An Introduction to *ABED*: Agent-Based Simulation of Evolutionary Game Dynamics[*]

Luis R. Izquierdo[†], Segismundo S. Izquierdo[‡], and William H. Sandholm[§]

October 2, 2019

Abstract

ABED is free and open-source software for simulating evolutionary game dynamics in finite populations. We explain how *ABED* can be used to simulate a wide range of dynamics considered in the literature and many novel dynamics. In doing so, we introduce a general model of revisions for dynamic evolutionary models, one that decomposes strategy updates into selection of candidate strategies, payoff determination, and choice among candidates. Using examples, we explore ways in which simulations can complement theory in increasing our understanding of strategic interactions in finite populations.

1. Introduction

Evolutionary game theory now encompasses a wide range of models that differ both in their basic assumptions and in the details of how shared assumptions are implemented. This diversity can make comparing different models and the predictions they generate a daunting task.

[†]Department of Civil Engineering, Edificio la Milanera, Calle de Villadiego, Universidad de Burgos, 09001 Spain. e-mail: lrizquierdo@ubu.es; website: www.luis.izqui.org.

[‡]BioEcoUva Research Institute on Bioeconomy. Department of Industrial Organization, Paseo del Cauce 59, Universidad de Valladolid, 47011 Spain. e-mail: segis@eis.uva.es; website: www.segis.izqui.org.

[§]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: www.ssc.wisc.edu/~whs.

ABED is free and open-source software for simulating evolutionary game dynamics in finite populations. It implements a wide range of alternative modeling assumptions within a single simple framework. This facilitates the comparison of different models, and aids in identifying the specific assumptions that underlie observed behaviors.

ABED follows an agent-based approach, with each agent in the model being represented as a distinct object within the code. It implements a variety of revision protocols (i.e., rules agents follow when updating their strategies), allowing it to simulate a substantial proportion of the dynamic evolutionary models appearing in the literature, as well as many models that have yet to be analyzed (see Table 1). Nearly all parameters in *ABED* can be modified while the program is running, making it easy to examine the consequences of changing assumptions about how agents make decisions, and facilitating the exploration of novel models.

To best relate *ABED* to existing formal models, we provide a general specification of stochastic evolutionary game dynamics for finite populations under random or complete matching. It is distinguished by a decomposition of decisions by revising agents into three stages: selection of candidate strategies, determination of their payoffs, and choice among the candidate strategies. This decomposition of revision protocols offers a unified framework for understanding a range of models that might otherwise not be easy to compare directly.

There is a large literature in physics that simulates evolutionary game models for which analytical results are not available. In contrast, *ABED* focuses on a simple class of models whose basic properties can be derived analytically in the large population limit. *ABED* thus has some important limitations. It only concerns 2-player games under random matching, so it cannot be used to simulate dynamics on networks or in spatial settings (cf. Szabó and Fáth (2007), Roca et al. (2009a), Szolnoki et al. (2014), and Perc (2018)), and interactions must be pairwise (cf. Perc et al. (2013)). Also, in *ABED* only the distribution of strategies evolves in time. *ABED* is thus unable to simulate coevolutionary game dynamics (cf. Pacheco et al. (2006) and Perc and Szolnoki (2010)), where the matching pattern or other aspects of the interactions adjust in tandem with agents' choices of strategies. Finally, agents in *ABED* do not have memory, i.e. their decisions only depend on the population distribution of strategies at the time of revision.

At the same time, *ABED* allows for some relatively unexplored modeling choices. Rather than requiring complete matching, *ABED* permits the user to specify the number of matches an agent engages in each period. A small but growing literature has shown that behavior under limited matching can be strikingly different from behavior that arises under the more common assumption of complete matching, and *ABED* is a powerful

I: Imitative/reproductive protocols

decision method		matching	
		complete matching	limited matching
	pairwise difference	replicator [a]	
	linear-*	replicator [b]	
	best	imitate the best [c]	imitate the best realization [d]
	logit	imitative logit [e]; Fermi process [f]	stochastic sampling Fermi process [g]
	positive proportional	Maynard Smith replicator [h]; freq. dependent Moran process [i]	stochastic sampling Moran process [j]

II: Direct protocols

decision method		matching/sampling		
		complete matching	limited matching	
			single sample	multiple samples
	pairwise difference	Smith [k]	—	—
	linear-*	—	—	—
	best	best response [l]	sample best response [m]	best experienced payoff [n]
	logit	logit [o]	sample logit [p]	—
	positive proportional	—	—	—

[a] Taylor and Jonker (1978); Helbing (1992); Schlag (1998); Hofbauer (1995a).

[b] Taylor and Jonker (1978); Weibull (1995); Björnerstedt and Weibull (1996).

[c] Hofbauer (1995a); Vega-Redondo (1997).

[d] Izquierdo and Izquierdo (2013); Loginov (2019).

[e] Weibull (1995).

[f] Traulsen and Hauert (2009).

[g] Woelfing and Traulsen (2009).

[h] Maynard Smith (1982).

[i] Binmore et al. (1995); Nowak et al. (2004); Taylor et al. (2004); Fudenberg et al. (2006); Fudenberg and Imhof (2008)

[j] Woelfing and Traulsen (2009).

[k] Smith (1984); Sandholm (2010a).

[l] Gilboa and Matsui (1991); Hofbauer (1995b); Kandori and Rob (1995).

[m] Sandholm (2001a); Kosfeld et al. (2002); Oyama et al. (2015).

[n] Sethi (2000); Sandholm et al. (2019, 2018).

[o] Blume (1997); Fudenberg and Levine (1998); Hofbauer and Sandholm (2007).

[p] Kreindler and Young (2013).

Table 1: Classes of revision protocols in *ABED* and principal instances of their dynamics.

tool for simulating these sometimes cumbersome models.[1] *ABED* also allows many agents to revise simultaneously, allowing one not only to test the robustness of standard approximation results to synchronized updating, but also to investigate applications in which simultaneous revisions play a basic role.

This paper also examines some roles for agent-based simulations in evolutionary game theory. Simulations are sometimes viewed with a skeptical eye in game and economic theory. They can be criticized on the grounds that any simulation run concerns a specific choice of parameter values, and that observing a certain outcome in a series of simulation runs does not imply that this outcome must always occur. But while simulations cannot prove theorems, they serve a variety of purposes that theorists should value. As theorem statements often invoke limits, simulations can be used to evaluate robustness to realistic choices of parameters, serving a similar role as small-sample robustness checks in econometrics. Simulations can help us interpret abstractly-defined solution concepts by allowing us to flesh out the implicit assumptions that underlie them. Finally, simulations are a powerful exploratory tool, one that can both suggest new models and provide conjectures for formal analysis. While one may object to some of the grander claims about the role of computer simulations in social science, such objections do not exclude many other ways that theory and simulation can serve complementary purposes. In what follows we aim to illustrate certain aspects of this symbiosis.[2]

1.1 Exact analysis, approximation, and simulation of population dynamics

The canonical model in evolutionary game theory is one in which a finite population of agents are recurrently matched to play a normal form game.[3] Each agent is occasionally given the opportunity to revise his pure strategy, basing his choice on the information about payoffs and current aggregate behavior he has obtained. Formulating such a model requires one to specify

  (i)   the number of agents N in the population,
  (ii) the n-strategy normal form game the agents are recurrently matched to play,

---

[1] See Osborne and Rubinstein (1998), Sethi (2000), Sandholm (2001b), Kosfeld et al. (2002), Oyama et al. (2015), Cárdenas et al. (2015), Mantilla et al. (2019), and Sandholm et al. (2019, 2018). *ABED* does not allow the number of matches to be stochastic (cf. Sánchez and Cuesta (2005), Roca et al. (2006, 2009b), Traulsen et al. (2007), and Oyama et al. (2015)).

[2] A recent paper by Adami et al. (2016a) has triggered a lively discussion about the role of agent-based simulations in evolutionary game theory—see the responses by Bellomo and Elaiw (2016), Hilbe and Traulsen (2016), Schuster (2016) and Tarnita (2016). These responses and the authors' reply (Adami et al. (2016b)) reveal a growing understanding of the complementary roles of simulation and analysis. Work by García and van Veelen (2016, 2018) nicely illustrates this complementarity.

[3] For a historical overview, see Sandholm (2010b, Ch. 1).

(iii) the rule describing how revision opportunities are assigned to the agents, and

(iv) the protocol according to which agents revise their strategies when opportunities to do so arise.

Of these items, the assignment rule (iii) and especially the revision protocol (iv) are particular to the evolutionary approach. In economic modeling, a revision protocol captures the information agents possess when revising, how that information is obtained, and the way that agents use the information to select a new strategy.[4] In biological contexts, revision is often interpreted as a death and birth event rather than as a conscious decision, and assignment rules and revision protocols are specified accordingly.[5]

The state of the resulting evolutionary process can be described in two ways. One can define the state to be the profile of strategies currently chosen by each of the N agents. Alternatively, one can specify the state to be the population distribution of the strategies chosen by these agents. As the name indicates, agent-based simulations start from the former specification, while formal analyses generally use the more parsimonious specification in terms of distributions over strategies, or *population states*.

In the latter approach, the analysis of the evolutionary process focuses on the behavior of a Markov chain $\{X_t^N\}$ on the finite set of population states, a grid $\mathcal{X}^N$ of mesh $\frac{1}{N}$ in the simplex $X \in \mathbb{R}^n$ of probability distributions over the n pure strategies. Under certain assumptions about the game and revision procedure, it is possible to analyze this Markov chain directly.[6] More typically, the Markov chain $\{X_t^N\}$ is not susceptible to direct analysis.

This difficulty can be circumvented by studying limiting versions of the Markov chain $\{X_t^N\}$. Here we are most interested in the large population limit. To understand the behavior of the process $\{X_t^N\}$ over some fixed time horizon $[0, T]$, we take the population size N to infinity to obtain limiting processes that are easier to analyze than the original Markov chain, and the analysis is then performed on the limiting processes so obtained. Because the limiting processes run in continuous-time on the continuous state space X, they can be studied using methods from calculus and analysis.

Focusing on these limiting processes is formally justified by suitable approximation results.[7] The most important of these is the finite-horizon deterministic approximation theorem (Benaïm and Weibull (2003), Roth and Sandholm (2013)). The theorem shows that over any fixed time horizon, as the population size N approaches infinity, the sample paths of the Markov chains $\{X_t^N\}$ converge to solutions of a deterministic *mean dynamic* defined

---

[4]See Weibull (1995), Björnerstedt and Weibull (1996), and Sandholm (2010b, 2015).

[5]See, e.g., Nowak et al. (2004) and Fudenberg et al. (2006).

[6]For instance, Monderer and Shapley (1996) and Blume (1997) on better response and logit dynamics in potential games, and Kandori and Rob (1995) on best response dynamics in supermodular games.

[7]We present these results in more detail in Appendices A.2 and A.3.

by the expected motion of $\{X_t^N\}$ from each state. This result formally links the wide array of deterministic evolutionary dynamics studied in the literature and the revision protocols that generate them (Table 1). To understand the behavior of the Markov chains $\{X_t^N\}$ near rest points x^* of the mean dynamic (in other words, to describe "equilibrium" behavior in the population), one can instead appeal to a diffusion approximation (Sandholm (2003)). This result shows that a "local behavior process", obtained by magnifying deviations of $\{X_t^N\}$ from x^* by a factor of $\sqrt{N}$, is approximated by the solution to a stochastic differential equation whose linear drift coefficient and constant diffusion coefficient are determined from the transition probabilities of $\{X_t^N\}$ near x^*.[8]

Of course, the population sizes relevant in applications are finite, and not always very large. Thus as a practical matter, one should ask to what extent results that rely on large population limits describe behavior in populations of moderate or small size. These questions are generally quite difficult to answer analytically. Simulations provide a simple way of evaluating the robustness of formal analyses whose tractability relies on sending parameter values to extremes.[9]

Mathematical tractability shapes not only the forms that questions about evolutionary models take, but also the specific assumptions that the models employ. As an example, most models of evolution for normal form games make the analytically convenient assumption that revising agents are able to evaluate strategies' expected payoffs. But in applications where agents from large populations engage in random matches, knowledge of expected payoffs would need to come from a complete matching (so that expected payoffs are just realized payoffs), from knowledge of the current population state (so that expected payoffs could be calculated), or from a central source. While these possibilities describe some applications, in others it seems likely that payoff information is more limited, with evaluations of strategies based on information obtained from samples, on direct but limited experiences, or from observations of others' experiences.[10] In some cases, this

---

[8]Other limiting analyses are used to understand the behavior of $\{X_t^N\}$ over longer time spans. *Stochastic stability analyses* consider behavior in the infinite horizon limit by means of the stationary distribution of $\{X_t^N\}$. Tractability is gained by considering the limit as the noise level in agents' decisions vanishes (see Foster and Young (1990); Kandori et al. (1993); Young (1993, 1998); Ellison (2000)), or as the population size grows large (see Binmore et al. (1995); Binmore and Samuelson (1997)), or both; see Sandholm (2010b, chapters 11 and 12) for an overview of this literature. Transitions between stable rest points in these limits are analyzed using methods from *large deviations theory*; see the previous references and Sandholm and Staudigl (2016) for the case of small noise limits, and Sandholm and Staudigl (2018) for the large population limit. Finally, *weak selection analyses* introduce a scaling parameter to study biological interactions with very small differences in fitness, so that the Markov process $\{X_t^N\}$ comes to resemble a perturbed random walk (see Nowak et al. (2004); Wu et al. (2010); Sample and Allen (2017)).

[9]Work in this spirit includes Wu et al. (2013), which uses simulations to assess the robustness of the weak selection limit.

[10]Early work emphasizing this point includes Robson and Vega-Redondo (1996), Schlag (1998, 1999),

additional level of detail can push models beyond the limits of mathematical tractability. One can then use simulations to assess whether assumptions imposed for simplicity are robust to more realistic alternatives, and, in cases where they are not, to look for alternative models and candidate theorems that seem likely to yield to formal analysis.[11]

1.2 Brief description of *ABED* and comparisons to other software

ABED[12] is free and open-source software for running agent-based simulations in the finite-population evolutionary game framework described at the start of Section 1.1. Its simple graphical user interface (Figure 1) allows one to implement a wide range of the revision protocols studied in the literature, including all of those appearing in Table 1, and to fine-tune many other aspects of the evolutionary process. At present, *ABED* consists of two separate computer programs: `ABED-1pop`, for single populations of agents matched to play symmetric two-player games; and `ABED-2pop`, for pairs of populations of agents matched to play (possibly asymmetric) two-player games.

ABED runs in *NetLogo* (Wilensky (1999)), a well-documented free and open-source platform for designing agent-based simulations and conducting automatic explorations of the parameter space. In the present context, the term "agent-based" means that each agent in an *ABED* simulation is represented in the code as a distinct object. This approach makes it relatively easy to modify *ABED*'s code to incorporate heterogeneity in agents' behaviors, locations, and interaction patterns.[13]

There are other useful software packages available for analyzing and simulating evolutionary game dynamics. *Dynamo* (Sandholm et al. (2012); Franchetti and Sandholm (2013)) is a suite of easy-to-use *Mathematica* notebooks for generating phase diagrams, vector fields, and other graphics related to mean dynamics from evolutionary game theory. *EvoDyn-3s* (Izquierdo et al. (2018)) is similar to *Dynamo* in spirit, but it only deals with one-population three-strategy games. On the positive side, it is very simple to use, requiring only a mouse, and performs basic computations using exact arithmetic. *BIRDS* (Dzonsons and Weibull (2013)) is a finite-population agent-based simulator that implements several revision protocols. *PDToolbox* (Barreto (2014)) is a set of functions coded in *Matlab* for analyzing both mean dynamics and finite-population agent-based models

---

Sethi (2000), and Sandholm (2001a).

[11]Working in the reverse direction, Izquierdo et al. (2013) show different ways in which mathematical analysis and computer simulation can be combined to better understand the dynamics of agent-based models.

[12]Downloadable from https://luis-r-izquierdo.github.io/abed/.

[13]For more on implementing, extending, and analyzing agent-based evolutionary models using NetLogo, see Izquierdo et al. (2019).

setup go ⟳ go once load parameters from file save parameters to file

ticks
4000

ticks-per-second
100

Strategy distributions (recent history)

1
2
3
4
5

1

Plot I-a

0

30000 milliseconds 40010

Strategy distributions (complete history)

1
2
3
4
5

1

Plot I-b

0

0 seconds 40

Game and population

payoff-matrix Change?

