

## 9.3 Evolutionary Game Theory

J. McKenzie Alexander

### Summary

Evolutionary game theory analyzes the outcomes of interactions between boundedly rational individuals in strategic settings. Originally developed in the field of population biology, economists and other social scientists became interested in evolutionary game theory because it provided a natural way of incorporating dynamics into game theory. Given the wide variety of different evolutionary dynamics and models, a number of interesting and surprising results have been proven. Under some dynamics, both weakly and strictly dominated strategies can persist in the limit. And, although it is typically the case that population states corresponding to Nash equilibria are fixed points of the evolutionary dynamics, some dynamics allow for fixed points that do *not* correspond to Nash equilibria. The connection between evolutionary dynamics and the traditional solution concept of game theory—the Nash equilibrium—is thus found to be subtle and complex.

### 1. The Origins of Evolutionary Game Theory

Evolutionary game theory emerged in the 1970s from the realization that frequency-dependent fitness introduced a strategic aspect to evolution, even in contexts where none of the agents were traditionally “rational.” The first work of modern evolutionary game theory is generally understood to be Maynard Smith and Price (1973), in which the concept of the *evolutionarily stable strategy* (ESS) was introduced.<sup>1</sup> Yet one can see the beginnings of game-theoretic approaches to the study of evolution in the work of R. A. Fisher, who, in *The Genetical Theory of Natural Selection* (1930), attempted to explain the approximate equality of the mammalian sex ratio using arguments that were essentially game-theoretic in nature.

After its introduction in population biology, evolutionary game theory became of interest to economists and social scientists more generally. Initially, it was hoped that, because the concept of an evolutionarily stable strategy

provided a proper strengthening of the Nash equilibrium solution concept, evolutionary game theory could offer insight into the equilibrium selection problem: when multiple Nash equilibria exist, which one rational players would settle upon. In addition, the fact that evolutionary game theory provided dynamical models of quasi-rational deliberation was seen to address a lacuna in the traditional theory of games, which largely provided a static analysis of strategic problems. In *Theory of Games and Economic Behavior*, von Neumann and Morgenstern wrote,

We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and therefore preferable. But there is ample evidence from other branches of science that it is futile to try to build one as long as the static side is not thoroughly understood. (§ 4.8.2)

The “most [emphatic]” reminder of the von Neumann and Morgenstern theory being static was needed because, although the extensive-form representation of games does make explicit the temporal relation between the choice nodes of players, individual strategies were represented atemporally, where a strategy specified the choice of a player, in advance, at all possible choice nodes.

Over time, it was recognized that evolutionary game theory was less successful in delivering on these aims than initially hoped. The equilibrium selection problem reappeared in the form of competing notions of evolutionary stability, and the plurality of dynamical models yielded varying results regarding the long-term survival of strategies. For example, whereas some concepts of evolutionary stability rule out even weakly dominated strategies, in the dynamical setting, it turns out that both weakly *and* strictly dominated strategies can persist for some families of dynamics. In addition, some families of dynamics—particularly imitative dynamics—have the property that not every fixed point of the dynamics corresponds to a Nash equilibrium of the underlying game.

Since its inception, evolutionary game theory has morphed into a more general study of the population

dynamics of boundedly rational agents (see chapter 8.5 by Hertwig & Kozyreva, this handbook), with the once-central concept of an evolutionarily stable strategy, and its subsequent refinements, no longer in central place. The family of models developed provides useful mathematical tools for modeling and understanding the behavior of quasi-rational individuals interacting in strategic contexts, but care needs to be exercised to ensure that the model and the target system are aligned before drawing conclusions.

## 2. Game-Theoretic Fundamentals

Consider a game played by a group of  $N$  individuals, where each player  $i$  has a set  $S_i$  of *pure strategies*. Each strategy  $s \in S_i$  can be thought of as an action available to player  $i$ ; the complete set of actions taken by all players determines the payoffs each player receives. In traditional game theory, it often proves useful to allow a player to choose a strategy at random, so as to keep other individuals uncertain as to what she will do; such strategies are called *mixed strategies* and correspond to probability distributions over the set of pure strategies. The set of mixed strategies for player  $i$  is denoted  $\Delta_i$ , and a particular mixed strategy for player  $i$  will be denoted  $\sigma_i \in \Delta_i$ . Mixed strategies, while indispensable for many important theoretical results in traditional game theory (see chapter 9.1 by Albert & Kliemt, this handbook), such as Nash's theorem that every finite  $n$ -player game has at least one equilibrium, play a less central role in evolutionary game theory.

If  $\vec{\sigma} = (\sigma_1, \dots, \sigma_N)$  denotes the particular strategies used by each player, also known as a *strategy profile*, the payoff to player  $i$  is given by  $\pi_i(\vec{\sigma})$ , where  $\pi_i$  is the *payoff function* for that player. In two-player games, we will use the notation " $\pi_i(\sigma | \mu)$ " to denote the fitness of player  $i$  when  $\sigma$  is played against  $\mu$ . A two-player game is *symmetric* when, for all strategies  $\sigma, \mu$ , the payoff of  $\sigma$  played against  $\mu$  does not depend on the identity of the player; for such games, the payoff function will simply be written " $\pi$ ."

The fundamental solution concept for traditional game theory is the *Nash equilibrium*, defined as follows:

*Definition.* A strategy profile  $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$  is a Nash equilibrium if and only if for each player  $i$  and all strategies  $\sigma_i \in \Delta_i$ ,

$$\pi_i(\sigma_i^*, \sigma_{-i}^*) \geq \pi_i(\sigma_i, \sigma_{-i}^*),$$

where " $(\sigma_i, \sigma_{-i}^*)$ " denotes the strategy profile identical to  $\sigma^*$  except that the strategy for player  $i$  is replaced by  $\sigma_i$ .

A Nash equilibrium is a set of mutual best responses for each player, where no individual can improve their payoff by adopting an alternative strategy. This property makes

it a natural minimal solution concept for traditional game theory. All games with a finite number of players and finitely many strategies have at least one Nash equilibrium.

