

Fast Convergence in Evolutionary Equilibrium Selection

Gabriel E. Kreindler and H. Peyton Young

September 26, 2011

Abstract. Stochastic learning models provide sharp predictions about equilibrium selection when the noise level of the learning process is taken to zero. The difficulty is that, when the noise is extremely small, it can take an extremely long time for a large population to reach the stochastically stable equilibrium. An important exception arises when players interact locally in small close-knit groups; in this case convergence can be rapid for small noise and an arbitrarily large population. We show that a similar result holds when the population is fully mixed and there is no local interaction. Selection is sharp and convergence is fast when the noise level is ‘fairly’ small but not extremely small.

1. Stochastic stability and equilibrium selection

Evolutionary models with random perturbations provide a useful framework for explaining how populations reach equilibrium from out-of-equilibrium conditions, and why some equilibria are more likely than others to be observed in the long run. Individuals in a large population interact with one another repeatedly to play a given game, and they update their strategies based on information about what others are doing. The updating rule is usually assumed to be a myopic best reply with a random component resulting from errors, unobserved utility shocks, or experiments. These assumptions lead to a stochastic dynamical system whose long-run behavior can be analyzed using the theory of large deviations (Freidlin and Wentzell 1984). The key idea is to examine the limiting behavior of the process when the random component becomes vanishingly small. This typically leads to powerful selection results, that is, the limiting ergodic distribution tends to be concentrated on particular equilibria (often a unique equilibrium) that are *stochastically stable* (Foster and Young 1990, Kandori, Mailath and Rob 1993, Young 1993).

This approach has a potential drawback however. When the noise level is very small, it can take a very long time for the process to reach the stochastically stable equilibrium. Another way of saying this is that the mixing time of the stochastic process becomes unboundedly large as the noise level is taken to zero. Nevertheless, this leaves open an important question: must the noise actually be close to zero in order to obtain sharp selection results? This assumption is needed to characterize the stochastically stable states theoretically, but it is not clear that it is a necessary condition for these states to occur with high probability. It could be that at intermediate levels of noise the learning process displays a fairly strong bias towards the stochastically stable states. If this is so, it

might not take very long for this bias to manifest itself, and the speed of convergence could be quite rapid.

A pioneering paper by Ellison (1993) shows that this is indeed the case when agents are situated at the nodes of a network and they interact ‘locally’ with small groups of neighbors. Ellison considers the specific situation where agents are located around a ring and they interact with agents who lie within a specified distance. The form of interaction is a symmetric 2×2 coordination game. The reason that this set-up leads to fast learning is local reinforcement. Once a small close-knit group of interconnected agents has adopted the stochastically stable action, say A , they continue to play A with high probability even when people outside the group are playing the alternative action B . When the noise level is sufficiently small (but not taken to zero), it takes a relatively short time in expectation for any such group to switch to A and to stay playing A with high probability thereafter. Since this occurs in parallel across the entire network, it does not take long in expectation until almost all of the close-knit groups have switched to A . Thus if everyone (or almost everyone) is in such a group, A spreads in short order to a large proportion of the population. In fact this argument is quite general and applies to a variety of network structures and stochastic learning rules, as shown in Young (1998, 2011).

The contribution of the present paper is to show that, even when the assumption of local interaction is dropped completely, rapid convergence still holds. To be specific, suppose that agents play a symmetric 2×2 coordination game and obtain information about the current state from a random sample of the population rather than from a fixed group of neighbors. Suppose further that they update based on the logit response function, which is a standard learning rule in this literature (Blume 1993, 1995). The expected time it takes to get to the state where most people are playing A depends on two parameters: the level of noise ε , and the difference in potential α between the two equilibria. We claim that when the noise level is small but not too small, selection is fast for all positive values of α . In other words, with high probability the process transits in relatively few steps from everyone playing B to a majority playing A , and the expected time T is bounded independently of the population size.

In addition, the dynamics exhibit a phase transition in the payoff gain: for any given level of noise ε there is a critical value $\alpha^*(\varepsilon)$ such that selection is fast (bounded in the population size) if α is larger than $\alpha^*(\varepsilon)$, and slow (unbounded in the population size) if α is at most α^* . For intermediate levels of noise it turns out that the critical value is zero: selection is fast for all games provided that there is *some* gain in potential between the two equilibria. Above the critical value, the process is able to escape from the all- B equilibrium quite quickly, and once it has reached a state where most agents are playing A . We provide a close estimate of this critical value, and we also derive an upper-bound on the number of steps. Simulations show that this upper bound is accurate over a wide range of parameter values. Moreover, the absolute magnitudes are small and surprisingly close to those obtained in local interaction models (Ellison 1993).