[[0 -1 1 0 0]
 [0 0 -1 1 0]
 [0 0 0 -1 1]
 [1 0 0 0 -1]
 [-1 1 0 0 0]]

1

On Off random-initial-condition?

initial-condition

[200 200 200 200 200]

n-of-agents 1000

Assignment of revision opportunities

On Off use-prob-revision?

prob-revision

n-of-revisions-per-tick 10

2

Plotting of output

plot-every-?-secs 0.01

duration-of-recent

On Off show-recent-history?

On Off show-complete-history?

6

Candidate selection

candidate-selection

direct ▼

n-of-candidates 2

Under imitative protocols, candidates are agents.
Under direct protocols, candidates are strategies.

3

Matching

On Off complete-matching?

for complete-matching

n-of-trials 9

for complete-matching = off
& candidate-selection = direct:

On Off single-sample?

4

Decision method

decision-method

best ▼

for decision-method = best:

tie-breaker

uniform ▼

mutations:

prob-mutation 0.000

for decision-method = logit:

log-noise-level 0.000

5

Auxiliary parameters

for candidate-selection = imitative:

On Off consider-imitating-s

On Off imitatees-with-repla

On Off self-matching?

for complete-matching = off:

On Off trials-with-replacement?

7

Strategies' expected payoffs (recent history)

1
2
3
4
5

0.39

Plot II-a

-0.435

30000 milliseconds 40010

Strategies' expected payoffs (complete history)

1
2
3
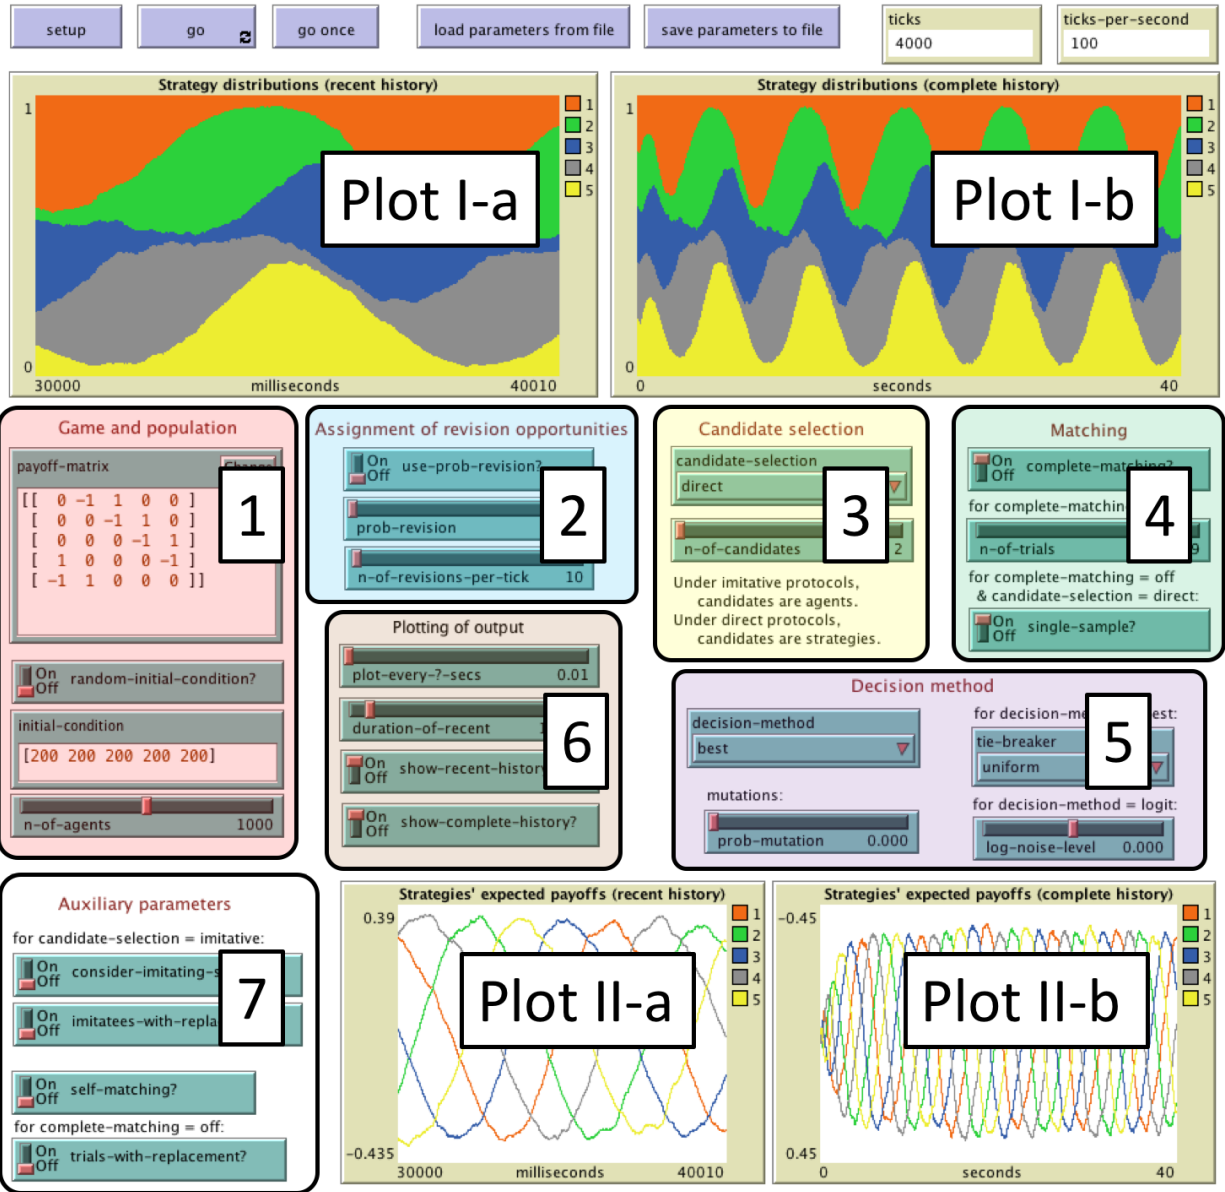4
5

-0.45

Plot II-b

0.45

0 seconds 40

Figure 1: Interface of ABED-1pop.

derived from built-in revision protocols.[14]

To start with basic comparisons, *ABED*, unlike *Dynamo* and *EvoDyn-3s* but like *BIRDS* and *PDToolbox*, simulates finite-population dynamics. Also, unlike *Dynamo*, *EvoDyn-3s* and *PDToolbox* but like *BIRDS*, *ABED* can be run without proprietary software. Turning to details, *ABED* is especially flexible in the specification of revision protocols, and incorporates revisions based on both complete and limited matching. Simulations in *ABED* can be inspected and parameters modified during runtime, making it easy to explore the effects of different modeling assumptions on population outcomes. Finally, *ABED* is exceptionally user-friendly, with a simple graphical user interface, and with extensive documentation and video tutorials available online.

1.3 Outline

The paper proceeds as follows. Section 2 provides a detailed description of *ABED*, focusing for simplicity on ABED-1pop. Section 2.1 explains *ABED*'s parameters, including a thorough account of those used to specify revision protocols. Section 2.2 (along with the Appendix) connects the parameter choices in *ABED* to a wide range of evolutionary dynamics studied in the literature, and suggests the broad scope for the simulation of unexplored models. The next three sections deal with *ABED*'s output: Section 2.3 describes *ABED*'s plots and monitors, Section 2.4 discusses the modification of parameters during runtime, and Section 2.5 indicates how to conduct automatic Monte Carlo explorations of *ABED*'s parameter space.

In Section 3 we present a variety of examples that show *ABED*'s capabilities and indicate some roles for simulation in evolutionary game theory. These examples also illustrate a number of basic features and applications of finite-population stochastic evolutionary models: the nature of "equilibrium" play in finite populations of revising agents (§3.1, §3.2); the sensitivity of the stability and location of equilibrium to modeling details (§3.2, §3.3, §3.8); the accuracy of deterministic approximations of the evolutionary process (§3.3, §3.4, §3.5); large deviations properties and waiting times until extinction (§3.2); and foundations for set-valued stability concepts (§3.6). Some concluding remarks are offered in Section 4.

The paper's formal analyses appear in the Appendix, which includes a formal framework for describing stochastic evolutionary processes in finite populations under random

---

[14]It is also worth mentioning the excellent websites *EvoLudo* (Hauert (2018)) and *Evolution and Games* (García and van Veelen (2012)), which include several interactive tutorials with accompanying software. Also, Friedman and Sinervo (2015) provide software written in R and three Excel spreadsheets that implement several evolutionary game dynamics described in Friedman and Sinervo (2016).
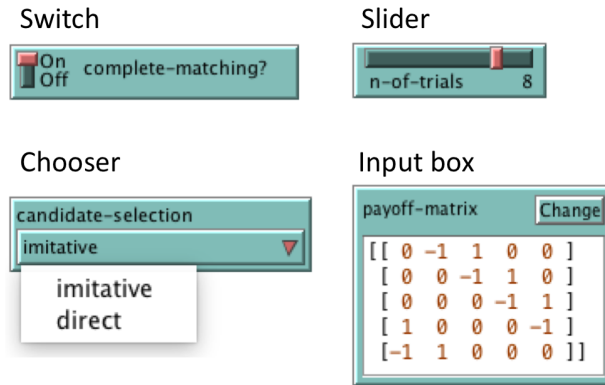
Figure 2: *NetLogo* interface elements used in *ABED* to input parameter values.

matching (including all those implemented in *ABED*), reviews the deterministic and diffusion approximations, and presents a number of examples. An online appendix provides detailed instructions to run computational experiments with *ABED* and collects the full parameter settings used in the examples discussed in Section 3.

2. Description of *ABED*

ABED's interface (Figure 1) contains buttons for running the simulation and for loading/saving parameter files, monitors that show the passage of time, interface elements organized into blocks for choosing simulation parameters, and plots that present simulation results.

The top row of the interface contains five blue buttons. The `setup` button initializes the model, the `go` button makes the model run indefinitely (until the same button is clicked again), and the `go once` button runs the model for only one period (or "tick" in *NetLogo* parlance). The two remaining buttons are used to load and save parameter files.

2.1 *ABED*'s parameters

In this section, we explain the parameters of *ABED*. In the course of doing so, we explain how simulations in *ABED* are executed, and we describe the range of models that *ABED* is able to enact. It is worth reiterating that nearly all parameter values can be modified while *ABED* is running, with immediate effects on how the simulation proceeds.[15]

Input in *ABED* is entered using *switches* (for binary parameters), *choosers* (dropdown menus for parameters with multiple named values), *sliders* (for numerical parameters),

---

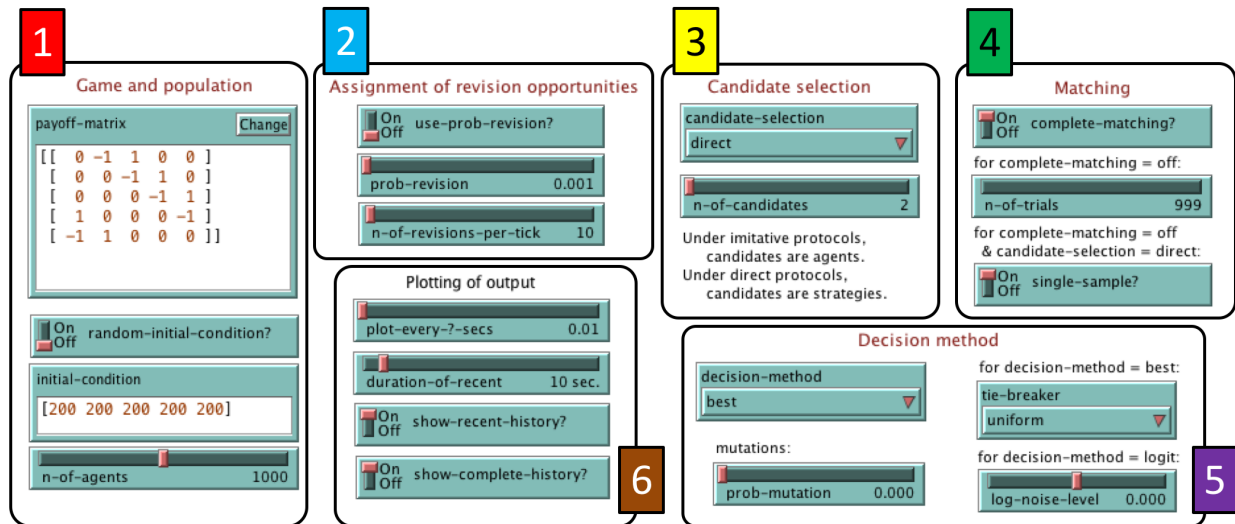[15]The payoff matrix and the initial conditions are the only two exceptions.

Figure 3: Main parameters in `ABED-1pop`.

and an *input box* (to enter the payoff matrix). Each is illustrated in Figure 2. In what follows we identify an interface item with the parameter it controls.

The main parameters that specify a simulation (Blocks 1–5) appear in the middle part of *ABED*'s interface (Figure 3). Block 6 contains parameters that control the plotting of output, and Block 7 includes a few secondary parameters.

2.1.1 The game, the population size, and the initial state

Block 1 contains the parameters used to specify the game, the population size, and the initial state.

The symmetric normal form game is entered into the `payoff-matrix` input box in the form of a square matrix. *ABED* deduces the number of strategies (henceforth `n-of-strategies`) from the size of this square matrix. How agents are matched to play this game is defined by parameters from later blocks.

The initial population state of the simulation is set as follows:

- If the switch `random-initial-condition?` is *on*, then the number of agents in the population is specified using the slider `n-of-agents`, and each of these agents is assigned one of the `n-of-strategies` strategies randomly.

- If `random-initial-condition?` is *off*, then the initial condition is read from the `initial-condition` input box. The input here is a list of `n-of-strategies` nonnegative integers whose *i*-th element specifies the number of agents using strategy *i* at

the beginning of the simulation. In this case `n-of-agents` is set to the sum of the elements of `initial-condition`.

As with the other parameters, the value of `n-of-agents` can be changed while the simulation is running. If the slider is moved to the left, randomly chosen agents are removed from the population; if it is moved to the right, randomly chosen agents are cloned. Thus on average, the proportions of agents playing each strategy remain the same, although the actual effect on these proportions is stochastic.

2.1.2 Assignment of revision opportunities to agents

The assignment of revision opportunities to agents during each period is controlled by the options in Block 2.

- If the switch `use-prob-revision?` is *on*, then assignments are stochastic and independent, with the probability that an agent is assigned an opportunity being specified using the slider `prob-revision`.

- If instead `use-prob-revision?` is *off*, then a fixed number of opportunities are allocated among the `n-of-agents` agents at random. This number is specified using the slider `n-of-revisions-per-tick`.

Since revising agents obtain information at the start of the simulation period, before any revisions have occurred, *ABED* can simulate anything from simultaneous updating (`prob-revision` = 1) to sequential updating (`n-of-revisions-per-tick` = 1).

2.1.3 Revision protocols

ABED allows the user considerable flexibility in specifying the agents' revision protocols. Here revision protocols are defined by three main elements: how candidate strategies are selected, how agents evaluate candidate strategies, and how the latter information is mapped to the probabilities of choosing each candidate strategy. These elements are defined using the parameters from Blocks 3–5 of the interface, along with some minor parameters in Block 7 whose default values (*off*) are usually satisfactory.

Selection of candidate strategies

The parameters in Block 3 determine how agents select and evaluate candidate strategies. The chooser `candidate-selection` specifies which of the two basic options is applied.

- *Direct protocols*. If the `candidate-selection` is set to *direct*, then the revising agent's additional candidate strategies—those besides his current strategy—are a random subset of the other strategies, with a strategy's current share having no bearing on whether it is chosen as a candidate. In this case, the slider `n-of-candidates` determines the total number of strategies the revising agent considers. The revising agent's current strategy is always part of the set of candidates. The agent evaluates a strategy's performance by explicitly testing it in matches with opponents, which are described in the next section.[16]

- *Imitative/reproductive protocols*. If `candidate-selection` is set to *imitative*, then an agent decides which strategies besides his current one to include as candidates by choosing other agents from the population at random. Specifically, the revising agent will compile a multiset of `n-of-candidates` agents to copy the strategy of one of them. The revising agent is always part of this multiset of agents. The selection of the other (`n-of-candidates` - 1) agents is conducted randomly, so popular strategies in the population are more likely to be observed than less popular ones.[17] All agents in the set of candidates play the game against a sample of opponents (as described below), and the revising agent observes the total payoffs obtained by each. In the end, the revising agent has a list of `n-of-candidates` candidate strategies with associated payoffs, where the same candidate strategy may appear several times on the list.