## 3. Evolutionarily Stable Strategies

To see that the concept of a Nash equilibrium is too *weak* for capturing the idea of evolutionary stability, consider the game shown in figure 9.3.1a. That game has two Nash equilibria in pure strategies:  $(S_1, S_1)$  and  $(S_2, S_2)$ . But imagine this game played in a population where individuals are paired at random and fitness corresponds to expected payoff. Informally, it is clear that a population consisting entirely of the  $S_2$  type would not be evolutionarily stable. In such a population, a novel  $S_1$  mutant receives the same payoff as the incumbent population and thus does not incur a fitness disadvantage. This means an  $S_1$  mutant can invade the population. If two  $S_1$  mutants appear, the payoff of an  $S_1$ - $S_1$  interaction exceeds that of an  $S_2$ - $S_2$  interaction, which means the  $S_1$  type has an expected fitness greater than the average of the population and thus would be expected to increase in number, potentially even driving the  $S_2$  type to extinction.<sup>2</sup> The problem stems from the fact that a Nash equilibrium only requires that deviant play *does no better* than the equilibrium strategy, while allowing for the possibility that deviant play may *do as well as* the equilibrium strategy.

A *strict Nash equilibrium* is one where any deviation from the equilibrium profile leaves players worse off. To see that the concept of a strict Nash equilibrium is too *strong* to capture the idea of evolutionary stability, consider the game in figure 9.3.1b, the Hawk-Dove of Maynard Smith and Price (1973).<sup>3</sup> In this game, two individuals compete over a resource with value  $V$ . The strategies individuals may adopt are "Hawk," which escalates to the point of fighting, if necessary, and "Dove," which makes a show of escalating initially but then backs down just before the point of fighting. Given this, if a Hawk meets a Dove, the Hawk gets the entire resource. If two Doves meet, they share the resource. If two Hawks meet, they escalate to the point of fighting until one is injured, with cost  $-C$ , and retreats.<sup>4</sup> If  $V < C$ , this game has no pure-strategy Nash equilibrium and one mixed-strategy Nash equilibrium where both individuals play Hawk with probability  $V/C$ . Denote this mixed strategy by  $\sigma^*$ . From the fundamental theorem of mixed-strategy Nash equilibria (see Gintis, 2009), it can be shown that

$$\pi(\text{Hawk} | \sigma^*) = \pi(\text{Dove} | \sigma^*) = \pi(\sigma^* | \sigma^*).$$

Therefore, for any other mixed strategy  $\mu \neq \sigma^*$ , it is the case that  $\pi(\mu | \sigma^*) = \pi(\sigma^* | \sigma^*)$ , and so  $\sigma^*$  is not a strict

	$S_1$	$S_2$
$S_1$	(2, 2)	(1, 1)
$S_2$	(1, 1)	(1, 1)

(a) A Nash equilibrium allows for evolutionary drift, and hence may not be evolutionarily stable.

	Hawk	Dove
Hawk	$\left(\frac{V-C}{2}, \frac{V-C}{2}\right)$	(V, 0)
Dove	(0, V)	$\left(\frac{V}{2}, \frac{V}{2}\right)$

(b) A game whose Nash equilibrium is not strict when  $V < C$ , yet is evolutionarily stable.

Figure 9.3.1

Why the concept of a Nash equilibrium is too weak to capture the idea of evolutionary stability.

Nash equilibrium strategy. However, unlike the game in figure 9.3.1a, here it can be shown that  $\pi(\sigma^* | \mu) > \pi(\mu | \mu)$ . This means that the mutant strategy  $\mu$  is at a fitness disadvantage in that it does less well *when it plays against itself* than the incumbent strategy  $\sigma^*$  does when it plays against  $\mu$ . This means that the mutant strategy  $\mu$  would be unable to invade and, hence, the incumbent strategy  $\sigma$  is evolutionarily stable.

These considerations led Maynard Smith (1982) to propose the following:

*Definition.* A strategy  $\sigma$  is an *evolutionarily stable strategy* (ESS) if and only if for all strategies  $\mu \neq \sigma$ , it is the case that

$$\text{either } \pi(\sigma | \sigma) > \pi(\mu | \sigma),$$

$$\text{or } \pi(\sigma | \sigma) = \pi(\mu | \sigma) \text{ and } \pi(\sigma | \mu) > \pi(\mu | \mu).$$

The first clause states that every strict Nash equilibrium strategy is an ESS; the second clause states that, if an ESS is not a strict best reply to itself, then the ESS must be a strict best reply to all *other* strategies. Although it is true that all strict Nash equilibrium strategies are evolutionarily stable, the converse does not hold. The second-order best-reply condition means that an ESS provides a real strengthening of the Nash equilibrium solution concept.

An alternative definition of an evolutionarily stable strategy, first introduced by Taylor and Jonker (1978) and used by Weibull (1995) and others, is as follows. Let  $\sigma$  and  $\mu$  be strategies, and let  $\varepsilon\mu + (1 - \varepsilon)\sigma$  for  $\varepsilon > 0$  denote the strategy, which  $\varepsilon$  of the time plays  $\mu$  and  $1 - \varepsilon$  of the time plays  $\sigma$ . For small  $\varepsilon$ , this can be thought of as either a small “tremble” around the base strategy  $\sigma$  or, equivalently, as the strategy employed by a randomly chosen opponent in a population of individuals where almost all play  $\sigma$  but a small number of mutants adopt  $\mu$ .

*Definition.* A strategy  $\sigma$  is an *evolutionarily stable strategy* if and only if for every strategy  $\mu \neq \sigma$ , there exists a  $0 < \varepsilon_\mu < 1$  such that

$$\pi(\sigma | \varepsilon\mu + (1 - \varepsilon)\sigma) > \pi(\mu | \varepsilon\mu + (1 - \varepsilon)\sigma)$$

for all  $0 < \varepsilon < \varepsilon_\mu$ .

It is easily shown that this definition, which can be interpreted as a mathematically precise formulation of the definition given by Maynard Smith and Price (1973), is equivalent to the one proposed by Maynard Smith (1982).

