- Kosfeld, Michael, Edward Droste, and Mark Voorneveld (2002), “A myopic adjustment process leading to best reply matching.” *Journal of Economic Theory*, 40, 270–298. [1351]
- Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson (1982), “Rational cooperation in the finitely repeated prisoner’s dilemma.” *Journal of Economic Theory*, 27, 245–252. [1351]
- Kubler, Felix, Philipp Renner, and Karl Schmedders (2014), “Computing all solutions to polynomial equations in economics.” In *Handbook of Computational Economics*, volume 3 (Karl Schmedders and Kenneth L. Judd, eds.), 599–652, Elsevier, Amsterdam. 11. [1350, 1374]
- Mantilla, César, Rajiv Sethi, and Juan Camilo Cárdenas (2019), “Efficiency and stability of sampling equilibrium in public good games.” *Journal of Public Economic Theory*, forthcoming. [1351]
- McKelvey, Richard D. and Thomas R. Palfrey (1992), “An experimental study of the centipede game.” *Econometrica*, 60, 803–836. [1348, 1352]
- McNamee, J. M. (2007), *Numerical Methods for Roots of Polynomials, Part I*, volume 14 of *Studies in Computational Mathematics*. Elsevier, Amsterdam. [1350, 1375]
- Merikoski, Jorma Kaarlo, Uoti Urpala, Ari Virtanen, Tin-Yau Tam, and Frank Uhlig (1997), “A best upper bound for the 2-norm condition number of a matrix.” *Linear Algebra and its Applications*, 254, 355–365. [1378]
- Osborne, Martin J. and Ariel Rubinstein (1998), “Games with procedurally rational players.” *American Economic Review*, 88, 834–847. [1348, 1351, 1354, 1356, 1364]
- Oyama, Daisuke, William H. Sandholm, and Olivier Tercieux (2015), “Sampling best response dynamics and deterministic equilibrium selection.” *Theoretical Economics*, 10, 243–281. [1351]
- Perea, Andrés (2014), “Belief in the opponents’ future rationality.” *Games and Economic Behavior*, 83, 231–254. [1348]
- Perko, Lawrence (2013), “Differential equations and dynamical systems.”, Third edition, volume 7 of *Texts in Applied Mathematics*. Springer Science & Business Media, New York. [1358, 1370, 1371]

- Ponti, Giovanni (2000), "Cycles of learning in the centipede game." *Games and Economic Behavior*, 30, 115–141. [1351]
- Radner, Roy (1980), "Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives." *Journal of Economic Theory*, 22, 136–154. [1352]
- Reny, Philip J. (1992), "Backward induction, normal form perfection, and explicable equilibria." *Econometrica*, 60, 627–649. [1348]
- Rosenthal, Robert W. (1981), "Games of perfect information, predatory pricing and the chain-store paradox." *Journal of Economic Theory*, 25, 92–100. [1347, 1354, 1355]
- Rustichini, Aldo (2003), "Equilibria in large games with continuous procedures." *Journal of Economic Theory*, 111, 151–171. [1351]
- Sandholm, William H. (2001), "Almost global convergence to p -dominant equilibrium." *International Journal of Game Theory*, 30, 107–116. [1351]
- Sandholm, William H. (2003), "Evolution and equilibrium under inexact information." *Games and Economic Behavior*, 44, 343–378. [1353]
- Sandholm, William H. (2010a), "Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium." *Games*, 1, 3–17. [1352]
- Sandholm, William H. (2010b), *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, Massachusetts. [1352, 1357]
- Sandholm, William H. (2015), "Population games and deterministic evolutionary dynamics." In *Handbook of Game Theory With Economic Applications*, volume 4 (H. Peyton Young and Shmuel Zamir, eds.), 703–778, Elsevier, Amsterdam. 13. [1352]
- Sandholm, William H., Segismundo S. Izquierdo, and Luis R. Izquierdo (2019), "Stability for best experienced payoff dynamics." Unpublished manuscript, University of Wisconsin, Universidad de Valladolid, and Universidad de Burgos. [1348, 1354]
- Sethi, Rajiv (2000), "Stability of equilibria in games with procedurally rational players." *Games and Economic Behavior*, 32, 85–104. [1348, 1351, 1354, 1364]
- Stalnaker, Robert (1996), "Knowledge, belief, and counterfactual reasoning in games." *Economics and Philosophy*, 12, 133–163. [1348]
- Taylor, Peter D. and Leo Jonker (1978), "Evolutionarily stable strategies and game dynamics." *Mathematical Biosciences*, 40, 145–156. [1351]
- von zur Gathen, Joachim and Jürgen Gerhard (2013), *Modern Computer Algebra*, Third edition. Cambridge University Press, Cambridge. [1350, 1374]
- Weibull, Jörgen W. (1995), *Evolutionary Game Theory*. MIT Press, Cambridge, Massachusetts. [1352]
- Xu, Zibo (2016), "Convergence of best-response dynamics in extensive-form games." *Journal of Economic Theory*, 162, 21–54. [1350, 1351, 1366]

Young, H. Peyton (1998), *Individual Strategy and Social Structure*. Princeton University Press, Princeton. [1348]

Co-editor Ran Spiegler handled this manuscript.

Manuscript received 14 December, 2018; final version accepted 19 April, 2019; available online 25 April, 2019.