  Biological models like the frequency-dependent Moran process (Section 3.7) assume that in each period, some randomly selected agents die and some of the remaining agents reproduce, with agents earning higher payoffs being more likely to have offspring. Although their interpretations are quite different, imitative and reproductive processes are formally very similar. When combined with what follows, the imitative procedure above can be used to define reproductive processes, with the various parameter choices determining the particular biological model to be simulated.[18]

---

[16]When this matching is complete, the payoff that an agent currently playing strategy i obtains from testing strategy j equals the expected payoff the agent would obtain in a random match from switching to j. But this payoff typically differs from the current payoff of an opponent now playing strategy j, who faces one additional i player and one fewer j player than the original agent. The latter payoff would be considered by the original agent if he were using an imitative protocol. The consequences of this distinction are usually minor unless the population size is small. For further discussion, see Sandholm (2010b, Section 11.4).

[17]Two Block 7 parameters influence this selection. Setting `consider-imitating-self?` to *on* allows an agent to choose himself as one or more of the additional (`n-of-candidates` - 1) candidates. Setting `imitatees-with-replacement?` to *on* has a self-explanatory effect.

[18]Helbing (1992), Weibull (1995), and Schlag (1998) are early works that develop this basic connection between reproductive and imitative models.

Matching and payoffs

The matching process, which determines the total payoff an agent obtains when playing or testing a strategy is governed by the three parameters in Block 4:

- If the switch `complete-matching?` is *on*, then an agent playing a strategy is matched against all other agents.[19]

- If `complete-matching?` is *off*, then the agent plays his current strategy against a random sample of opponents.[20] The size of this sample is specified using the slider `n-of-trials`.[21] If in addition the protocol is *direct*, the switch `single-sample?` determines how many samples of opponents the revising agent faces during testing:

  - If `single-sample?` is *on*, the revising agent draws a single sample of `n-of-trials` opponents and tests each of his candidate strategies against this sample.
  - If `single-sample?` is *off*, the agent tests each of his candidate strategies against distinct, independent samples of `n-of-trials` opponents. Thus, different strategies are tested against potentially different opponents.

Decision method

The procedures described in the previous two sections leave a revising agent with an assignment of a payoff to each strategy in a multiset of `n-of-candidates` elements. How the agent responds to this information—the probabilities with which he winds up switching to each strategy—is specified in Block 5. The main parameter here is the `decision-method`:

- If the `decision-method` is *best*, then the agent chooses a strategy whose payoff was highest. The chooser `tie-breaker` includes various basic options for handling multiple maxima.[22] The final option, *random walk*, has agents randomize among optimal strategies with probabilities determined by an auxiliary random walk, allowing the population to wander around a component of Nash equilibria (see Section 3.6 and Appendix A.4).

---

[19]For either setting of `complete-matching?`, setting `self-matching?` (Block 7) to *on* allows an agent to be matched against himself.

[20]Setting `trials-with-replacement?` (Block 7) to *off* has the anticipated effect on this random sample.

[21]Under an imitative protocol, during any period in which multiple agents revise, a single sequence of trials is conducted for each agent who is either revising or a candidate for imitation, and this sequence determines the average payoff recorded for the agent during the period. Payoffs an agent obtains when he participates in another agent's sequence are not counted toward his own average payoff.

[22]These are *uniform* (uniform randomization among the highest-payoff strategies), *min* (choose the optimal strategy with the smallest index), and *stick-uniform* and *stick-min* (stick with the current strategy if it is an optimal one; otherwise apply the second criterion).

- If the `decision-method` is *logit*, then the choice probabilities are determined by applying the logit choice rule to the collection of *average* payoffs. The logit noise level is specified using the slider `log-noise-level` (see eq. (12) in Appendix A.1).[23]

- If the `decision-method` is *positive-proportional*, then the agent chooses a strategy randomly with probabilities proportional to *average* payoffs.[24] Thus, when using this decision method, the payoffs specified in `payoff-matrix` should be non-negative.

In the remaining three decision methods, the revising agent considers exactly two strategies, his current one and a single (random) alternative.[25] In all cases, payoff differences are converted into probabilities by dividing by the largest payoff difference that is possible in the game.

- If the `decision-method` is *pairwise-difference*, the agent switches to the alternative strategy only if that strategy yielded a higher average payoff than his current strategy, and in this case he switches with probability proportional to the payoff difference.

- If the `decision-method` is *linear-**, the agent switches to the alternative strategy with probability proportional to the difference between the maximum possible payoff in the game and the revising agent's average payoff (under *linear-dissatisfaction*), or between the alternative strategy's average payoff and the minimum possible payoff in the game (under *linear-attraction*).

Finally, the parameter `prob-mutation` specifies the probability with which a revising agent opts out of the process described above, instead choosing a strategy uniformly at random from the `n-of-strategies` available strategies.

2.2 Markov chains, mean dynamics, and diffusion approximations

A simulation run in *ABED* is a sample path of the Markov chain whose transition law is induced by the game, the population size, the method of assigning revision opportunities to agents, and the revision protocol. Focusing on the main qualitative options that determine the revision protocol, Table 1 distinguishes 25 trios of choices for candidate selection, matching, and decision method: 10 imitative and 15 direct. Of these 25 trios of

---

[23]When used with `candidate-selection` = *imitative* and `n-of-candidates` = 2, this decision method leads to the so-called Fermi rule (see Roca et al. (2009b), Traulsen and Hauert (2009), Perc and Szolnoki (2010), and Adami et al. (2016a)).

[24]This `decision-method` can be used to simulate frequency-dependent Moran and Wright-Fisher processes; see Section 3.7.

[25]Thus, `n-of-candidates` is automatically set to 2.

main options, there are 16 (to our knowledge) that have been explored in the literature, either in infinite-population models in the form of mean dynamics, or in finite population processes with explicitly stated revision protocols. Table 1 reports the names of these mean dynamics and processes, and provides references to papers introducing either a mean dynamic or a finite population process corresponding to the trio in question.[26]

The Appendix provides a formal framework for describing such finite-population Markov processes of the sort simulated by *ABED*. This framework is based on a decomposition of revision protocols into the three stages described above: selection of candidate strategies, determination of payoffs through matching, and choice among the candidate strategies. It also shows how mean dynamic and local diffusion approximations are constructed from the transition probabilities of the Markov chain, and derives the former from the latter in a variety of examples. In addition to providing a structured and broadly applicable framework for defining finite-population game dynamics, this presentation demonstrates how the mean dynamics appearing in the interior of Table 1 are derived from the ingredients of revision protocols that label the rows and columns.

2.3 *ABED* plots and monitors

The results of an *ABED* simulation run are presented in two pairs of plots. The upper pair (Plots I-a and I-b in Figure 1) show the time path of the population state, and the lower pair (Plots II-a and II-b in Figure 1) show the time path of each strategy's expected payoff. The first plot in each pair (labeled "recent history") shows only recent data from the simulation, while the second (labeled "complete history") shows all data starting from the beginning of the simulation. The data represented in any plot can be easily exported to a *csv* file at any time, simply by right-clicking on the plot and selecting "Export" on the contextual menu. The parameters that control plotting are contained in Block 6 of the interface.

The passage of time in the simulation is tracked in the two monitors in the upper right corner of the interface. The first, labeled `ticks`, tallies the number of periods ("ticks" in *NetLogo*) that have elapsed in the current simulation run. The second, labeled `ticks-per-second`, displays the number of periods that elapse per unit of clock time (here taken to be seconds). This number is computed from the parameters concerning the assignment of revision opportunities from Block 2 of the interface (see Section 2.1.2), with

---

[26]We should emphasize that the trios by themselves only partially describe the revision protocol. The choices of number of opponents to observe, alternate strategies to test, trials playing a strategy, noise settings, and other parameters are needed to complete the specification. The mean dynamics and processes named in Table 1 sometimes restrict these other parameters as well, so even the trios that have been studied may only have been analyzed in some instances.

the presumption that each agent expects to receive one revision opportunity per unit of clock time.[27] The conversion of ticks to seconds determines the horizontal scaling of the plots: in Plots I-a and II-a, time is measured in milliseconds, and in Plots I-b and II-b it is measured in seconds.

2.4 Using *ABED*: Modifying assumptions during runtime

To illustrate one use of *ABED*, we show how the evolution of behavior under the *pairwise-difference* decision method in the standard Rock-Paper-Scissors game is affected by switching from *imitative* to *direct* candidate selection. We initialize the simulation with the parameter settings in Table 2, and change `candidate-selection` from *imitative* to *direct* in the middle of the run. In addition to the parameter choices already mentioned, Table 2 specifies that `complete-matching?` is *on*, so that agents' experienced payoffs coincide with their expected payoffs, as traditionally assumed in models of deterministic game dynamics.[28]

Game, population size and initial state		
Game	**Initial state**	**Population size**
payoff-matrix = $\begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$	random-initial-condition? = *off* initial-condition = [800 100 100]	(n-of-agents ← 1000)

Assignment of revision opportunities	
use-prob-revision? = *off*	n-of-revisions-per-tick = 1

Revision protocol		
Candidate selection	**Matching**	**Decision method**
candidate-selection = *imitative*	complete-matching? = *on*	decision-method = *pairwise-difference*
(n-of-candidates ← 2)	(n-of-trials ← 999)	prob-mutation = 0

Table 2: Parameter settings for Figure 4. The parameter `candidate-selection` is switched from *imitative* to *direct* mid-run. Parameters stored in `rock-paper-scissors.csv`.

The mean dynamic for the initial parameter settings is the replicator dynamic (Example A.1), which in standard Rock-Paper-Scissors exhibits a continuum of closed orbits surrounding the unique Nash equilibrium at $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The later parameter settings instead generate the Smith (1984) dynamic. Under this deterministic dynamic, the Nash

---

[27]Thus, if `use-prob-revision?` is *on*, then `ticks-per-second = 1 / prob-revision`, while if `use-prob-revision?` is *off*, then `ticks-per-second = n-of-agents / n-of-revisions-per-tick`.

[28]In this table and those in Appendix II (online), we use gray text and symbol ← to denote parameter assignments that are automatically set as a consequence of other parameter values.

Figure 4: Time series of the population state in Rock-Paper-Scissors under the *pairwise-difference* decision method. `candidate-selection` is changed from *imitative* to *direct* mid-run.

equilibrium is a global attractor, although the equilibrium is approached extremely slowly once the state becomes close to x^*.[29]

Running *ABED* with the parameter values from Table 2 simulates the evolution of the population shares when 1000 agents follow the two revision protocols described above. Initially, 800 agents are assigned to Rock, 100 to Paper, and 100 to Scissors.[30] While candidate selection is imitative, the population consistently follows cycles of large amplitude among the three strategies. When candidate selection is switched to direct, the cycles quickly become much smaller in magnitude, but persist indefinitely. These behaviors are largely consistent with predictions from the mean dynamics.

2.5 Using *ABED* for Monte Carlo experiments

Besides running individual simulations using the interface, it is also possible to conduct several runs for different combinations of parameter values in a systematic fashion. This automatic exploration of the parameter space can be easily conducted using *BehaviorSpace* (Wilensky and Shargel, 2002), a software tool integrated within *NetLogo* which greatly facilitates running a model many times, systematically varying the desired parameter values, and keeping a record of the results of each run. Thus, together with *BehaviorSpace*, one can use *ABED* not only for exploratory or illustrative purposes, but

---

[29]See Hofbauer et al. (1979), Zeeman (1980), and Hofbauer and Sandholm (2009). For phase diagrams of the two mean dynamics, see Sandholm (2010b, Figure 5.3).

[30]As noted in Table 2, setting both `random-initial-condition?` = *off* and `initial-condition` = [800 100 100] causes *ABED* to automatically set `n-of-agents` to 1000.

also to conduct statistical analyses of various revision protocols for any 2-player game in finite populations. We explain how to perform Monte Carlo experiments in *ABED* using *BehaviorSpace* in Appendix I (online).

3. Examples

We now present *ABED* simulations that illustrate a variety of properties of finite-population game dynamics. The parameter specifications for each example can be found in Appendix II (online), as well as in parameter files included with the *ABED* release.

3.1 "Equilibrium" behavior in finite-population dynamics

As explained above, deterministic evolutionary dynamics are often derived by first defining finite-population stochastic models and then taking large population limits. A rest point x^* of the resulting mean dynamics is a state at which the *expected* motion of the finite-population stochastic processes equals (or converges to) zero. If multiple strategies are used at x^*, and if agents base their decisions on information obtained from limited numbers of matches, then behavior at x^* is not at rest in the finite-population processes. Instead, there is variation in the use of each strategy, with agents switching among the strategies in the support of x^*. If choice probabilities vary smoothly with the population state, then in the large population limit, this order $1/\sqrt{N}$ variation of the state around x^* can be approximated by a diffusion process (see Appendix A.3).

We now illustrate this random variation near "equilibrium points", and assess how well predictions based on the limiting diffusion approximation describe behavior in small populations. We focus on a simple model in which agents are randomly matched once in each discrete time period to play a Hawk-Dove game. One agent is randomly chosen to revise her strategy, and she does so using *imitate the best realization*: the revising agent compares the payoff he obtained in his match to the payoff a randomly selected individual obtained in her match, adopting her strategy if and only if her realized payoff was larger than his.[31] Since agents are matched just once, the payoff comparison is very simple: when a Hawk and a Dove each participate in a match, the Hawk will have the higher payoff if and only if he encountered a Dove.

In Example A.6, we show that the expected motion of this process from interior population states is always toward $x^* = (\frac{1}{2}, \frac{1}{2})$, the only interior rest point of the mean dynamic.

---

[31]The dynamics of this model are analyzed in detail by Izquierdo and Izquierdo (2013, case 5.1), who also provide convergence results of this revision protocol to its mean dynamics in general n-player games.

Figure 5: Time series of the population state in Hawk-Dove game under "imitate the best realization". (Simulation parameters: Table 1 in the online appendix; `hawk-dove.csv`.)

Population size	SD from local diffusion approximation	Empirical SD in the simulation
20	$\frac{1}{2\sqrt{20}} \approx .1118$	.1131
200	$\frac{1}{2\sqrt{200}} \approx .03536$	.0367
2000	$\frac{1}{2\sqrt{2000}} \approx .01118$	.0110

Table 3: Dispersion in behavior at an "equilibrium" state: simulations vs. diffusion approximations.

We then present the local diffusion approximation, which shows that in the large N limit, fluctuations of the process about x^* are approximately Gaussian with variance $\frac{1}{4N}$.

In the simulation run shown in Figure 5, the finite population process described above is initialized with a population size of 20. After 1000 units of clock time,[32] the population size is increased to 200; then after another 1000 units of clock time, to 2000. The figure makes plain both the random variation in the proportions using each strategy, and the fact that this variation diminishes when the population size is larger.

Table 3 compares the empirical standard deviations in the proportion of agents playing Hawk with predictions based on the local diffusion approximation. Although the approximation is only guaranteed for large population sizes, Table 3 shows that in the present example, the predictions about the level of dispersion hold up quite well for small

---

[32]Recall that a unit of clock time is defined in such a way that each agent expects to receive one revision opportunity per unit of clock time (Section 2.3). In the present example, since `n-of-revisions-per-tick = 1`, one unit of clock time corresponds to `n-of-agents` ticks.

–20–

populations.[33] The next example examines this approximation more closely.

3.2 Infinite- and finite-horizon behavior in good Rock-Paper-Scissors

To allow for less trivial mean dynamics, we now turn to games with more than two strategies. We start by considering *Rock-Paper-Scissors* games

(1)
$$A = \begin{bmatrix} 0 & -\ell & w \\ w & 0 & -\ell \\ -\ell & w & 0 \end{bmatrix},$$

where $w > 0$ represents the benefit of winning a match and $\ell > 0$ the cost of losing.

Many finite-population models have finite-horizon behavior that differs dramatically from their infinite-horizon behavior. There are formal results that describe behavior on both time scales and that specify the order of magnitude of the wait until the transition between them. As the statements of these results concern large population limits, they do not directly address what behavior should be expected for particular population sizes and other specific parameter values of interest.

As an illustration, we consider the evolution of play in a *good Rock-Paper-Scissors* game. (equation (1) with $2 = w > \ell = 1$). We assume that matching is complete, and that agents use an imitative protocol and the linear-dissatisfaction decision method.

Because the finite-population process is a finite-state Markov chain, it must ultimately converge to one of its recurrent classes, which here are the three monomorphic states. Nevertheless, the time needed to reach one of these absorbing states may be extremely long for large populations, and behavior in the interim may be better approximated by solutions to the protocol's mean dynamic, which here is the replicator dynamic (Taylor and Jonker (1978)) modulo a state-independent change of speed (see Example A.7). It is well known that state $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is an interior evolutionarily stable strategy (ESS) of good RPS. Thus classic results show that all interior solution trajectories of the replicator dynamic converge to this state; the local diffusion approximation describes the fluctuations of the process around this state. All this suggests that simulations with large finite populations will initially approach the ESS and linger around it for some time, but will ultimately converge to a monomorphic state.