The above definition of an ESS states that the expected payoff of an ESS in a mixed population consisting of the ESS and a small number of mutants is strictly greater than the expected payoff of the mutant in that same mixed population. However, Hofbauer, Schuster, and Sigmund (1979) give the following characterization of an ESS, which not only has the benefit of being conceptually clearer but also proves to be more useful when we move to considering evolutionary dynamics in the next section:

*Definition.* A strategy  $\sigma$  is *locally superior* if there exists an  $\varepsilon$ -neighborhood around  $\sigma$  such that, for all strategies  $\mu \neq \sigma$  in that neighborhood, it is the case that  $\pi(\sigma | \mu) > \pi(\mu | \mu)$ .

From this, one can then prove the following:

*Theorem 1.* A strategy  $\sigma$  is an ESS if and only if  $\sigma$  is locally superior.

There are a number of theorems one can prove about ESSs. For example:

*Theorem 2.* No weakly dominated strategy is an ESS.

This is of interest because, as shown in figure 9.3.2, some weakly dominated strategies constitute Pareto-optimal Nash equilibria. In such cases, it can be rational to play a weakly dominated strategy even though it is not evolutionarily stable.

The *support* of a mixed strategy  $\sigma$  is the set of pure strategies to which  $\sigma$  assigns nonzero probability. One can show that if  $\mu$  and  $\sigma$  are ESSs with identical supports, then  $\mu = \sigma$ . From this, it immediately follows that

	$S_1$	$S_2$
$S_1$	(1, 1)	(100, 0)
$S_2$	(0, 100)	(100, 100)

Figure 9.3.2

The fact that no weakly dominated strategy is an ESS means that some Pareto-optimal Nash equilibria are excluded, as shown here.

	Rock	Paper	Scissors
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissors	(-1, 1)	(1, -1)	(0, 0)

**Figure 9.3.3**  
The game of Rock–Paper–Scissors, which has no ESS.

*Theorem 3.* The number of ESSs is finite (and possibly zero).

The latter half of the above theorem is easily proven by considering the well-known game of Rock–Paper–Scissors, shown in figure 9.3.3. The only Nash equilibrium is in mixed strategies and assigns equal probability to all three pure strategies. The zero-sum nature of the game means that neither of the two conditions in the definition of an ESS are met, and so this game has no evolutionarily stable strategy.

#### 4. Dynamics for Evolutionary Games

The main shortcoming of the evolutionarily-stable-strategy concept is that it, like that of a Nash equilibrium, is *static*. It would be useful if we could say something about the behavior of a population in an out-of-equilibrium environment, especially in games where multiple ESSs exist or where no ESS exists. For this reason, we now turn to dynamical models of evolutionary games. As we shall see, there is an imperfect match between static and dynamical notions of evolutionary stability, and although there are some stability concepts that possess a useful degree of generality, considerable variation exists in the behavior of a population between continuous and discrete models, as well as between unstructured and structured population models.

To begin, let  $S = \{S_1, \dots, S_n\}$  denote the set of pure strategies and let  $p_i(t) \in [0, 1]$  denote the proportion of the population using strategy  $i$  at time  $t$ . It is assumed that no player employs a mixed strategy and that  $\sum_{i=1}^n p_i(t) = 1$ . Let  $\Delta_n$  denote the set of all possible frequency distributions of strategies over the population. If the only thing that matters, from an evolutionary point of view, is the frequency with which a strategy is used in the population, then  $\vec{p}(t) = \langle p_1(t), \dots, p_n(t) \rangle \in \Delta_n$  represents the total state of the population at time  $t$ . A game, in this context, is a function  $W: \Delta_n \rightarrow \mathbb{R}^n$ , which

maps states of the population to a vector of expected payoffs—which will typically be thought of as expected fitnesses—for each pure strategy. In what follows, we use  $W_i(\vec{p}(t))$  to denote the expected payoff to strategy  $i$  when the population is in state  $\vec{p}$  at time  $t$ .

Sandholm (2010a, 2010b) provides a useful framework for thinking about evolutionary dynamics, demonstrating how continuous dynamics at the population level can be derived from particular learning rules—also known as *revision protocols*—used by the individuals comprising the population. This in effect provides the “microfoundations for existing dynamics” (Sandholm, 2010b, p. 31), the evolutionary game theory analogue of providing microfoundations for macroeconomics.

Sandholm models a revision protocol as a map from the current expected payoffs of each strategy in the population, as well as its present composition, to a transition matrix expressing the rate at which strategies switch from one type to another. More formally, a revision protocol is a map, generally assumed to be Lipschitz-continuous, as follows:

$$\rho: \mathbb{R}^n \times \Delta_n \rightarrow \mathbb{R}_+^{n \times n}.$$

Given a particular vector of expected payoffs  $\pi \in \mathbb{R}^n$  and a population state  $\vec{p} \in \Delta_n$ , the  $ij$ th entry of  $\rho(\pi, \vec{p})$ , written  $\rho_{ij}(\pi, \vec{p})$ , is the rate at which individuals switch from strategy  $i$  to strategy  $j$ .

Once a revision protocol is specified, the evolutionary dynamics for the population are derived by substituting the revision protocol into the following schema of differential equations defined on  $\Delta_n$  (see equation 1 below). For the precise mathematical details, see Sandholm (2010b).