---

[33]The diffusion approximation is stated in terms of a linear approximation of the mean dynamic around the rest point, and a constant approximation of the dispersion and comovements of increments based on values at the rest point of the mean dynamic. In the present example, these two approximations remain accurate at population states outside of small neighborhoods of the rest point, and this helps to account for the quality of the predictions in Table 3. See Example A.6 for details.

To examine how the population size influences the actual course of play, we run simulations of the process in *ABED* for a variety of population sizes. Figure 6 shows representative runs over 500 units of clock time, so that each agent expects to receive 500 revision opportunities. The different panels show runs with populations of 50, 100, 250, and 1000 agents. Under the smaller population sizes, an absorbing monomorphic state is reached within the allotted interval, very quickly with 50 agents, and somewhat less so with 100. This agrees with the infinite horizon prediction. For the larger population sizes, all strategies remain in use throughout over the 500 time units, with the proportions using each strategy varying randomly around the ESS. The sizes of the fluctuations are large with 250 agents, and smaller but far from negligible with 1000; more on this below.
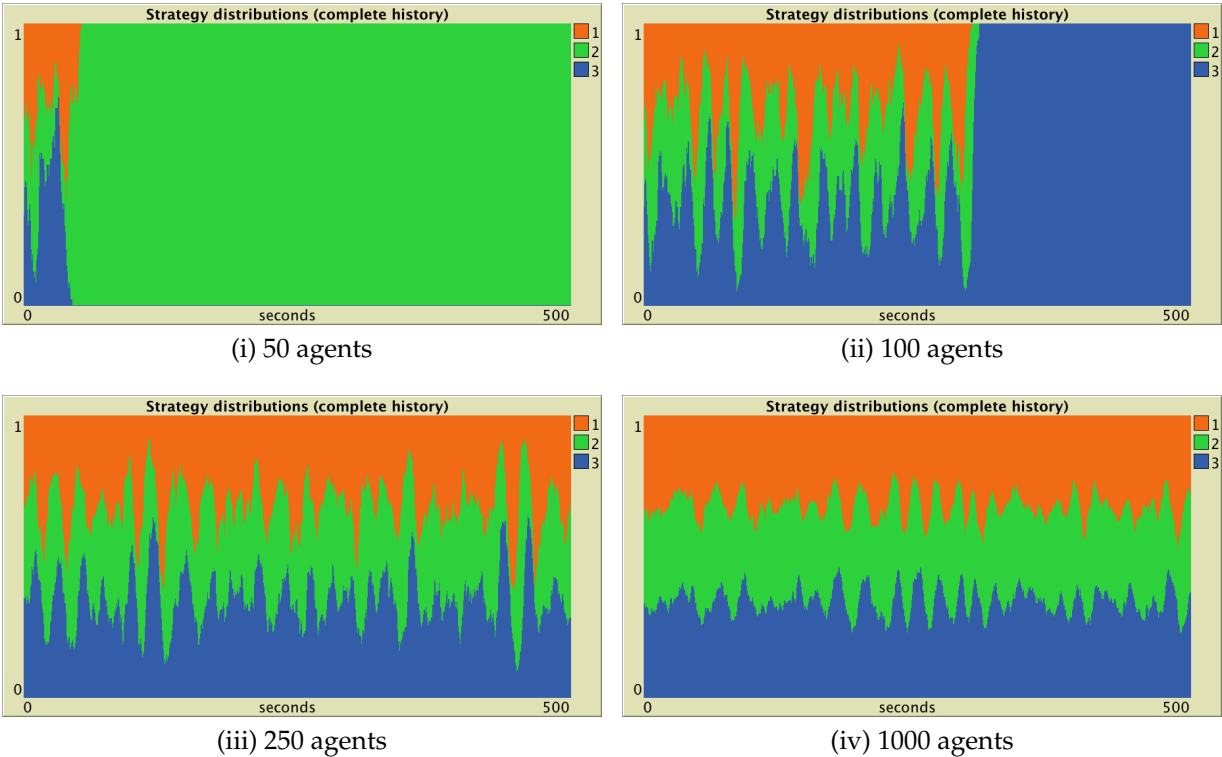

Figure 6: Imitative linear-dissatisfaction in good Rock-Paper-Scissors. (Simulation parameters: Table 2 in the online appendix, `good-rock-paper-scissors.csv`.)

The time before the infinite horizon prediction is likely to hold force can be investigated theoretically using tools from large deviations theory, which can be used to evaluate the probabilities of seeing sample paths that move against the flow of the mean dynamic. Results of Sandholm and Staudigl (2018) imply that for large population sizes N, the time required for an excursion to the boundary to occur is of an exponential order in N, with the rate of growth determined by the solution to a certain optimal control problem. To
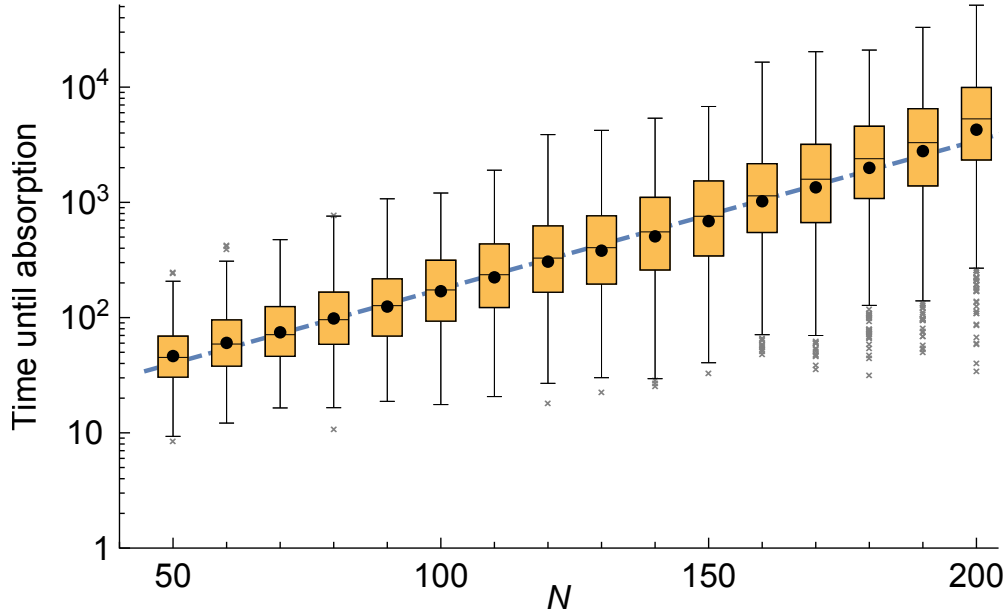
Figure 7: Time until absorption for imitative linear-dissatisfaction in good Rock-Paper-Scissors as a function of the number of agents N. The dots in each column are subpopulation means of the log time until absorption, and the dashed line is the semilog regression line.

roughly compare this prediction to simulated behavior in populations of moderate size, we ran 1000 simulations of the evolutionary process for each population size from 50 to 200, starting each from a state as close as possible to x^*, and recorded the times at which the boundary was reached. Figure 7 presents a semilog plot of the results, along with the semilog regression line. Clearly, the regression line fits the subpopulation mean points very well, and the subpopulation dispersions grow rather slowly as N increases. Thus the prediction that the time to reach the boundary grows at an exponential rate appears to be reasonably accurate even for small population sizes.

For population sizes large enough that absorption to the boundary does not occur quickly, we can describe fluctuations of the population state around the ESS using a diffusion approximation, as in Section 3.1. Example A.7 shows that for large enough population sizes, the time t distributions of the evolutionary process settle down exponentially quickly to a quasi-stationary distribution (cf. Darroch and Seneta (1965)) in which the standard deviation in the proportions of agents using each strategy is $\sqrt{20/9N} \approx 1.491 \frac{1}{\sqrt{N}}$. This prediction of dispersion agrees reasonably well with the simulation runs with 250 agents and 1000 agents from Figure 6.[34]

---

[34]For $N = 250$, the predicted standard deviation is $\sqrt{20/(9 \cdot 250)} = .0943$, compared to empirical standard deviations of .1029, .1031, and .1009 for the three strategies. For $N = 1000$, the prediction is $\sqrt{20/(9 \cdot 1000)} = .0471$, and the empirical standard deviations are .0509, .0506, and .0496.

Strategy distributions (complete history)

(i) time series

(ii) marginal empirical distributions (theoretical predictions in dashed black)
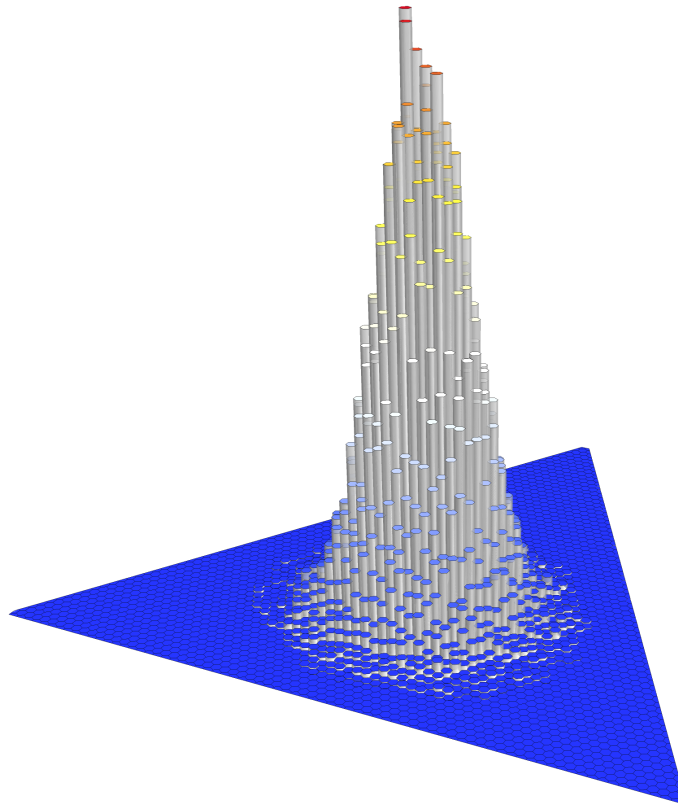
(iii) joint empirical distribution (triangular base centered at $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, boundaries $x_i \geq \frac{1}{3} - .10$)

Figure 8: Imitative linear-dissatisfaction in good Rock-Paper-Scissors, 5000 agents, 10000 units of clock time.

The diffusion approximation also provides a prediction of the full empirical distribution: namely, that it is approximately Gaussian with mean x^* and covariance matrix $\frac{10}{3N}\Phi$, where the matrix $\Phi \in \mathbb{R}^{3\times3}$, defined by $\Phi_{ii} = \frac{2}{3}$ and $\Phi_{ij} = -\frac{1}{3}$ for $i \neq j$, orthogonally projects $\mathbb{R}_0^3$ onto the tangent space of the simplex.[35] To evaluate this prediction, we simulate an evolutionary process with 5000 agents over 10000 units of clock time (Figure 8). Since $N = 5000$, we predict that the empirical distributions of the proportions of agents using each strategy should be approximately normal with mean $\frac{1}{3}$ and standard deviation $\sqrt{20/(9 \cdot 5000)} = .0211$. Figure 8(ii) shows that this prediction is accurate. Thus with 5000 agents, fluctuations in the proportions using each strategy throughout the range $\frac{1}{3} \pm .05$ are common. Figure 8(iii) shows the empirical joint distribution of strategies on the portion of the simplex with $x_i \geq \frac{1}{3} - .10$ for all strategies i. Although there is not enough data to describe the quasi-stationary distribution with great accuracy, it appears that the level sets in Figure 8(iii) are roughly spherical, which is just how the Gaussian distribution from the diffusion approximation should look: see Example A.7 for an explanation.

3.3 Imitation with complete matching and single matches

Generally speaking, the form that mean dynamics take depends on the nature of the process through which agents are matched to play the game. One important exception to this rule, noted in Table 1, arises in the case of the replicator dynamic, which appears as a mean field in imitative models both when there is a complete matching of agents in every period and when agents are matched just once each period (see Examples A.1 and A.2, and Remark A.3). Thus in the large population limit, the finite-horizon evolution of population shares under both matching regimes should be indistinguishable.[36] For moderate population sizes, however, differences in behavior could in principle be easier to discern.

To investigate this question, we use *ABED* to simulate populations of 500 agents playing *standard Rock-Paper-Scissors* (equation (1) with $w = \ell = 1$) using an imitative protocol with the *pairwise-difference* decision method. The panels of Figure 9 show the results of two simulation runs. In panel (i), payoffs were generated by complete matching; in panel (ii), payoffs were generated by matching each agent once.

In the complete-matching simulation run shown in panel (i), the population cycles around the mixed equilibrium x^* at varying distances for 1000 time units. Repeating this

---

[35]This is the source of the predicted standard deviation above: $\sqrt{\frac{20}{9N}} = \sqrt{\frac{10}{3N} \cdot \frac{2}{3}}$.

[36]However, processes with the same mean dynamics may be distinguishable by looking at magnified deviations from rest points (Appendix A.3), or at infinite-horizon behavior (Sandholm (2010b, Section 12.4.3)).

(i) complete matching	(ii) single matches

Figure 9: Standard Rock-Paper-Scissors under imitative pairwise-difference protocol, starting at `initial-condition` = [166 167 167]. (Other simulation parameters as in Table 2.)

simulation 100 times, we observed cycling among interior states 84 times. In the single-matching run shown in panel (ii), the population reaches a monomorphic state within 1000 time units; in 100 simulation runs, this behavior was observed 85 times. Because the single-matching specification includes randomness in payoff realizations, it results in a stochastic process whose increments exhibit greater dispersion around their expectations. This additional variation is the likely source of the qualitative difference in medium-run behaviors.

3.4 Best response dynamics in bad Rock-Paper-Scissors

Bad Rock-Paper-Scissors (equation (1) with $\ell > w$) is the simplest example in which a variety of standard deterministic population dynamics exhibit an attracting closed orbit in the interior of the simplex. Under the best response dynamic (Gilboa and Matsui (1991), Hofbauer (1995b)), play converges to a collection of population states called a *Shapley triangle*. When $w = 1$, the triangle's vertices are the cyclically symmetric states $\frac{1}{\ell^2+\ell+1}(\ell^2, \ell, 1)$, $\frac{1}{\ell^2+\ell+1}(\ell, 1, \ell^2)$, and $\frac{1}{\ell^2+\ell+1}(1, \ell^2, \ell)$. Motion toward and within this cycle is determined by the exponential decay in the use of suboptimal strategies: if strategy i is suboptimal over time interval $[s, t]$, then the proportion x_i playing strategy i in the limiting process satisfies $x_i(t) = e^{t-s} x_i(s)$.[37]

How much of this structure is preserved with finite population sizes? Figure 10 shows that with a population size is 100, cycling is somewhat irregular, both in the rate of decay in the use of suboptimal strategies and in the fractions using each strategy when the optimal strategy changes. But with a population size of 1000, the deterministic

---

[37]Gaunersdorfer and Hofbauer (1995) provide a comprehensive analysis of this example. For phase diagrams of a variety of dynamics in bad Rock-Paper-Scissors, see Sandholm (2015, Ex. 13.26).

Figure 10: The best response protocol in bad Rock-Paper-Scissors with $w = 1$ and $\ell = 2$. (Simulation parameters: Table 3 in the online appendix; `bad-rock-paper-scissors.csv`.)

approximation appears very accurate: the prevalence of each strategy at the moment it becomes suboptimal is nearly the same at the start of each cycle, and the abandonment of the strategy while it is suboptimal is an almost deterministic exponential decay.

3.5 Mutations and the hypercycle system

The *hypercycle system* is a differential equation model introduced by Eigen and Schuster (1979) to model cyclical catalysis in mixtures of polynucleotides. It also arises as the replicator dynamic for the *hypercycle game*, which has a simple cyclical payoff structure—specifically, one in which each strategy yields a payoff of 1 against the preceding strategy (modulo n (= `n-of-strategies`)) and a payoff of 0 against the others (Schuster and Sigmund (1983)). When $n \geq 5$, the unique Nash equilibrium $x^* = (\frac{1}{n}, \ldots, \frac{1}{n})$ and the boundary of the simplex are both repelling (Schuster et al. (1978), Hofbauer et al. (1981)), and the dynamics admit a stable periodic orbit (Hofbauer et al. (1991)). Along this orbit, each strategy becomes very scarce at its low ebb, so much so that it might vanish entirely if the population were of moderate size.

With this question in mind, we use *ABED* to simulate a finite-population imitative process whose mean dynamic is the hypercycle system with $n = 5$, but with the addition of rare mutations to prevent strategies from going extinct. Figure 11(i) shows a typical simulation run for a population of size 1000 and a mutation probability of 10^{-3}. After random variation moves the population away from the Nash equilibrium, play follows a path consistent with the hypercycle system, with the use of each strategy varying from about 800 agents to no agents at all. Mutations prevent the latter situation from persisting, and indeed the cycling in Figure 11(i) looks almost completely regular. If the mutation probability is reduced to 10^{-4}, then extinctions persist for long enough that

(i) 1000 agents, `prob-mutation` = 10^{-3} (ii) 1000 agents, `prob-mutation` = 10^{-4}
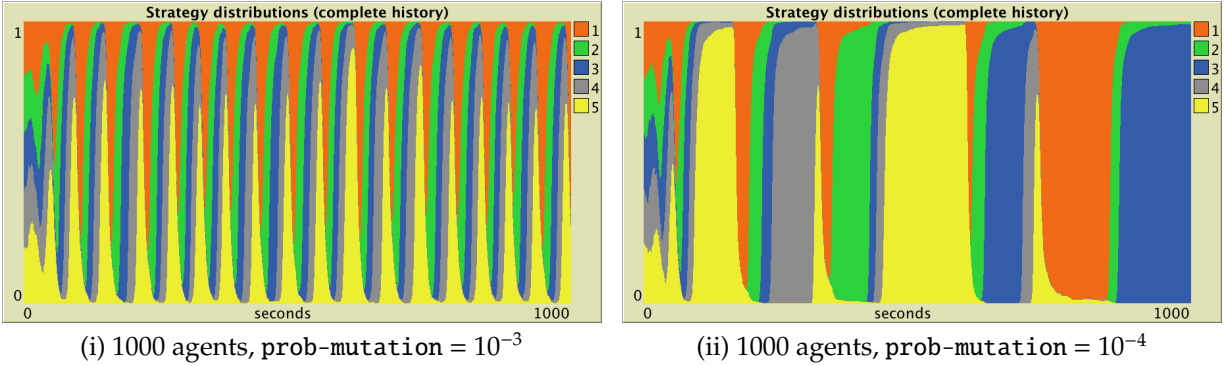
Figure 11: The five-strategy hypercycle game under imitation/pairwise-difference with mutations. (Simulation parameters: Table 4 in the online appendix; `hypercycle-5.csv`.)

regular cycling is destroyed. The population still passes through regimes in which each of the five strategies predominates, but the stochastic process no longer resembles the classic hypercycle system.

3.6 Random-walk tie breaking and setwise stability

In evolutionary game models in which agents optimize exactly, one must specify what they do when choosing among multiple optimal strategies. Most stochastic models posit that the evolutionary process is Markov with respect to the population state. When the population is large, large-numbers arguments then imply that the proportions of revising agents choosing each of the optimal strategies are fixed, leading to motion in a nearly deterministic direction. In other words, combining the usual Markov assumption with large population sizes creates strong implicit restrictions on adjustments in regions with multiple best responses.

An important possibility that these restrictions rule out is wandering of the population state around and possibly out of a component of Nash equilibria, a phenomenon with deep connections to set-valued versions of the ESS concept. Building on work of Thomas (1985) and Balkenborg and Schlag (2001), van Veelen (2012) characterizes sets that are evolutionarily stable as those that are *robust against indirect invasions* (*RAII*). Roughly speaking, this definition requires a set of states to be mutually accessible through sequences of neutral mutations, but not to be susceptible to mutations that lead out of the set and then outperform the incumbents in the post-entry state.

To allow for wandering within a Nash equilibrium component, *ABED* includes the option of *random-walk* tie breaking when `decision-method` = *best*. As explained in Appendix A.4, this option initiates an auxiliary random walk that runs on a discrete set of interior

Figure 12: Wandering out of a Nash component in van Veelen's game (2) under imitation/best with random-walk tie-breaking, 100 agents, and mutation probability 10^{-3}. (Simulation parameters: Table 5 in the online appendix; `random-walk-tie-breaking.csv`.)

population states and whose stationary distribution is uniform on this set. The value of the auxiliary variable determines the relative probabilities with which currently optimal strategies are chosen.

As an illustration, we consider the following game of van Veelen (2012, Example 4):

$$(2) \qquad A = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix}.$$

In this game, any state at which only strategies 1 and 2 are used is neutrally stable (Maynard Smith (1982)). If the population reaches the state at which only population 2 is used, then strategy 3 can invade as a neutral mutant. Under imitative protocols, the population may then wander among states at which only strategies 2 and 3 are used, and may reach the state at which only strategy 3 is used. There strategy 4 is also a best response, and a mutation introducing it can quickly lead strategy 4 to predominate.

Figure 12 presents a typical simulation run of a process in which agents are recurrently completely matched to play game (2). When revising, an agent observes his own payoff and that of a randomly chosen opponent, adopting the strategy whose payoff is higher, and resolving ties using the *random-walk* tie-breaker. Because the protocol is imitative, strategies that go extinct can only be reintroduced via mutations, allowing the stepping-stone path described in the previous paragraph to be traversed. Figure 12 shows that
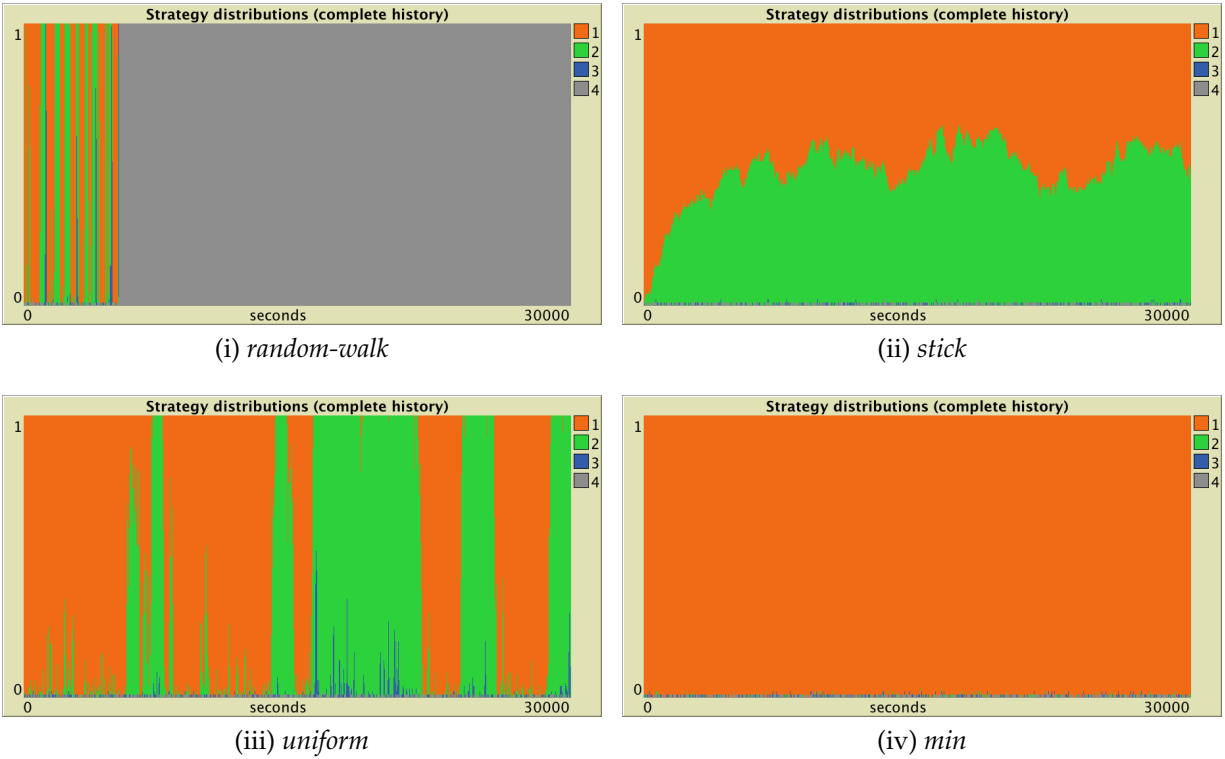
	(i) *random-walk*	(ii) *stick*

	(iii) *uniform*	(iv) *min*

Figure 13: Alternate tie-breaking rules in van Veelen's game (2) under imitation/best.

starting from the state at which all agents play strategy 1, the state at which all agents play strategy 4 is reached at the time that each agent has had approximately 2000 revision opportunities. Figure 13 shows simulation runs of other processes that differ only in the choice of tie-breaking rule. While the sequence of events that allows the initial equilibrium component to be escaped is always possible, without persistence in tie-breaking probabilities it is unlikely to occur within a reasonable amount of time.

3.7 The frequency-dependent Moran process

The frequency-dependent Moran process (Nowak et al. (2004); Taylor et al. (2004)) is a finite-population stochastic process named for Moran (1958, 1962) and commonly studied in the biology literature.[38] During each time step of this process, one individual is randomly selected to be cloned, with probability proportional to payoffs from a complete matching. The newborn replaces another randomly selected individual, thus keeping the population size constant.[39] Moran processes are often studied under *weak selection*, with

---

[38]See Nowak (2006) and Traulsen and Hauert (2009) for overviews and references.

[39]Closely related is the frequency-dependent Wright-Fisher process (Imhof and Nowak (2006)), which uses the same updating rule, but with synchronized reproduction (i.e., non-overlapping generations). As

x

the payoffs of an *i* player facing a *j* player set to $1 - w + wA_{ij}$ for some small intensity of selection $w > 0$. Analyses of these processes commonly focus on very long run behavior, often in the limit as the population size becomes large or the mutation rate becomes small.[40]

ABED can be used to examine the behavior of the Moran process away from these limits. We describe the appropriate parameter settings in Example A.4 of Appendix A.2, and show there that the form of the mean dynamic of the Moran process depends on whether standard or weak selection is assumed. In the former case, one obtains the Maynard Smith replicator dynamics (Maynard Smith (1982)), while the latter yields the imitative logit dynamics (Weibull (1995)) with a large noise level on a slightly modified payoff matrix.

3.8 Sample best response vs. best experienced payoff in Centipede

To illustrate how stochastic evolutionary dynamics in two-player asymmetric games can be simulated with `ABED-2pop`, we consider the Centipede game of Rosenthal (1981). This extensive form game has ten decision nodes. Player 1 decides at odd nodes while player 2 decides at even nodes. Players' choices at each node are to Stop or to Continue. The payoffs obtained by each player at every possible outcome are shown in Figure 14. For each player, strategy *i* is the plan to Continue for $i - 1$ decision nodes and stop at his *i*th decision node. Strategy 6 for both players is the plan to always Continue.


Figure 14: The Centipede game (Rosenthal (1981)).

Here we explore the behavior of populations matched to play this game under two direct protocols based on selecting best-performing strategies, where performances are computed using samples of size 1 (i.e. `n-of-trials = 1`). Under a *sample best response protocol* (Sandholm (2001a), Kosfeld et al. (2002), Oyama et al. (2015)), a revising agent observes the strategy of a single opponent and plays a best response. Under a *best experienced payoff protocol* (Osborne and Rubinstein (1998), Sethi (2000), Sandholm et al.

---

noted in Section 2.1.2, this process can be simulated in *ABED* by setting `prob-revision = 1`.

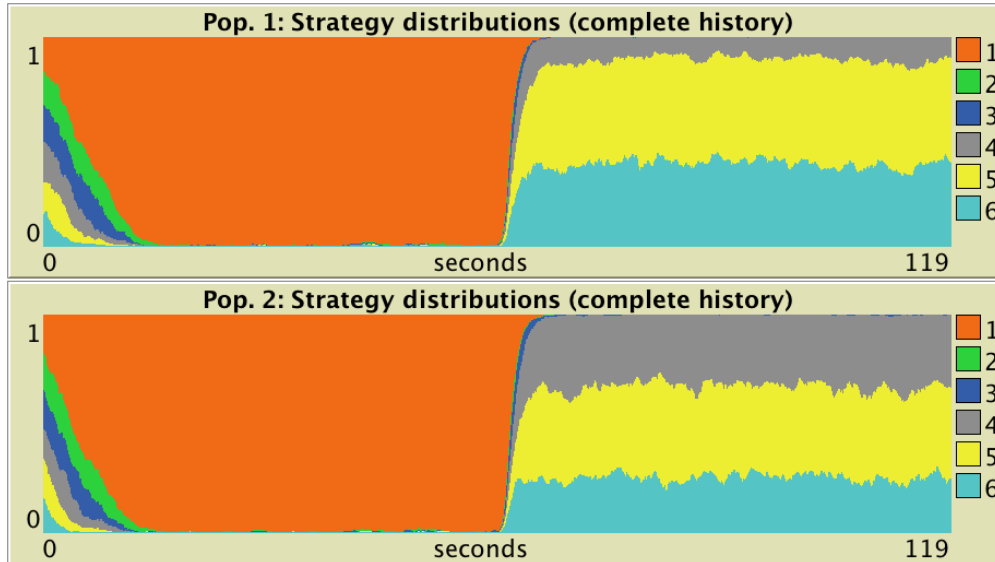[40]See Fudenberg et al. (2006), Imhof et al. (2005), Imhof and Nowak (2006), and Fudenberg and Imhof (2006, 2008).

Figure 15: Best response to samples in Centipede, with `single-sample?` switched from *on* to *off* mid-run. (Simulation parameters: Table 6 in the online appendix; `centipede-10-nodes-sampleBR-min.csv`.)

(2019, 2018)), a revising agent plays each of his strategies once against a single randomly-drawn opponent and proceeds by playing the strategy that performed best. In both cases, ties are broken in favor of the least cooperative strategy.

Figure 15 shows the time series of the strategy shares in a representative run of `ABED-2pop`, with agents initially using a sample best response protocol, and then, starting halfway through the run, a best experienced payoff protocol. The parameter settings are presented in Table 1. Switching between the two protocols is accomplished by changing one parameter: `single-sample?`, from *on* to *off*. Figure 15 shows that this change has a dramatic effect on the level of cooperation in Centipede. When agents play a best response to a simple estimate of aggregate behavior in the opposing population, play rarely moves beyond the initial decision node. If instead agents choose the strategy that performed best during testing against randomly chosen opponents, play is instead concentrated on the last few nodes of the game.

This simulation was among the first carried out by the authors using *ABED*, and we did not anticipate its outcome. Taking these simulation results as our starting point, we were later able to establish the main properties observed in the simulation analytically, and to show that these properties are robust to many variations in how agents make decisions (Sandholm et al. (2019)). We hope that other researchers will also put *ABED* to work both as an exploratory tool and as a platform for robustness testing.

4. Concluding remarks

Evolutionary game theory provides a powerful set of analytical tools for studying behavior in populations of myopic, strategically interacting agents, tools which have been deployed in wide-ranging applications.[41] As with any formal theory, progress in evolutionary game theory has relied on assumptions and approximations that make its models tractable. The examples presented above illustrate how *ABED* can be used to see how theoretical predictions stand up in settings where these assumptions are relaxed—for instance, when populations are not especially large, when matching of agents is incomplete, and when agents' information about strategies' performances comes from their own and others' direct experiences. By encompassing a great variety of specifications of revision processes within a single framework, *ABED* not only facilitates comparisons between different models from the literature, but also provides powerful tools for exploring formulations of agents' choice processes.

Because *ABED* is thoroughly documented and runs in the user-friendly *NetLogo* platform, it is not overly difficult to add new parameters, new protocols, or more elaborate alterations (including agent heterogeneity) by directly modifying the *ABED* code. Future work will substantially expand *ABED*'s capabilities, and we gladly accept implementation requests directed towards this end.

Appendix. A framework for finite-population evolutionary game dynamics

In this appendix, we present a framework for specifying finite-population evolutionary game dynamics and use it to formally describe the Markov chains simulated in *ABED*. The main novelty is a general definition of revision protocols that decomposes them into three stages: candidate selection, matching and determination of payoffs, and choices of new strategies (Section A.1). The description of Markov chains covers a range of parameter choices wide enough to allow all of the combinations presented in Table 1, but for simplicity certain parameter choices are assumed fixed. We also review results on deterministic (Section A.2) and diffusion (Section A.3) approximations for these Markov chains and exhibit examples of each. Finally, Section A.4 presents the random-walk tie-breaker mentioned in Section 2.1.3 and applied in Section 3.6.

---

[41]For a thorough survey of recent applications, see Newton (2018).