Much of the early literature in evolutionary game theory focused on the *replicator dynamics*, the first dynamical model of evolutionary games, proposed by Taylor and Jonker (1978). The core idea behind the replicator dynamics is that the proportion of the population that follows a given strategy increases if the fitness of that strategy is higher than the average fitness of the population, which is denoted by  $\phi(\vec{p})$ . The average fitness of the population is just the weighted average of the individual-strategy fitnesses, with the proportion of the population following the respective strategies used as weights:

$$\phi(\vec{p}) = \sum_{i \in S} p_i \cdot W_i(\vec{p}).$$

The *continuous replicator dynamics* is the system of differential equations given by

$$\frac{dp_i}{dt} = \left( \begin{array}{c} \text{Rate at which the population} \\ \text{switches to strategy } i \end{array} \right) - \left( \begin{array}{c} \text{Rate at which the population} \\ \text{switches from strategy } i \end{array} \right). \tag{1}$$



$$\frac{dp_i}{dt} = p_i \cdot [W_i(\vec{p}) - \phi(\vec{p})], \text{ for } i \in S.$$

Note that if  $p_i(0) > 0$ , then  $p_i(t) > 0$  for all  $t$ . This is of course compatible with  $p_i(t) \rightarrow 0$  as  $t \rightarrow \infty$ ; the point is that, under the replicator dynamics, strategies can be driven to extinction in the limit, but they cannot actually be driven out in finite time. A related point is that the replicator dynamics cannot *introduce* strategies into the population if they are originally absent: if  $p_i(0) = 0$ , then  $p_i(t) = 0$  for all  $t$ .

In the simplest case, if population members interact at random and the probability of interacting with strategy  $j$  equals  $p_j(t)$ , then the fitness of strategy  $i$  is  $W_i(\vec{p}(t)) = \sum_j p_j(t) \cdot \pi(S_i | S_j)$ . This is the continuous *linear* replicator dynamics (henceforth, simply “the replicator dynamics”).

Schlag (1998) showed that the replicator dynamics can be derived from the following revision protocol: suppose that each player in the population selects someone at random and compares their individual payoff with that of the person selected. If the person selected received a higher payoff, then the player adopts the strategy of the person selected with a probability proportional to the difference in payoffs. That is, if  $\vec{\pi}$  denotes the vector of expected payoffs for the population in state  $\vec{p}$ , then  $\rho_{ij}(\vec{\pi}, \vec{p}) = p_j \cdot [\pi_j - \pi_i]_+$ , where  $[x]_+ = \max(x, 0)$ . The replicator dynamics, then, is the population-level description of the outcome when people *imitate with probability proportional to success*.

At this point, a number of interesting differences already emerge between the static ESS concept and the outcomes

of dynamical models: no weakly dominated strategy can be an ESS, but the replicator dynamics does not necessarily eliminate weakly dominated strategies. Indeed, in the case of the game of figure 9.3.4, almost all evolutionary trajectories converge to a state in which some proportion of the population follows a weakly dominated strategy.

Although the replicator dynamics need not eliminate weakly dominated strategies, Akin (1980) showed that, in most cases, the replicator dynamics will eliminate *strictly* dominated strategies:


*Theorem 4.* Let  $\vec{p} \in \Delta_n$  be a point in the interior of the simplex. If  $S_i$  is strictly dominated, then  $\lim_{t \rightarrow \infty} p_i(t) = 0$  under the replicator dynamics.

The requirement that all strategies must be represented in the initial state is easily seen to be necessary: consider the prisoner’s dilemma, and suppose that the population begins in the state where everyone follows the strategy Cooperate. Although Cooperate is strictly dominated by Defect, the fact that the replicator dynamics cannot introduce absent strategies means that the strictly dominated state of All-Cooperate will endure.

Although the replicator dynamics generally eliminates strictly dominated strategies, the same is not true for other types of evolutionary dynamics. Indeed, Hofbauer and Sandholm (2011) show that many families of evolutionary dynamics have the property of not eliminating strictly dominated strategies. This result shows how classical conceptions of rational behavior, at the individual level, can become decoupled from population-level

	$S_1$	$S_2$	$S_3$
$S_1$	(1, 1)	(1, 1)	(1, 0)
$S_2$	(1, 1)	(1, 1)	(0, 0)
$S_3$	(0, 1)	(0, 0)	(0, 0)

(a) A game where strategy  $S_2$  is weakly dominated by  $S_1$ .



(b) In almost all cases, the replicator dynamics converges to a polymorphic state featuring a mix of both  $S_1$  and  $S_2$ .

**Figure 9.3.4**

Weakly dominated strategies may not be eliminated by the replicator dynamics. Shading indicates the relative speed of the evolutionary trajectories.

outcomes. In what follows, two additional evolutionary dynamics are introduced, which—despite the plausible nature of the underlying revision protocol—do not, in general, eliminate strictly dominated strategies.

As we’ve seen, the replicator dynamics can be derived from a revision protocol that imitates strategies with probability proportional to success. One might wonder whether imitation is a reliable guide to future payoffs. Imitation certainly suffers from an inability to introduce new strategies into the population. Perhaps it would be better to adopt a strategy only taking into account its performance compared to aggregate properties of the overall population. One such revision protocol, which yields an interesting evolutionary dynamic, is the following: suppose that the rate at which players switch to strategy  $j$  does not depend on what strategy they currently employ, or what their expected payoff is in the current population state, but rather only on whether the expected payoff of strategy  $j$  exceeds the average payoff of the population. Sandholm (2010a) shows that when

$$\rho_{ij}(W(\vec{p}), \vec{p}) = [W_j(\vec{p}) - \phi(\vec{p})]_+, \text{ for all } j \in S,$$

the evolutionary dynamic obtained is the Brown–Nash–von Neumann (BNN) dynamic (Brown & von Neumann, 1950). This dynamic belongs to a more general class of dynamics known as *separable excess payoff dynamics*. The BNN dynamic takes the following form:

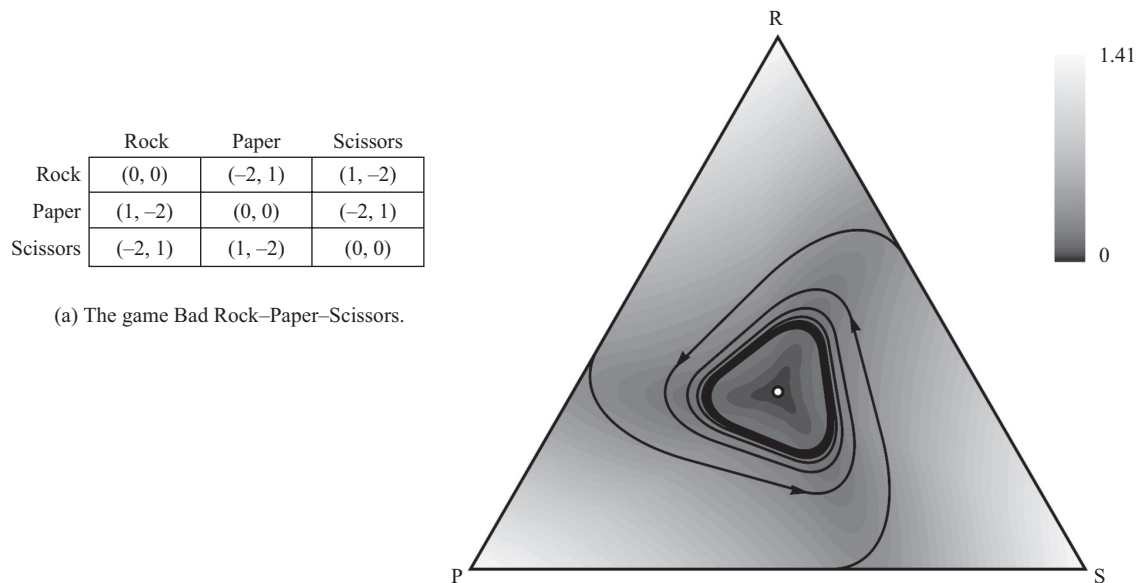
$$\frac{dp_i}{dt} = [W_i(\vec{p}) - \phi(\vec{p})]_+ - p_i \sum_{j \in S} [W_j(\vec{p}) - \phi(\vec{p})]_+, \text{ for all } i \in S.$$

One interesting difference between the BNN dynamic and the replicator dynamic is that the BNN dynamic can introduce *new* strategies into the population. Figure 9.3.5 shows this for the game Bad Rock–Paper–Scissors. The phase diagram shows the evolutionary trajectories for three different initial population states, each beginning at vertices of the simplex. The BNN dynamic moves the population along each of the respective faces in the direction of the dominating strategy. At some point, there is a fitness advantage to playing the third, unrepresented, strategy, and so the dynamics moves into the interior of the simplex and converges to an orbit.

The revision protocol that yields the BNN dynamics compares alternative strategies to the average payoff of the population. In a sense, this is a strange motivation because the average payoff of the population, in most cases, does not correspond to the payoff obtained by any particular strategy. As an alternative, consider the revision protocol that selects a new strategy by comparing the expected payout of alternative strategies to the expected payout of one’s current strategy, where only those strategies with higher expected payoffs have a nonzero chance of being adopted. That is,

$$\rho_{ij}(W(\vec{p}), \vec{p}) = [W_j(\vec{p}) - W_i(\vec{p})]_+.$$

This revision protocol generates an evolutionary dynamic first discussed by Smith (1984) and hence is known as the *Smith dynamic*. This dynamic is a member of a more general family of evolutionary dynamics



**Figure 9.3.5** The Brown–Nash–von Neumann dynamic can introduce new strategies into the population that are not initially present.

known as *impartial pairwise comparison dynamics*. (It is “impartial” in that the switch rate does not depend upon the player’s current strategy  $i$ .) The Smith dynamic takes the following form:

$$\frac{dp_i}{dt} = \sum_{j \in S} p_j \cdot [W_i(\vec{p}) - W_j(\vec{p})]_+ - p_i \sum_{j \in S} [W_j(\vec{p}) - W_i(\vec{p})]_+.$$

Both the Smith dynamic and the BNN dynamic allow strictly dominated strategies to persist in the limit. To see this, consider the game of figure 9.3.6a. Here, the game of Bad Rock–Paper–Scissors is augmented with a fourth strategy: a “feeble twin.” The feeble twin strategy is essentially the same as Scissors except that all of its payoffs are decremented by an additional amount  $\varepsilon > 0$ . Figure 9.3.6b shows how, under the Smith dynamic, populations beginning at the four vertices of the simplex (states such as All-Rock, All-Paper, etc.) follow trajectories such that all four strategies are played with positive probability, *including* the strictly dominated feeble twin strategy.

The phenomenon illustrated in figure 9.3.6 is a specific instance of the following general result:


*Theorem 5 (Hofbauer & Sandholm, 2011).* For all evolutionary dynamics that belong to either the family of impartial pairwise comparison dynamics or the family of separable excess payoff dynamics, there exists a game such that, from most initial conditions, there is a strictly dominated strategy played by a proportion of the population that is bounded away from 0 and exceeds  $\frac{1}{2}$  infinitely often as time approaches infinity.

One interesting point about the result of theorem 5 is that both the BNN dynamics and the Smith dynamics are generated by revision protocols that are not obviously implausible. Thus, irrational outcomes at the population level—the survival of strictly dominated strategies—can be generated by revision protocols that, at the individual level, may appear sensible.

All of the above three dynamics are *single-population* models, in that they assume the underlying game is

	Rock	Paper	Scissors	Twin
Rock	(0, 0)	(-2, 1)	(1, -2)	(1, -2 - $\varepsilon$ )
Paper	(1, -2)	(0, 0)	(-2, 1)	(-2, 1 - $\varepsilon$ )
Scissors	(-2, 1)	(1, -1)	(0, 0)	(0, - $\varepsilon$ )
Twin	(-2 - $\varepsilon$ , 1)	(1 - $\varepsilon$ , -1)	(- $\varepsilon$ , 0)	(- $\varepsilon$ , - $\varepsilon$ )

(a) In this version of Bad Rock–Paper–Scissors, the fourth strategy Twin is the same as Scissors except that all of its payoffs are decreased by  $\varepsilon$ .