A.1 The Markov chain

When the agents in the population play a game with strategy set $S = \{1, \ldots, n\}$, strategy distributions are elements of the simplex $X = \{x \in \mathbb{R}^n_+ : \sum_{i=1}^n x_i = 1\}$, whose ith vertex is denoted e_i. If the population is of size N (= n-of-agents), the population state is an element of the grid $\mathcal{X}^N = \{x \in X : Nx_i \in \mathbb{Z} \text{ for all } i\}$.

We define the Markov chain as a discrete-time process on the set of population states $\mathcal{X}^N$.[42] We assume during each period of duration $\frac{1}{N}$, exactly one agent receives a revision opportunity (i.e., use-prob-revision? = *off* and n-of-revisions-per-tick = 1). The transition law of the Markov chain $\{X_t^N\}$ is then described by

(3) $\mathbb{P}\left(X_{k+1/N}^N = x + \frac{1}{N}(e_j - e_i) \,\middle|\, X_k^N = x\right) = x_i \rho_{ij}^A(x).$

For an i player to switch to strategy j during the next period, the player who receives the next revision opportunity must be an i player, which happens with probability x_i, and he must switch to strategy j, which happens with probability $\rho_{ij}^A(x)$. The *revision protocol* ρ captures both the information that agents have when revising due to matching, and the decision rule they employ. Choice probabilities under the revision protocol depend on the game $A \in \mathbb{R}^{n \times n}$ (= payoff-matrix) being played and the population state x. In general, choice probabilities in *ABED* can depend directly on the population size N, but in what follows we will focus on parameter values that preclude such dependence.

We first provide a general formulation (4) of revision protocols, and then explain each component of (4) in detail.

Let c (= n-of-candidates) denote the number of observations about strategies' performances that a revising agent obtains.[43] The total information a revising agent obtains is described by a *record* (s, π), a $c \times 2$ matrix whose first column s is the *strategy record* and whose second column π is the *payoff record*. The hth row of the record is a pair (s_h, π_h) that describes the strategy played and the payoff obtained in the revising agent's hth observation.

Revision protocols implemented in *ABED* can be expressed in the form

(4) $\rho_{ij}^A(x) = \sum_{(s,\pi)} p_i(s \,|\, x) \, q(\pi \,|\, s, x, A) \, \sigma_{ij}(s, \pi).$

The summand in (4) expresses the probability that a revising strategy i player obtains

---

[42]This description does not correspond to the agent-based specification of the process within *ABED*'s code.

[43]Here a "performance" is an average payoff generated by the matching procedure in place.

record (s, π) and switches to strategy j. Summing this over all payoff records yields the probability that a revising strategy i player switches to j.

The summand in (4) is the product of three terms. The first, $p_i(s \mid x)$, describes candidate selection. It is the probability that a revising agent playing strategy i obtains strategy record s. This probability will depend on the population state x when the protocol is imitative, but not when it is direct.

The second term, $q(\pi \mid s, x, A)$ describes matching and sampling. It is the probability that a revising agent with strategy record s obtains corresponding payoff record π. This probability depends on the population state x and on the payoff matrix A.

The third term, $\sigma_{ij}(s, \pi)$, describes the decision method. It is the probability that a revising strategy i player with record (s, π) chooses strategy j. This probability only depends on x and A by way of the record (s, π).[44]

Candidate selection and strategy records

In constructing a strategy record $s = \{s_h\}_{h=1}^c$ for a revising agent playing strategy i, we let $s_1 = i$ be the revising agent's strategy, and view subsequent strategies as being added sequentially.[45]

Under a *direct protocol*, strategies $s_2, \ldots, s_c$ are drawn without replacement from $S \setminus \{i\}$. Thus the probability of any strategy record with $s_1 = i$ and all strategies distinct is the same, and equal to the inverse of the number of such records:

$$(5) \qquad p_i(s \mid x) = \frac{(n-c)!}{(n-1)!}.$$

Thus when $c = 2$, the $n - 1$ feasible strategy records are of the form (i, j) with $j \neq i$, and each has probability $\frac{1}{n-1}$. When $c = n$, the $(n-1)!$ feasible strategy records start with i and contain all n strategies, and each has probability $\frac{1}{(n-1)!}$

How an *imitative* protocol and the population state x determine the probabilities of strategy lists depends on the values of the auxiliary parameters `imitatees-with-replacement?` and `consider-imitating-self?`. If both are *on*, then the probability of a strategy record s with $s_1 = i$ is

---

[44] A minor exception is noted below.

[45] The ordering of the pairs (s_h, π_h) is irrelevant in subsequent parts of the protocol. One could alternatively express revision protocols directly in terms of classes of records that are equivalent up to permutation. Then instead of there being distinct summands for equivalent records (see Examples A.4 and A.5), larger combinatorial terms would appear in the function p. On balance the approach we follow here is notationally simpler.

(6) $\qquad p_i(s\,|\,x) = \displaystyle\prod_{h=2}^{c} x_{s_h}.$

If both parameters are *off* and $c = N$, then every strategy record $s \in S^N$ with $s_1 = i$ and with Nx_j entries equal to j for all $j \in S$ has equal probability:

(7) $\qquad p_i(s\,|\,x) = \dfrac{(Nx_i - 1)!\,\prod_{k\neq i}(Nx_k)!}{(N-1)!}.$

Matching and sampling parameters and payoff records

Given a strategy record $s = \{s_h\}_{h=1}^c$, a population state x, and the payoff matrix A, the matching and sampling parameters determine the probabilities of each payoff record $\pi = \{\pi_h\}_{h=1}^c$. To rule out finite-population effects we assume here that `self-matching?` is *on*.

If matching is complete, then each agent gets the expected payoff of the strategy he plays:

(8) $\qquad q\left(((Ax)_{s_1}, \ldots (Ax)_{s_c})\,|\,s, x, A\right) = 1.$

For cases where matching is not complete, we assume here for simplicity that `n-of-trials` $= 1$. If the protocol is imitative, or if it is direct and `single-sample?` is *off*, then the payoff associated with each strategy in the record is determined by a distinct random match. If we assume for convenience that elements within each row of A are distinct,[46] then

(9) $\qquad q\left((A_{s_1 t_1}, \ldots A_{s_c t_c})\,|\,s, x, A\right) = \displaystyle\prod_{h=1}^{c} x_{t_h}.$

If instead the protocol is direct and `single-sample?` is *on*, then all strategies are evaluated by taking expectations with respect to the empirical distribution of the single sample. The case with `n-of-trials` $= 1$ is very simple:[47] if A has distinct rows, then

(10) $\qquad q\left((A_{s_1 k}, \ldots A_{s_c k})\,|\,s, x, A\right) = x_k.$

---

[46]If rows may have repeated elements, we must account for the fact that the payoff value obtained by an agent playing strategy j does not uniquely determine the strategy of his match partner. See Example A.2.

[47]For an example with larger numbers of trials, see Oyama et al. (2015).

Decision methods

Each decision method defines a collection of functions σ_{ij}, where $\sigma_{ij}(s, \pi)$ is the probability that a revising i player with record (s, π) chooses strategy j. Under `decision-method` = *best*,

$$\text{(11)} \qquad \text{if } \left[\pi_h = \max_{h'} \pi_{h'} \implies s_h = j \right], \text{ then } \sigma_{ij}(s, \pi) = 1.$$

If the antecedent condition fails (i.e., if the maximum payoff is achieved by more than one strategy in the record), then choice probabilities are determined by the value of `tie-breaker`, and may condition on the revising agent's strategy i.

Under `decision-method` = *logit*,

$$\text{(12)} \qquad \sigma_{ij}(s, \pi) = \frac{\sum_{h:\, s_h=j} \exp(\eta^{-1}\pi_h)}{\sum_{k \in S} \sum_{h:\, s_h=k} \exp(\eta^{-1}\pi_h)},$$

where the noise level η is set using *ABED* parameter `log-noise-level` = $\log_{10} \eta$. Similarly, `decision-method` = *positive-proportional* is defined for positive payoff records π by

$$\text{(13)} \qquad \sigma_{ij}(s, \pi) = \frac{\sum_{h:\, s_h=j} \pi_h}{\sum_{k \in S} \sum_{h:\, s_h=k} \pi_h}.$$

Under the remaining decision methods, $c = 2$, $s_1 = i$ is the revising agent's strategy, and if $s_2 = i$, then $\sigma_{ii}(s, \pi) = 1$. Let $A^{\max}$ and $A^{\min}$ be the largest and smallest payoff entries in A, and let $\bar{\Delta}A = A^{\max} - A^{\min}$. If `decision-method` = *linear-dissatisfaction*, then

$$\text{(14)} \qquad \sigma_{ij}((i, \pi_1), (j, \pi_2)) = \frac{A^{\max} - \pi_1}{\bar{\Delta}A} \text{ and } \sigma_{ii}((i, \pi_1), (j, \pi_2)) = 1 - \frac{A^{\max} - \pi_1}{\bar{\Delta}A}.$$

If `decision-method` = *linear-attraction*, then

$$\text{(15)} \qquad \sigma_{ij}((i, \pi_1), (j, \pi_2)) = \frac{\pi_2 - A^{\min}}{\bar{\Delta}A} \text{ and } \sigma_{ii}((i, \pi_1), (j, \pi_2)) = 1 - \frac{\pi_2 - A^{\min}}{\bar{\Delta}A}.$$

And if `decision-method` = *pairwise-difference*, then

$$\text{(16)} \qquad \sigma_{ij}((i, \pi_1), (j, \pi_2)) = \frac{[\pi_2 - \pi_1]_+}{\bar{\Delta}A} \text{ and } \sigma_{ii}((i, \pi_1), (j, \pi_2)) = 1 - \frac{[\pi_2 - \pi_1]_+}{\bar{\Delta}A}.$$

A.2 Finite-horizon deterministic approximation: the mean dynamic

The mean dynamic of the Markov chain described by (3) is

$$(17) \qquad \dot{x}_i = V_i^{\rho,A}(x) = \sum_{j \in S} x_j \rho_{ji}^A(x) - x_i.$$

Results of Benaïm and Weibull (2003) (for Lipschitz continuous dynamics) and Roth and Sandholm (2013) (for differential inclusions) imply that over finite time spans, sample paths are very likely to be closely approximated by solutions to (17) if the population size N is large enough. Next, we present four revision protocols and their mean dynamics.

Example A.1. The replicator dynamic from complete matching. Suppose candidate selection is imitative as in (6) with $c = 2$, matching is complete (8), and the decision method is *pairwise-difference* (16). To derive the mean dynamics for protocols using *pairwise-difference*, it is convenient to express the formula (17) for the mean dynamic without the diagonal terms ρ_{ii} of the revision protocol:

$$(18) \qquad \dot{x}_i = V_i^{\rho,A}(x) = \sum_{j \neq i} \left(x_j \rho_{ji}^A(x) - x_i \rho_{ij}^A(x) \right).$$

The off-diagonal terms of the revision protocol are

$$\begin{aligned}
\rho_{ij}^A(x) &= \sum_{(s,\pi)} p_i(s \mid x)\, q(\pi \mid s, x, A)\, \sigma_{ij}(s, \pi) \\
&= x_j \times 1 \times \sigma_{ij}((i, (Ax)_i), (j, (Ax)_j)) \\
&= x_j \frac{[(Ax)_j - (Ax)_i]_+}{\bar{\Delta}A},
\end{aligned}$$

where the second equality uses the fact that the only strategy record that allows a revising i player to switch to j is $(s_1, s_2) = (i, j)$. Using (18), we compute the mean dynamic as

$$\begin{aligned}
\dot{x}_i &= \sum_{j \neq i} \left(x_j \rho_{ji}^A(x) - x_i \rho_{ij}^A(x) \right) \\
&= \sum_{j \neq i} \left(x_j \left(x_i \frac{[(Ax)_i - (Ax)_j]_+}{\bar{\Delta}A} \right) - x_i \left(x_j \frac{[(Ax)_j - (Ax)_i]_+}{\bar{\Delta}A} \right) \right) \\
&= \frac{1}{\bar{\Delta}A} x_i \sum_{j \neq i} x_j ((Ax)_i - (Ax)_j) \\
&= \frac{1}{\bar{\Delta}A} x_i \sum_{j \in S} x_j ((Ax)_i - (Ax)_j) \\
&= \frac{1}{\bar{\Delta}A} x_i ((Ax)_i - x'Ax).
\end{aligned}$$

This is the replicator dynamic, sped up by a constant factor of $\frac{1}{\bar{\Delta}A}$. ◆

Example A.2. The replicator dynamic from single matches. Repeat the setup from Example A.1, but with complete matching replaced by limited matching with `n-of-trials = 1`, as in (9). Then for $j \neq i$,

$$
\begin{aligned}
\rho_{ij}^A(x) &= \sum_{(s,\pi)} p_i(s\,|\,x)\, q(\pi\,|\,s,x,A)\, \sigma_{ij}(s,\pi) \\
&= \sum_{k\in S} \sum_{\ell\in S} x_j \times x_k x_\ell \times \sigma_{ij}((i,A_{ik}),(j,A_{j\ell})) \\
&= x_j \sum_{k\in S} \sum_{\ell\in S} x_k x_\ell \frac{[A_{j\ell} - A_{ik}]_+}{\bar{\Delta} A},
\end{aligned}
$$

where the sums are over the strategies of the match partners of the revising i player and the j player that he observes. The mean dynamic is

$$
\begin{aligned}
\dot{x}_i &= \sum_{j\neq i} \left(x_j \rho_{ji}^A(x) - x_i \rho_{ij}^A(x) \right) \\
&= \sum_{j\neq i} \left(x_j \left(x_i \sum_{k\in S} \sum_{\ell\in S} x_k x_\ell \frac{[A_{ik} - A_{j\ell}]_+}{\bar{\Delta} A} \right) - x_i \left(x_j \sum_{k\in S} \sum_{\ell\in S} x_k x_\ell \frac{[A_{j\ell} - A_{ik}]_+}{\bar{\Delta} A} \right) \right) \\
&= \frac{1}{\bar{\Delta} A} x_i \sum_{j\neq i} x_j \sum_{k\in S} \sum_{\ell\in S} x_k x_\ell (A_{ik} - A_{j\ell}) \\
&= \frac{1}{\bar{\Delta} A} x_i \sum_{j\neq i} x_j ((Ax)_i - (Ax)_j) \\
&= \frac{1}{\bar{\Delta} A} x_i ((Ax)_i - x'Ax).
\end{aligned}
$$

Thus the mean dynamic is again the replicator dynamic, as in the previous example. ◆

Remark A.3. Calculations similar to those from Examples A.1 and A.2 show that imitation using the decision methods *linear-dissatisfaction* and *linear-attraction* also generate the replicator dynamic under both complete matching and single matches, up to a state-independent change of speed; compare Example A.7.

Example A.4. The Maynard Smith replicator dynamic, the imitative logit dynamic, and frequency-dependent Moran processes. Suppose candidate selection is imitative with $c = N$ and without replacement, as in (7), that matching is complete (8), and that the decision method is *positive-proportional* (13) (with the entries of A assumed to be positive). Let $S^N(x,i)$ denote the set of strategy records $s \in S^N$ with $s_1 = i$ and $\#\{h: s_h = k\} = Nx_k$ for all $k \in S$. Since matching is complete, one of these records must occur, and the probability $p_i(s\,|\,x)$ of each such record is the same (and given by (7)), implying that $p_i(s\,|\,x) = 1/\#S^N(x,i)$ for

all $s \in S^N(x, i)$. We can thus compute the revision protocol as

$$\rho_{ij}^A(x) = \sum_{(s,\pi)} p_i(s \mid x) \, q(\pi \mid s, x, A) \, \sigma_{ij}(s, \pi)$$

$$= \sum_{s \in S^N(x,i)} p_i(s \mid x) \times 1 \times \frac{\sum_{h:\, s_h=j} (Ax)_j}{\sum_{k \in S} \sum_{h:\, s_h=k} (Ax)_k}$$

(19)
$$= \frac{x_j (Ax)_j}{\sum_{k \in S} x_k (Ax)_k}.$$

Using (17), we compute the mean dynamic:

(20) $$\dot{x}_i = \sum_{j \in S} x_j \rho_{ji}^A(x) - x_i = \frac{x_i (Ax)_i}{\sum_{k \in S} x_k (Ax)_k} - x_i = \frac{x_i((Ax)_i - x'Ax)}{x'Ax}.$$

(20) is the Maynard Smith replicator dynamic.

If the only change we make to the above is to use the *logit* decision method with noise level $\eta \in (0, \infty)$, then the revision protocol and the mean dynamic become

(21) $$\rho_{ij}^A(x) = \frac{\sum_{h:\, s_h=j} \exp(\eta^{-1}(Ax)_j)}{\sum_{k \in S} \sum_{h:\, s_h=k} \exp(\eta^{-1}(Ax)_k)} = \frac{x_j \exp(\eta^{-1}(Ax)_j)}{\sum_{k \in S} x_k \exp(\eta^{-1}(Ax)_k)} \quad \text{and}$$

(22) $$\dot{x}_i = \frac{x_i \exp(\eta^{-1}(Ax)_i)}{\sum_{k \in S} x_k \exp(\eta^{-1}(Ax)_k)} - x_i.$$

(22) is the imitative logit dynamic (Weibull (1995)).