**Figure 9.3.6**

The Smith dynamic allows strictly dominated strategies to be preserved.

symmetric. Since there is no way of distinguishing row and column players, it must be the case that  $\pi(S_i|S_j) = \pi(S_j|S_i)$ . Asymmetric games are typically analyzed using *multipopulation* models, one population for Row and another population for Column, which give rise to a coupled system of differential equations. Suppose, for example, that the base game is the two-player asymmetric game *Battle of the Sexes*, shown in figure 9.3.7a. The two-population replicator dynamics treats the Row and Column players as being drawn from separate subpopulations, where the expected fitness of a Row player depends on the distribution of strategies found in the Column subpopulation, and vice versa, as indicated in figure 9.3.7b. (The superscripts *R* and *C*

denote the Row and Column subpopulation, respectively.) In this case, the replicator dynamics takes the following form (again, the explicit time dependencies are suppressed for clarity):

$$\frac{dp_i^R}{dt} = p_i^R \cdot [W_i^R(\vec{p}^C) - \phi(\vec{p}^R)], \text{ for } i = 1, 2,$$


and

$$\frac{dp_i^C}{dt} = p_i^C \cdot [W_i^C(\vec{p}^R) - \phi(\vec{p}^C)], \text{ for } i = 1, 2.$$


For the payoff matrix given in figure 9.3.7a, the specific system of coupled differential equations is obtained by substituting the following values into the above:

		Column	
		$S_1$	$S_2$
Row	$S_1$	(1, 2)	(0, 0)
	$S_2$	(0, 0)	(2, 1)

(a) Battle of the Sexes.



(b) A two-population model with the interaction probabilities indicated.



(c) Phase diagram for the two-population, continuous linear replicator dynamics for Battle of the Sexes. The two stable rest points are  $\langle(1, 0), (1, 0)\rangle$  and  $\langle(0, 1), (0, 1)\rangle$ , denoted by solid circles; unstable rest points are indicated by unfilled circles.

**Figure 9.3.7**

A two-population version of the replicator dynamics.



$$\begin{aligned}
 W_1^R(\vec{p}^C) &= p_1^C, & W_1^C(\vec{p}^R) &= 2p_1^R, \\
 W_2^R(\vec{p}^C) &= 2p_2^C, & W_2^C(\vec{p}^R) &= p_2^R, \\
 \phi(\vec{p}^R) &= p_1^R p_1^C + 2p_2^R p_2^C, & \phi(\vec{p}^C) &= 2p_1^R p_1^C + p_2^R p_2^C.
 \end{aligned}$$

Figure 9.3.7c illustrates the phase diagram for Battle of the Sexes under the two-population replicator dynamics. Cressman (2003) provides a detailed investigation of multipopulation models with reference to extensive-form games.

### 5. Stability

One natural question to ask is what connection—if any—exists between rest points of the various dynamical systems under consideration and the Nash equilibria (and/or evolutionarily stable strategies) of the underlying game. As a preliminary, an important conceptual distinction needs to be made: the various dynamics considered model the evolutionary flows over population states, where a population state is a *distribution of strategies* over the individual players. In contrast, a Nash equilibrium strategy (or an evolutionarily stable strategy) is a *probability distribution* over pure strategies adopted by a player. These are very different entities. But, as we will see, as long as we are careful about the conceptual differences, important connections exist.

Inspection of figure 9.3.7c suggests one possible connection. The underlying game has three Nash equilibria:  $\langle(1,0), (1,0)\rangle$ ,  $\langle(0,1), (0,1)\rangle$ , and  $\langle\frac{1}{3}, \frac{2}{3}\rangle, \langle\frac{2}{3}, \frac{1}{3}\rangle$ . Of the three Nash equilibria, the first two feature evolutionarily stable strategies. If all three of these Nash equilibria are interpreted as *population states*, all correspond to rest points of the dynamics, with the two states corresponding to evolutionarily stable strategies being stable points of the dynamics. (The exact meaning of “stability” will be specified below.) However, the fact that additional rest points exist shows that, at least for the replicator dynamics, it is not the case that the rest points correspond to Nash equilibria.<sup>5</sup> A precise statement of the relationship between the two concepts requires introducing some definitions from the theory of dynamical systems.

A *rest point* of a dynamical system (also known as a *fixed point*) is a point  $\vec{p}$  such that  $dp_i/dt = 0$  for all  $i = 1, \dots, n$ . A stability concept provides a characterization of the behavior of a dynamical system in the local region around a fixed point. There are several natural stability concepts used in evolutionary game theory.

*Definition.* Let  $\vec{p} \in \Delta_n$  be a rest point of the replicator dynamics. Then  $\vec{p}$  is said to be *Lyapunov-stable* if for every  $\varepsilon > 0$ , there exists a  $\delta > 0$  such that  $|\vec{p}(0) - \vec{p}| < \delta$  implies  $|\vec{p}(t) - \vec{p}| < \varepsilon$  for all  $t \geq 0$ . That is,

every trajectory that is sufficiently close to the fixed point  $\vec{p}$  (i.e., no more than  $\delta$  away) remains close to  $\vec{p}$  (i.e., no more than  $\varepsilon$  away) for all future times.

The intuition behind Lyapunov stability is that there is not a local “push” away from the rest point. Another intuition about stability is that there is a local “pull” toward the rest point. This motivates the following:

*Definition.* Let  $\vec{p}$  be a fixed point of the replicator dynamics. Then  $\vec{p}$  is *asymptotically stable* if  $\vec{p}$  is Lyapunov-stable and, in addition, there exists a  $\delta > 0$  such that if  $|\vec{p}(0) - \vec{p}| < \delta$ , then  $\lim_{t \rightarrow \infty} \vec{p}(t) = \vec{p}$ .

Asymptotic stability captures one aspect of evolutionary stability: that the population is not easily displaced from the rest point corresponding to a particular distribution of strategies. A displacement from an asymptotically stable rest point, provided it is not too great, will return to the rest point in the limit.

*Definition.* If  $\vec{p}$  is Lyapunov-stable but not asymptotically stable, then  $\vec{p}$  is said to be *neutrally stable*.