In the biology literature, the stochastic evolutionary process induced by revision protocol (19) is called the frequency-dependent Moran process (Nowak (2006)). This process is typically studied in the weak selection limit, meaning that the payoff matrix with entries A_{ij} is replaced with one with entries $1 - w + wA_{ij}$, where the intensity of selection $w > 0$ is taken to 0. To study the weak selection limit in the frequency-dependent Moran process (19), one can also use the *logit* decision method (21) with a large choice of noise level η and modified payoffs $A_{ij} - 1$. A large value of η means that η^{-1} is close to zero, so it follows from the Taylor approximation $\exp(y) \approx 1 + y$ for $y \approx 0$ that doing this leads to an excellent approximation of the frequency-dependent Moran process with $w = \eta^{-1}$. ◆

Example A.5. Best experienced payoff dynamics. Suppose candidate selection is direct with $c = n$ as in (5), that matching is based on limited multiple samples with n-of-trials = 1, as in (9), and that the decision method is *best*, as in (11). To derive the revision protocol, we introduce the following notations: Let $S^n(i)$ denote the set of strategy records $s \in S^n$ that are a permutation of the elements of $S = \{1, \ldots, n\}$ with $s_1 = i$. For $s, t \in S^n$ and payoff

matrix A, define the payoff record $\check{\pi}(s, t, A) \in \mathbb{R}^n$ by $\check{\pi}_h(s, t, A) = A_{s_h t_h}$. Define $s^\circ \in S^n$ by $s_h^\circ = h$ for all h. And let $S_*^n(j, A) = \{t \in S^n : A_{j t_j} > A_{k t_k} \text{ for all } k \in S\}$.

   If all entries of A are distinct we can derive the revision protocol as follows:

$$\rho_{ij}^A(x) = \sum_{(s,\pi)} p_i(s \mid x)\, q(\pi \mid s, x, A)\, \sigma_{ij}(s, \pi)$$

$$= \sum_{s \in S^n(i)} \sum_{t \in S^n} \frac{1}{(n-1)!} \times \prod_{h=1}^{n} x_{t_h} \times \sigma_{ij}(s, \check{\pi}(s, t, A))$$

$$= \sum_{t \in S^n} \prod_{h=1}^{n} x_{t_h} \sigma_{ij}(s^\circ, \check{\pi}(s^\circ, t, A))$$

$$= \sum_{t \in S_*^n(j, A)} \prod_{h=1}^{n} x_{t_h}.$$

The third equality uses the invariance of σ_{ij} to the arrangement of the pairs in (s, π) and the fact that the sum is taken over all t in S^n. The mean dynamic that this protocol generates,

$$\dot{x}_i = \sum_{t \in S_*^n(i, A)} \prod_{h=1}^{n} x_{t_h} - x_i,$$

is an instance of a best experienced payoff dynamic. ◆

A.3 Diffusion approximation near equilibria

Suppose that x^* is a rest point of the mean dynamic (17): $V^{\rho, A}(x^*) = \mathbf{0}$. We define the *local behavior process* near x^*, denoted $\{Z_t^N\}_{t \geq 0}$, by

$$(23) \qquad Z_t^N = \sqrt{N}(\bar{X}_t^N - x^*)$$

where $\{\bar{X}_t^N\}_{t \geq 0}$, defined by $\bar{X}_t^N = X_{\lfloor Nt \rfloor / N}^N$ is the piecewise constant interpolation of $\{X_t^N\}_{t \in \{0, 1/N, \ldots\}}$, and where $X_0^N \in \mathcal{X}^N$ converges to x^* at rate $o(1/\sqrt{N})$. Let $DV_i^{\rho, A}(x^*)$ denote the Jacobian of $V^{\rho, A}$ evaluated at x^*, and define the *local covariance matrix* $\Sigma^* \in \mathbb{R}^{n \times n}$ by

$$(24) \qquad \Sigma_{ij}^* = \begin{cases} -\left(x_i^* \rho_{ij}^A(x^*) + x_j^* \rho_{ji}^A(x^*) \right) & \text{if } i \neq j, \\ \sum_{k \neq i} \left(x_i^* \rho_{ik}^A(x^*) + x_k^* \rho_{ki}^A(x^*) \right) & \text{if } i = j. \end{cases}$$

Now consider the stochastic differential equation

(25) $$dZ_t = DV^{\rho,A}(x^*)Z_t\,dt + \sqrt{\Sigma^*}\,dB_t,$$

where $\sqrt{\Sigma^*}$ is the symmetric positive semidefinite square root of Σ^*, and B_t denotes an n-dimensional Brownian motion. Applying results of Stroock and Varadhan (1979), Sandholm (2003) shows that over finite time spans, the local behavior process (23) converges in distribution (with respect to the Skorokhod topology) to the solution to (25) with initial condition $\mathbf{0}$ as the population size N grows large.[48]

When the process defined by (25) has a limiting stationary distribution, this distribution describes the persistent order $\frac{1}{\sqrt{N}}$ fluctuations of the original evolutionary process around the rest point x^*. Such a distribution will exist whenever the rest point x^* is linearly stable under the mean dynamic—that is, whenever the eigenvalues of $DV^{\rho,A}(x^*)$ corresponding to eigenvectors tangent to the simplex have negative real part.[49] To derive this distribution, from (25), write $M^* = DV^{\rho,A}(x^*)$, and use Ito's lemma to show that the solution to (31) with $Z_0 = 0$ is the zero-mean Gaussian process

(26) $$Z_t = \int_0^t \exp\left(M^*(t-s)\right) \sqrt{\Sigma^*}\,dB_s,$$

whose time t covariance matrix is

(27) $$\mathrm{Cov}(Z_t) = \int_0^t \exp(M^*s)\,\Sigma^*\exp(M^*s)'\,ds$$

(see Karatzas and Shreve (1991, Sec. 5.6.A)). When the relevant eigenvalues of M^* have negative real part, the limit of (27) exists, and it is the covariance matrix of the (zero-mean Gaussian) limiting distribution of the solution to (25).

The following examples use the diffusion approximation to describe "equilibrium" behavior in a Hawk-Dove game under the "imitate the best realization" protocol (see Section 3.1) and in a good Rock-Paper-Scissors game under the "linear dissatisfaction" protocol (see Section 3.2).

Example A.6. Local diffusion approximation for "imitate the best realization" in Hawk-Dove.

Suppose that agents use imitative candidate selection (6) with $c = 2$, that matching is limited matching with n-of-trials $= 1$, as in (9), and that the decision method is *best*, as in (11). In addition, suppose that agents play a Hawk-Dove game with strategy set

---

[48]The result in Sandholm (2003) is stated for sequences of continuous-time processes, but also holds for sequences of discrete-time processes as defined here. Compare Durrett (1996, Sec. 8.7).

[49]At boundary rest points, it is enough to consider eigenvalues corresponding to eigenvectors that are tangent to the face of the smallest face of the simplex containing the rest point, since only strategies in the support of the rest point exhibit nonnegligible fluctuations in use; see Sandholm (2003).

$S = \{H, D\}$ and a payoff matrix $\mathcal{A} \in \mathbb{R}^{S \times S}$ with $\mathcal{A}_{HD} > \mathcal{A}_{DD} > \mathcal{A}_{DH} > \mathcal{A}_{HH}$. Then in a record containing both strategies, the payoff associated with H is higher if and only if it is $\mathcal{A}_{HD}$. Thus for a revising H player, (11) can be summarized as follows:

$$\text{in game } \mathcal{A}, \quad \sigma_{HD}\big((H, \pi_H), (D, \pi_D)\big) = 1_{\pi_H = \mathcal{A}_{HH}}.$$

Thus

$$\rho^{\mathcal{A}}_{HD}(x) = \sum_{(s,\pi)} p_i(s \mid x)\, q(\pi \mid s, x, A)\, \sigma_{ij}(s, \pi) = \sum_{\ell \in S} x_D \times x_H x_\ell \times 1 = x_D x_H,$$

and similarly,

$$\rho^{\mathcal{A}}_{DH}(x) = \sum_{(s,\pi)} p_i(s \mid x)\, q(\pi \mid s, x, A)\, \sigma_{ij}(s, \pi) = \sum_{k \in S} x_H \times x_k x_D \times 1 = x_H x_D.$$

It is convenient to change the state variable from $x = (x_H, x_D) \in X$ to $y \in [0, 1]$ by letting $y = x_H = 1 - x_D$. The mean dynamic (18) generated by ρ and $\mathcal{A}$ is then

(28) $\qquad \dot{y} = V^{\rho, \mathcal{A}}(y) = (1 - y)\, \rho^{\mathcal{A}}_{DH}(y) - y\, \rho^{\mathcal{A}}_{HD}(y) = y(1 - y)(1 - 2y).$

The rest points of (28) are 0, 1, and $y^* = \frac{1}{2}$.

To describe the diffusion approximation at interior rest point $y^* = \frac{1}{2}$, let $\{Y^N_t\}_{t \in \{0, 1/N, \ldots\}}$ be the Markov chain (3) expressed in terms of the new state variable, let $\{\bar{Y}^N_t\}_{t \geq 0}$ be its piecewise constant interpolation, and let

(29) $\qquad Z^N_t = \sqrt{N}(\bar{Y}^N_t - y^*)$

be the local behavior process near y^*. To use approximation (25), we calculate the derivative

$$DV^{\rho, \mathcal{A}}(y^*) = 6y^2 - 6y - 1\Big|_{y=y^*} = -\frac{1}{2}$$

and the local variance (24):

(30) $\qquad \Sigma^* = (1 - y^*)\, \rho^{\mathcal{A}}_{DH}(y^*) + y^*\, \rho^{\mathcal{A}}_{HD}(y^*) = y^*(1 - y^*) = \frac{1}{4}.$

Then (29) is approximated in distribution by the solution to

Figure 16: The mean dynamic $V^{\rho,\mathcal{A}}(y) = y(1-y)(1-2y)$ (bottom) and the local variance function $\Sigma^{\rho,\mathcal{A}}(x) = y(1-y)$ (top) in Example A.6.

(31) $\qquad dZ_t = DV_i^{\rho,A}(y^*)Z_t \, dt + \sqrt{\Sigma^*} \, dB_t = -\frac{1}{2}Z_t \, dt + \frac{1}{2} \, dB_t.$

with $Z_0 = 0$. Since $M^* \equiv DV^{\rho,A}(y^*) = -\frac{1}{2}$, the limiting distribution of $\{Z_t\}$ exists; it is normal with mean 0, and by (27) its variance is

$$\int_0^\infty \exp(2M^* s) \, \Sigma^* \, ds = -\frac{\Sigma^*}{2M^*} = \frac{1}{4}.$$

Using (29) to return to the scale of the original process, the diffusion approximation (31) estimates the variance of the fluctuations of $\{Y_t^N\}$ about y^* as $(\frac{1}{\sqrt{N}})^2 \cdot \frac{1}{4} = \frac{1}{4N}$, and the standard deviation of these fluctuations as $\frac{1}{2\sqrt{N}}$.

Dispersion estimates based on the diffusion approximation need not be accurate at moderate population sizes if the linear approximation of expected motion and the constant approximation of the dispersion of increments are inaccurate for states outside a small neighborhood of y^*. But in the present example, the latter approximations are accurate even at small population sizes. Figure 16 graphs the mean dynamic (28) and the local variance function $\Sigma^{\rho,\mathcal{A}}(x) = y(1-y)$ (cf. equation (30)); clearly, the linear approximation $V^{\rho,\mathcal{A}}(y) \approx DV^{\rho,\mathcal{A}}(y^*)(y-y^*)$ and the constant approximation $\Sigma^{\rho,\mathcal{A}}(x) \approx \Sigma^*$ lead only to small errors for a wide interval of states around $y^* = \frac{1}{2}$. This explains why the diffusion approximation estimate of the empirical standard deviation is accurate even at small population sizes (Table 3). ◆

Example A.7. Local diffusion approximation for "linear dissatisfaction" in good Rock-Paper-Scissors.

Suppose candidate selection is imitative as in (6) with $c = 2$, matching is complete (8), and the decision method is *linear-dissatisfaction* (14). Following the approach from Example A.1, we compute the off-diagonal terms of the revision protocol:

$$\rho_{ij}^A(x) = \sum_{(s,\pi)} p_i(s \mid x) \, q(\pi \mid s, x, A) \, \sigma_{ij}(s, \pi)$$

$$= x_j \times 1 \times \sigma_{ij}((i, (Ax)_i), (j, (Ax)_j))$$

$$= x_j \frac{A^{\max} - (Ax)_i}{\bar{\Delta}A},$$

where the second equality uses the fact that the only strategy record that allows a revising i player to switch to j is $(s_1, s_2) = (i, j)$. Then using (18), we compute the mean dynamic as

$$\dot{x}_i = \sum_{j \neq i} \left(x_j \rho_{ji}^A(x) - x_i \rho_{ij}^A(x) \right)$$

$$= \sum_{j \neq i} \left(x_j \left(x_i \frac{A^{\max} - (Ax)_j}{\bar{\Delta}A} \right) - x_i \left(x_j \frac{A^{\max} - (Ax)_i}{\bar{\Delta}A} \right) \right)$$

$$= \frac{1}{\bar{\Delta}A} x_i \sum_{j \in S} x_j ((Ax)_i - (Ax)_j)$$

$$(32) \qquad = \frac{1}{\bar{\Delta}A} x_i ((Ax)_i - x' Ax).$$

As in Example A.1, we obtain the replicator dynamic sped up by a constant factor of $\frac{1}{\bar{\Delta}A}$.

To obtain the diffusion approximation at interior rest point $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, we do not need to completely compute the Jacobian $DV^{\rho,A}(v)$. Instead, letting $\mathbf{1}$ denote the column vector of ones, and letting $\Phi = I - \frac{1}{n}\mathbf{11}'$ denote the orthogonal projection onto the set $\mathbb{R}_0^n = \{z \in \mathbb{R}^n : \mathbf{1}'z = 0\}$, the tangent space of the simplex, it is enough to compute $DV^{\rho,A}(x^*)\Phi$, which describes the action of $V^{\rho,A}(x^*)$ on vectors in $\mathbb{R}_0^n$.[50] A direct computation (or see Sandholm (2010b, equation (8.29))) shows that

$$(33) \qquad DV^{\rho,A}(x^*)\Phi = \frac{1}{\bar{\Delta}A} \left(\text{diag}(x^*) - x^*(x^*)' \right) A\Phi = \frac{1}{3\bar{\Delta}A} \Phi A\Phi,$$

where the second equality uses the fact that $x^* = \frac{1}{3}\mathbf{1}$.

In the case of good Rock-Paper-Scissors (equation (1) with $w = 2$ and $\ell = 1$), the projected Jacobian (33) and the local covariance matrix (24) become

---

[50]By construction, the Jacobian $DV^{\rho,A}(x^*)$ maps $\mathbb{R}_0^n$ to itself, and the local covariance matrix Σ^* from (24) and its square root $\sqrt{\Sigma^*}$ have rows and columns in $\mathbb{R}_0^n$. Writing the stochastic differential equation (25) in its integral form and using these facts shows that if $Z_0 \in \mathbb{R}_0^n$, then $Z_t \in \mathbb{R}_0^n$ for all $t \geq 0$. We can therefore replace Z_t with ΦZ_t in (25), after which we can associate Φ with the Jacobian appearing before it, obtaining the matrix product suggested in the text. The same conclusion can be reached by applying a similar argument to the explicit solution (26) of (25).

$$M^* \equiv DV^{\rho,A}(x^*)\Phi = \frac{1}{27}\begin{pmatrix} -1 & -4 & 5 \\ 5 & -1 & -4 \\ -4 & 5 & -1 \end{pmatrix} \quad \text{and} \quad \Sigma^* = \frac{10}{27}\Phi$$

The two eigenvalues of M^* corresponding to eigenvectors in $\mathbb{R}_0^n$ are $\frac{5}{27}(-11 \pm \sqrt{3}i)$. Thus the local diffusion process $\{Z_t\}$ defined by $dZ_t = M^* Z_t \, dt + \sqrt{\Sigma^*}\, dB_t$ has a limiting distribution, which by (27) is zero-mean Gaussian with covariance matrix

$$\int_0^\infty \exp(M^* s)\, \Sigma^* \exp(M^* s)'\, ds = \frac{10}{3}\Phi.$$

Returning to the scale of the original process, the approximate covariance matrix describing the fluctuations of $\{X_t^N\}$ around x^* is $\frac{10}{3N}\Phi$, so the standard deviation of the fluctuations in each component is approximately $\sqrt{\frac{20}{9N}} \approx 1.491\frac{1}{\sqrt{N}}$.