In light of the Hofbauer et al. (1979) result that an evolutionarily stable strategy is one that is locally superior, we can introduce an analogous definition in the context of evolutionary games. In particular, Hofbauer and Sigmund (2002) propose the following:

*Definition.* Let  $\vec{p} \in \Delta_n$ , and let  $A$  be the payoff matrix for the underlying evolutionary game. Then  $\vec{p}$  is a *Nash equilibrium* if, for all  $\vec{q} \in \Delta_n$ , it is the case that

$$\vec{p} \cdot A \vec{p} \geq \vec{q} \cdot A \vec{p}.$$

The state  $\vec{p}$  is an *evolutionarily stable state* if, for all  $\vec{q} \neq \vec{p}$  in a neighborhood around  $\vec{p}$ ,

$$\vec{p} \cdot A \vec{q} > \vec{q} \cdot A \vec{q}.$$

Note that, despite the mathematical similarity between the various definitions, the term “evolutionarily stable state” is used to stress the difference between the population and the strategy interpretation. But, in effect, an evolutionarily stable state is the distribution corresponding to an evolutionarily stable strategy interpreted as the frequency with which pure strategies appear in a population. For example, in the case of the Hawk–Dove game, the evolutionarily stable strategy is the mixed strategy  $\sigma = \frac{1}{3}\text{ Hawk} + (1 - \frac{1}{3})\text{ Dove}$ , and the evolutionarily stable state is the population where  $\frac{1}{3}$  of the population follows Hawk and  $1 - \frac{1}{3}$  follows Dove.

Returning to the game of Rock–Paper–Scissors from figure 9.3.3, the phase diagram for the replicator dynamics is shown in figure 9.3.8. The rest point corresponding to the single Nash equilibrium of the game,  $\langle\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\rangle$ ,

although not an evolutionarily stable state, is neutrally stable. Any small displacement from the rest point in the center of the simplex will result in the population cycling around the rest point in a stable orbit.

These two observations provide an illustration of the following theorem, proven by Hofbauer and Sigmund (2002):

*Theorem 6.* Let  $A$  be the payoff matrix for a symmetric two-player game. Then

if  $\vec{p} \in \Delta_n$  is a Nash equilibrium of  $A$ , then  $\vec{p}$  is a rest point of the replicator dynamics;

if  $\vec{p}$  is Lyapunov-stable under the replicator dynamics, then it is a Nash equilibrium of the underlying game.

One can also prove (see Hofbauer & Sigmund, 2002) that the following connection holds between evolutionarily stable states and asymptotic stability, for the replicator dynamics:

*Theorem 7.* Let  $A$  be the payoff matrix for a symmetric two-player game. If  $\vec{p} \in \Delta_n$  is an evolutionarily stable state for  $A$ , then  $\vec{p}$  is asymptotically stable under the replicator dynamics with  $A$  as the underlying game.

It turns out that if one adopts a slightly stronger stability concept than that of an evolutionarily stable state, the above result for the replicator dynamics can be extended to multiple families of dynamics. Taylor and Jonker (1978) define a *regular ESS* as follows:

*Definition.* Let  $S = \{1, \dots, n\}$  be a set of pure strategies for a symmetric two-player game with payoff matrix  $A$ . Let  $\sigma$  be a Nash equilibrium for  $A$ . Then  $\sigma$  is a *regular ESS* if  $\pi(i|\sigma) < \pi(\sigma|\sigma)$  whenever  $i \notin \text{supp}(\sigma)$ , and  $\vec{x} \cdot A\vec{x} < 0$  whenever  $\text{supp}(x) \subset \text{supp}(\sigma)$ ,  $\vec{x} \neq \vec{0}$  and  $\sum x_i = 0$ .

The concept of a regular ESS is a slight strengthening of the ESS concept. Every regular ESS is an ESS, but not

conversely. However, Taylor and Jonker note that *most* ESSs are regular, and if  $\pi$  is a payoff function for which a nonregular ESS exists, one can find another payoff function  $\pi'$ , arbitrarily close to  $\pi$ , that only has regular ESSs. Intuitively, a regular ESS is an ESS  $\sigma$  where the support of  $\sigma$  contains *all* of the pure strategies that are best responses to  $\sigma$ .


One can then prove the following:

*Theorem 8 (Hofbauer & Sandholm, 2011).* Let  $\sigma$  be a regular ESS for an evolutionary game  $G$ . Then,  $\sigma$  is asymptotically stable under

- any impartial pairwise comparison dynamic for  $G$ ;
- any separable excess payoff dynamic for  $G$ ;
- the best response dynamic for  $G$ .

While the generality of this result is noteworthy, it must be stressed that what this establishes is a validation of “the use of regular ESS as a blanket sufficient condition for local stability under evolutionary game dynamics” (Sandholm, 2010b). Local stability is not necessarily a good guide to the expected behavior of a population in an evolutionary game. Consider, for example, the game shown in figure 9.3.9: although  $S_3$  is a regular ESS and is asymptotically stable, the *basin of attraction* for  $S_3$ —that is, the set of all points that converge to  $S_3$  as  $t \rightarrow \infty$ —can be made arbitrarily small by adjusting the value of  $\varepsilon$  accordingly.

Furthermore, it can also be the case that the basin of attraction for a population state is the entire interior of the simplex space, even though the population state is *not* asymptotically stable. Figure 9.3.10 illustrates a game where, under the replicator dynamics, the long-term behavior is such that the population will always converge to  $A$ . Yet  $A$  is not asymptotically stable, as small




**Figure 9.3.8**

For the game of Rock–Paper–Scissors, the population state in which all three strategies are equally represented is a neutrally stable state of the replicator dynamics.

	$S_1$	$S_2$	$S_3$
$S_1$	(0, 0)	(1, 0)	( $\epsilon$ , 0)
$S_2$	(0, 1)	(1, 1)	(-0, 0)
$S_3$	(0, $\epsilon$ )	(0, 0)	(2 $\epsilon$ , 2 $\epsilon$ )

(a) The payoff matrix.




(b) The basin of attraction for  $S_3$  appears on the right of the simplex and can be made arbitrarily small by adjusting the value of  $\epsilon$  accordingly. Nevertheless,  $S_3$  is a regular ESS and is asymptotically stable.

**Figure 9.3.9**

An ESS can have an arbitrarily small basin of attraction.

	$S_1$	$S_2$	$S_3$
$S_1$	(0, 0)	(0, 1)	(2, 0)
$S_2$	(1, 0)	(0, 0)	(0, 0)
$S_3$	(0, 2)	(0, 0)	(1, 1)

(a) The payoff matrix.



(b) The basin of attraction for  $S_2$  is the whole interior of the simplex, even though  $S_2$  is not asymptotically stable. Small displacements from  $S_2$  result in the population following a lengthy trajectory, which carries it very far away from the all- $S_2$  state before eventually returning. (The trajectories ultimately converge to the all- $S_2$  point; the trajectories appear to stop in the middle of the simplex due to numeric limits on the differential equation solver.)

**Figure 9.3.10**

A population state can be the expected evolutionary outcome without being asymptotically stable.

displacements may result in the population being carried very far away before eventually returning via a circuitous route.

Evolutionary game theory provides a great variety of tools for modeling the boundedly rational behavior of populations in strategic contexts. Yet, as we have seen, the evolutionary outcomes may not always be rational: strictly dominated strategies may be played with positive probability, as figure 9.3.7b shows. The fact that a population state is a (regular) evolutionarily stable state does not guarantee that it is, in any way, the expected or even likely evolutionary outcome, as seen in figure 9.3.9. And the long-term convergence behavior of a population can also be decoupled from asymptotic stability, as in figure 9.3.10. With great variety comes great complexity and hence the need to exercise care when drawing inferences about expected outcomes.

### Acknowledgments

I thank William Sandholm, Emin Dokumaci, and Francisco Franchetti for their open-source software suite *Dynamo*, which was used to create many figures in this article (see Franchetti & Sandholm, 2013).

### Notes

1. It should be noted that, in his collection of essays *On Evolution*, published in 1972, Maynard Smith gives the same definition of an ESS in the essay “Game Theory and the Evolution of Fighting.” But he clearly notes in the introduction to the volume that “I would probably not have had the idea for this essay if I had not seen an unpublished manuscript on the evolution of fighting by Dr George Price. . . . Unfortunately, Dr Price is better at having ideas than at publishing them. The best I can do therefore is to acknowledge that if there is anything in the idea, the credit should go to Dr Price and not to me.”
2. Whether this would, in fact, happen depends on specific features of the dynamics.
3. In their 1973 paper, the game is called the “Hawk–Mouse” game as a result of George Price’s objection to using the term “Dove,” due to its religious connotations.
4. It is assumed that the losing Hawk is selected at random with equal probability, which is why the payoffs in the upper-left cell of the payoff matrix take the form that they do.
5. The two additional rest points are where the entire Row population follows  $S_2$  and the Column population follow  $S_1$ , and vice versa. These are unstable rest points, but they exist because the replicator dynamics cannot introduce strategies into the population if they are originally absent.

### References

- Akin, E. (1980). Domination or equilibrium. *Mathematical Biosciences*, 50(3–4), 239–250.
- Brown, G. W., & von Neumann, J. (1950). Solutions of games by differential equations. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory of games* (Vol. 1, Annals of Mathematics Studies, Nr. 24, pp. 73–80). Princeton, NJ: Princeton University Press.
- Cressman, R. (2003). *Evolutionary dynamics and extensive form games*. Cambridge, MA: MIT Press.
- Fisher, R. A. (1930). *The genetical theory of natural selection*. Oxford, England: Oxford Clarendon Press.
- Franchetti, F., & Sandholm, W. H. (2013). An introduction to *Dynamo*: Diagrams for evolutionary game dynamics. *Biological Theory*, 8, 167–178.
- Gintis, H. (2009). *Game theory evolving* (2nd ed.). Princeton, NJ: Princeton University Press.
- Hofbauer, J., & Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6(3), 341–377.
- Hofbauer, J., Schuster, P., & Sigmund, K. (1979). A note on evolutionary stable strategies and game dynamics. *Journal of Theoretical Biology*, 81, 609–612.
- Hofbauer, J., & Sigmund, K. (2002). *Evolutionary games and population dynamics*. Cambridge, England: Cambridge University Press.
- Maynard Smith, J. (1972). *On evolution*. Edinburgh, Scotland: Edinburgh University Press.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge, England: Cambridge University Press.
- Maynard Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature*, 246, 15–18.
- Sandholm, W. H. (2010a). Local stability under evolutionary game dynamics. *Theoretical Economics*, 5(1), 27–50.
- Sandholm, W. H. (2010b). *Population games and evolutionary dynamics*. Cambridge, MA: MIT Press.
- Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78(1), 130–156.
- Smith, M. J. (1984). The stability of a dynamic model of traffic assignment—An application of a method of Lyapunov. *Transportation Science*, 18(3), 245–252.
- Taylor, P. D., & Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1–2), 145–156.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Weibull, J. W. (1995). *Evolutionary game theory*. Cambridge, MA: MIT Press.

This is a section of [doi:10.7551/mitpress/11252.001.0001](https://doi.org/10.7551/mitpress/11252.001.0001)

# The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn

## Citation:

*The Handbook of Rationality*

Edited by: Markus Knauff, Wolfgang Spohn


DOI: 10.7551/mitpress/11252.001.0001

ISBN (electronic): 9780262366175

Publisher: The MIT Press

Published: 2021

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.



The MIT Press



© 2021 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Knauff, Markus, editor. | Spohn, Wolfgang, editor.

Title: The handbook of rationality / edited by Markus Knauff and Wolfgang Spohn.

Description: Cambridge : The MIT Press, 2021. | Includes bibliographical references and index.

Identifiers: LCCN 2020048455 | ISBN 9780262045070 (hardcover)

Subjects: LCSH: Reasoning (Psychology) | Reason. | Cognitive psychology. | Logic. | Philosophy of mind.

Classification: LCC BF442 .H36 2021 | DDC 153.4/3—dc23

LC record available at <https://lcn.loc.gov/2020048455>