The prediction about the appearance of Figure 8(iii) is based on the following claim: if we map the simplex $X \subset \mathbb{R}^3$ isometrically into the plane, then the quasi-stationary distribution on X is mapped to a Gaussian distribution on the plane with independent components, and hence circular level sets. To establish this claim, let X_∞^N be a random vector with mean vector x^* and covariance matrix $\frac{10}{3N}\Phi$. Let $T\colon X \to \mathbb{R}^2$ be any transformation that maps X isometrically into $\mathbb{R}^2$ and (for definiteness) sends $x^* \in \mathbb{R}^3$ to $\mathbf{0} \in \mathbb{R}^2$. Then we can write $T(x) = R(x - x^*)$, where $R \in \mathbb{R}^{2\times 3}$ is a matrix whose two rows form an orthonormal basis for $\mathbb{R}_0^3$. Now let $Y_\infty^N = T(X_\infty^N)$. By standard properties of Gaussian distributions, Y_∞^N is zero-mean Gaussian with covariance matrix

$$\mathrm{Cov}(Y_\infty^N) = R\,\mathrm{Cov}(X_\infty^N)R' = R\left(\frac{10}{3N}\Phi\right)R' = \frac{10}{3N}RR' = \frac{10}{3N}I,$$

and so has independent components, each with variance $\frac{10}{3N}$. $\blacklozenge$

A.4 Random-walk tie-breaking

Here we describe the *random-walk* tie-breaking option for decision method *best* (11). This option makes use of an auxiliary Markov chain $\{Y_t^N\}_{t\in\{0,1/N,\dots\}}$ that runs on the state space $\mathcal{X}_+^N = \{y \in \mathcal{X}^N\colon y_i > 0 \text{ for all } i \in S\}$ according to the following transition law:[51]

---

[51]To interpret this process, imagine that there is an auxiliary population of $N+n$ agents choosing strategies in S. For each strategy in S, the population has one agent who always plays that strategy. The process $\{Y_t^N\}$ describes the behavior of the remaining N agents. In each period, one of these agents is chosen at random; this agent then chooses one of the remaining $N+n-1$ agents at random and switches to that agent's strategy.

$$\mathbb{P}\left(Y_{k+1/N}^N = y + \tfrac{1}{N}(e_j - e_i)\,\middle|\, Y_k^N = y\right) = y_i\,\frac{Ny_j + 1}{N + n - 1}.$$

Sandholm (2010b, Example 11.4.3) shows that the stationary distribution of this irreducible Markov chain is the uniform distribution on $\mathcal{X}_+^N$.

Returning to the original model, let (s, π) be the record of a revising agent, and let $m_j(s, \pi) = \#\{h \colon s_h = j \text{ and } \pi_h = \max_{h'} \pi_{h'}\}$ be the number of pairs in the record comprised of strategy j and the highest payoff in the record. Let y denote the current state of the auxiliary process. Then decision method *best* with *random-walk* tie-breaking is defined by

(34) $$\sigma_{ij}(s, \pi, y) = \frac{m_j(s, \pi)\, y_j}{\sum_{k \in S} m_k(s, \pi)\, y_k}.$$

Thus strategy j is chosen with probability proportional both to the number of times in the record that it is associated with the highest payoff and to the jth component of the auxiliary Markov chain.[52]

---

[52]Note that (34) agrees with (11) when the antecedent of (11) is satisfied (i.e., when tie-breaking is unnecessary).

References

Adami, C., Schossau, J., and Hintze, A. (2016a). Evolutionary game theory using agent-based methods. *Physics of Life Reviews*, 19:1–26.

Adami, C., Schossau, J., and Hintze, A. (2016b). The reasonable effectiveness of agent-based simulations in evolutionary game theory: Reply to comments on "Evolutionary game theory using agent-based methods". *Physics of Life Reviews*, 19:38–42.

Balkenborg, D. and Schlag, K. H. (2001). Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595.

Barreto, C. (2014). Population dynamics toolbox (PDToolbox). Software. https://github.com/carlobar/PDToolbox_matlab.

Bellomo, N. and Elaiw, A. (2016). Space dynamics and stochastic features of evolutionary game dynamics. comment on "Evolutionary game theory using agent-based methods". *Physics of Life Reviews*, 19:27–28.

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.

Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35. Erratum, 21 (1997), 325.

Björnerstedt, J. and Weibull, J. W. (1996). Nash equilibrium and evolution by imitation. In Arrow, K. J. et al., editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin's Press, New York.

Blume, L. E. (1997). Population games. In Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, Reading, MA.

Cárdenas, J. C., Mantilla, C., and Sethi, R. (2015). Stable sampling equilibrium in common pool resource games. *Games*, 6:299–317.

Darroch, J. N. and Seneta, E. (1965). On quasi-stationary distributions in absorbing discrete-time finite Markov chains. *Journal of Applied Probability*, 2(1):88–100.

Durrett, R. (1996). *Stochastic Calculus: A Practical Introduction*. CRC Press, Boca Raton, FL.

Dzonsons, K. and Weibull, J. (2013). Birds: Game theory simulator. Software. http://www.kcons.eu/birds.

Eigen, M. and Schuster, P. (1979). *The Hypercycle: A Principle of Natural Self-Organization*. Springer, Berlin.

Ellison, G. (2000). Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.

Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232. Corrigendum, 51 (1997), 77-78.

Franchetti, F. and Sandholm, W. H. (2013). An introduction to *Dynamo*: Diagrams for evolutionary game dynamics. *Biological Theory*, 8(2):167–178.

Friedman, D. and Sinervo, B. (2015). Evolutionary game theory: Simulations. Software. https://doi.org/10.7291/D1MW2X.

Friedman, D. and Sinervo, B. (2016). *Evolutionary Games in Natural, Social, and Virtual Worlds*. Oxford University Press.

Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262.

Fudenberg, D. and Imhof, L. A. (2008). Monotone imitation dynamics in large populations. *Journal of Economic Theory*, 140:229–245.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge.

Fudenberg, D., Nowak, M. A., Taylor, C., and Imhof, L. A. (2006). Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theoretical Population Biology*, 70:352–363.

García, J. and van Veelen, M. (2012). Evolution and games. Interactive tutorials. http://evolutionandgames.nl.

García, J. and van Veelen, M. (2016). In and out of equilibrium I: Evolution of strategies in repeated games with discounting. *Journal of Economic Theory*, 161:161 – 189.

García, J. and van Veelen, M. (2018). No strategy can win in the repeated prisoner's dilemma: Linking game theory and computer simulations. *Frontiers in Robotics and AI*, 5:102.

Gaunersdorfer, A. and Hofbauer, J. (1995). Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11:279–303.

Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.

Hauert, C. (2018). Evoludo. Interactive tutorials in evolutionary games. https://wiki.evoludo.org.

Helbing, D. (1992). A mathematical model for behavioral changes by pair interactions. In Haag, G., Mueller, U., and Troitzsch, K. G., editors, *Economic Evolution and Demographic Change: Formal Models in Social Sciences*, pages 330–348. Springer, Berlin.

Hilbe, C. and Traulsen, A. (2016). Only the combination of mathematics and agent-based simulations can leverage the full potential of evolutionary modeling. comment on "Evolutionary game theory using agent-based methods". *Physics of Life Reviews*, 19:29–31.

Hofbauer, J. (1995a). Imitation dynamics for games. Unpublished manuscript, University of Vienna.

Hofbauer, J. (1995b). Stability for the best response dynamics. Unpublished manuscript, University of Vienna.

Hofbauer, J., Mallet-Paret, J., and Smith, H. L. (1991). Stable periodic solutions for the hypercycle system. *Journal of Dynamics and Differential Equations*, 3:423–436.

Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.

Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144:1665–1693.

Hofbauer, J., Schuster, P., and Sigmund, K. (1979). A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81:609–612.

Hofbauer, J., Schuster, P., and Sigmund, K. (1981). Competition and cooperation in catalytic self-replication. *Journal of Mathematical Biology*, 11:155–168.

Imhof, L. A., Fudenberg, D., and Nowak, M. A. (2005). Evolutionary cycles of cooperation and defection. *Procedings of the National Academy of Sciences*, 31:10797–10800.

Imhof, L. A. and Nowak, M. A. (2006). Evolutionary game dynamics in a wright-fisher process. *Journal of Mathematical Biology*, 52(5):667–681.

Izquierdo, L. R., Izquierdo, S. S., Galán, J. M., and Santos, J. I. (2013). Combining mathematical and simulation approaches to understand the dynamics of computer models. In Edmonds, B. and Meyer, R., editors, *Simulating Social Complexity: A Handbook*, chapter 11, pages 235–271. Springer, Berlin.

Izquierdo, L. R., Izquierdo, S. S., and Sandholm, W. H. (2018). Evodyn-3s: A mathematica computable document to analyze evolutionary dynamics in 3-strategy games. *SoftwareX*, 7:226 – 233.

Izquierdo, L. R., Izquierdo, S. S., and Sandholm, W. H. (2019). *Agent-Based Evolutionary Game Dynamics*. University of Wisconsin Pressbooks.

Izquierdo, S. S. and Izquierdo, L. R. (2013). Stochastic approximation to understand simple simulation models. *Journal of Statistical Physics*, 151(1):254–276.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.

Kandori, M. and Rob, R. (1995). Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414.

Karatzas, I. and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*. Springer, New York, second edition.

Kosfeld, M., Droste, E., and Voorneveld, M. (2002). A myopic adjustment process leading to best reply matching. *Journal of Economic Theory*, 40:270–298.

Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67.

Loginov, G. (2019). Ordinal imitative dynamics. Unpublished manuscript. https://arxiv.org/abs/1907.04272.

Mantilla, C., Sethi, R., and Cárdenas, J. C. (2019). Efficiency and stability of sampling equilibrium in public good games. *Journal of Public Economic Theory*, forthcoming.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Moran, P. A. P. (1958). Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(1):60–71.

Moran, P. A. P. (1962). *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford.

Newton, J. (2018). Evolutionary game theory: a renaissance. *Games*, 9(2).

Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap/Harvard, Cambridge.

Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, 88:834–847.

Oyama, D., Sandholm, W. H., and Tercieux, O. (2015). Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics*, 10:243–281.

Pacheco, J. M., Traulsen, A., and Nowak, M. A. (2006). Coevolution of strategy and structure in complex networks with dynamical linking. *Phys. Rev. Lett.*, 97:258103.

Perc, M. (2018). Stability of subsystem solutions in agent-based models. *European Journal of Physics*, 39(1):014001.

Perc, M., Gómez-Gardeñes, J., Szolnoki, A., Floría, L. M., and Moreno, Y. (2013). Evolutionary dynamics of group interactions on structured populations: a review. *Journal of The Royal Society Interface*, 10(80):20120997.

Perc, M. and Szolnoki, A. (2010). Coevolutionary games—a mini review. *Biosystems*, 99(2):109–125.

Robson, A. and Vega-Redondo, F. (1996). Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory*, 70:65–92.

Roca, C. P., Cuesta, J. A., and Sánchez, A. (2006). Time scales in evolutionary dynamics. *Phys. Rev. Lett.*, 97:158701.

Roca, C. P., Cuesta, J. A., and Sánchez, A. (2009a). Effect of spatial structure on the evolution of cooperation. *Phys. Rev. E*, 80:046106.

Roca, C. P., Cuesta, J. A., and Sánchez, A. (2009b). Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Physics of Life Reviews*, 6(4):208–249.

Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25:92–100.

Roth, G. and Sandholm, W. H. (2013). Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51:525–555.

Sample, C. and Allen, B. (2017). The limits of weak selection and large population size in evolutionary game theory. *Journal of Mathematical Biology*, 75(5):1285–1317.

Sánchez, A. and Cuesta, J. A. (2005). Altruism may arise from individual selection. *Journal of Theoretical Biology*, 235(2):233–240.

Sandholm, W. H. (2001a). Almost global convergence to p-dominant equilibrium. *International Journal of Game Theory*, 30:107–116.

Sandholm, W. H. (2001b). Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108.

Sandholm, W. H. (2003). Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378.

Sandholm, W. H. (2010a). Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. *Games*, 1:3–17.

Sandholm, W. H. (2010b). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.

Sandholm, W. H. (2015). Population games and deterministic evolutionary dynamics. In Young, H. P. and Zamir, S., editors, *Handbook of Game Theory with Economic Applications*, volume 4, chapter 13, pages 703–778. Elsevier, Amsterdam.

Sandholm, W. H., Dokumacı, E., and Franchetti, F. (2012). Dynamo: Diagrams for evolutionary game dynamics. Software. http://www.ssc.wisc.edu/$\sim$whs/dynamo.

Sandholm, W. H., Izquierdo, S. S., and Izquierdo, L. R. (2018). Stability for best experienced payoff dynamics. Unpublished manuscript, University of Wisconsin, Universidad de Valladolid, and Universidad de Burgos.

Sandholm, W. H., Izquierdo, S. S., and Izquierdo, L. R. (2019). Best experienced payoff dynamics and cooperation in the Centipede game. *Theoretical Economics*, forthcoming.

Sandholm, W. H. and Staudigl, M. (2016). Large deviations and stochastic stability in the small noise double limit. *Theoretical Economics*, 11:279–355.

Sandholm, W. H. and Staudigl, M. (2018). Sample path large deviations for stochastic evolutionary game dynamics. *Mathematics of Operations Research*, 43(4):1348–1377.

Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156.

Schlag, K. H. (1999). Which one should I imitate? *Journal of Mathematical Economics*, 31:493–522.

Schuster, P. (2016). Models of evolution and evolutionary game theory. a comment on "evolutionary game theory using agent based models". *Physics of Life Reviews*, 19:32–35.

Schuster, P. and Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538.

Schuster, P., Sigmund, K., and Wolff, R. (1978). Dynamical systems under constant organization I: Topological analysis of a family of nonlinear differential equations—a model for catalytic hypercycles. *Bulleting of Mathematical Biology*, 40:743–769.

Sethi, R. (2000). Stability of equilibria in games with procedurally rational players. *Games and Economic Behavior*, 32:85–104.

Smith, M. J. (1984). The stability of a dynamic model of traffic assignment—an application of a method of Lyapunov. *Transportation Science*, 18:245–252.

Stroock, D. W. and Varadhan, S. R. S. (1979). *Multidimensional Diffusion Processes*. Springer, New York.

Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, 446:97–216.

Szolnoki, A., Mobilia, M., Jiang, L.-L., Szczesny, B., Rucklidge, A. M., and Perc, M. (2014). Cyclic dominance in evolutionary games: a review. *Journal of The Royal Society Interface*, 11(100):20140735.

Tarnita, C. E. (2016). Mathematical approaches or agent-based methods? comment on "evolutionary game theory using agent-based methods". *Physics of Life Reviews*, 19:36–37.

Taylor, C., Fudenberg, D., Sasaki, A., and Nowak, M. A. (2004). Evolutionary game dynamics in finite populations. *Bulletin of Mathematical Biology*, 66:1621–1644.

Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.

Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115.

Traulsen, A. and Hauert, C. (2009). Stochastic evolutionary game dynamics. In Schuster, H. G., editor, *Reviews of Nonlinear Dynamics and Complexity*, volume 2, pages 25–61. Wiley, New York.

Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2007). Stochastic payoff evaluation increases the temperature of selection. *Journal of Theoretical Biology*, 244(2):349–356.

van Veelen, M. (2012). Robustness against indirect invasions. *Games and Economic Behavior*, 74:382–393.

Vega-Redondo, F. (1997). The evolution of Walrasian behavior. *Econometrica*, 65:375–384.

Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge.

Wilensky, U. (1999). Netlogo. Software. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Wilensky, U. and Shargel, B. (2002). BehaviorSpace. Software. http://ccl.northwestern.edu/netlogo/behaviorspace.html. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Woelfing, B. and Traulsen, A. (2009). Stochastic sampling of interaction partners versus deterministic payoff assignment. *Journal of Theoretical Biology*, 257(4):689 – 695.

Wu, B., Altrock, P. M., Wang, L., and Traulsen, A. (2010). Universality of weak selection. *Phys. Rev. E*, 82:046106.

Wu, B., García, J., Hauert, C., and Traulsen, A. (2013). Extrapolating weak selection in evolutionary games. *PLOS Computational Biology*, 9(12):1–7.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.

Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.

Zeeman, E. C. (1980). Population dynamics from game theory. In Nitecki, Z. and Robinson, C., editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in Lecture Notes in Mathematics, pages 472–497, Berlin. Springer.