The paper unfolds as follows. We begin with a review of related literature in section 2. Section 3 sets up the model, and section 4 contains the first main result, namely the existence and estimation of a critical payoff gain when agents have perfect information. We derive an upper bound on the number of steps to get close to equilibrium in section 5. Section 6 extends the results in the previous two sections to the imperfect information case, and section 7 concludes.

2. Related literature

The rate at which a coordination equilibrium becomes established in a large population (or whether it becomes established at all) has been studied from a variety of perspectives. To understand the connections with the present paper we shall divide the literature into several parts, depending on whether interaction is assumed to be global or local, and on whether the learning dynamics are deterministic or stochastic. In the latter case we shall also distinguish between those models in which the stochastic perturbations are taken to zero (*low noise dynamics*), and those in which the perturbations are maintained at a positive level (*noisy dynamics*).

To fix ideas, let us consider the situation where agents interact in pairs and play a fixed 2×2 symmetric pure coordination game G of the following form:

$$\begin{array}{cc}
 & \begin{array}{c} A \\ B \end{array} \\
 \begin{array}{c} A \\ B \end{array} & \begin{array}{cc} 1 + \alpha, 1 \\ + \alpha & 0,0 \end{array} \\
 & \begin{array}{c} 0,0 \\ 1,1 \end{array}
 \end{array}
 , \quad \alpha \geq 0$$

We can think of B as the “status quo”, of A as the “innovation” and of α as the *payoff gain* of the innovation relative to the status quo (this representation is in fact without loss of generality, as we show in section 3).

Local interaction refers to the situation where agents are located at the nodes of a network and they interact only with their neighbors. *Global interaction* refers to the situation where agents react to the distribution of actions in the entire population, or to a random sample of such actions.

Virtually all of the results about waiting times and rate of convergence can be discussed in this setting. The essential question is how long it takes to transit from the all- B equilibrium to a state where most of the agents are playing A .

Deterministic dynamics, local interaction

Morris (2000) studies deterministic best-response dynamics on an infinite network. Each node of the network is occupied by an agent, and in each period all agents myopically best respond to their neighbors’ actions. (Myopic best response, either with or without perturbations, is assumed throughout this literature.) Morris studies the threshold $\bar{\alpha}$ such that, for any payoff gain $\alpha > \bar{\alpha}$, there exists some finite group of initial adopters from which the innovation spreads by ‘contagion’ to the entire population. He provides several characterizations of $\bar{\alpha}$ in terms of topological properties of the network, such as the existence of cohesive (inward looking) subgroups, the uniformity of the local interaction, and the growth rate of the number of players who can be reached in k steps. We note that Morris does not address the issue of waiting times as such, rather only when and whether full adoption can occur.

Deterministic dynamics, global interaction

López-Pintado (2006) considers a class of models with deterministic best-response dynamics and a continuum of agents. She studies a deterministic (mean-field) approximation of a large, finite system where agents are linked by a random network with a given degree distribution, and in each time period agents best respond to their neighbors’ actions. López-Pintado proves that, for a given distribution of sample sizes, there exists a minimum threshold value of α above which, for any initial

fraction of A -players (however small), the process evolves to a state in which a positive proportion of the population plays A forever. This result is similar in spirit to Morris', but it employs a mean-field approach and global sampling rather than a fixed network structure.

Jackson and Yariv (2007) use a similar mean-field approximation technique to analyse models of innovation diffusion. They identify two types of equilibrium adoption levels: stable levels of adoption and unstable ones (tipping points). A smaller tipping point facilitates diffusion of the innovation, while a larger stable equilibrium corresponds to a higher final adoption level. Jackson and Yariv derive comparative statics results on how the network structure changes the tipping points and stable equilibria.

Watts (2002) and Lelarge (2010) study deterministic best-response dynamics on large random graphs with a specified degree distribution. In particular, Lelarge analyzes large but finite random graphs and characterizes in terms of the degree distribution the threshold value $\bar{\alpha}$ such that, for any payoff gain $\alpha > \bar{\alpha}$, with high probability, a single player who adopts A leads the process to a state in which a positive proportion of the population is playing A .

Stochastic dynamics, local interaction

Ellison (1993) and Young (1998, 2011) study adoption dynamics when agents are located at the nodes of a network. Whenever agents revise, they best respond to their neighbors' actions, with some random error. Unlike most models discussed above, the population is *finite* and the learning process is *stochastic*. The aim of the analysis is to characterize the ergodic distribution of the process rather than to approximate it by a mean-field deterministic dynamic, and to study whether convergence to this distribution occurs in a 'reasonable' amount of time.

Ellison examines the case where agents best respond with a uniform probability of error (choosing a nonbest reply), while Young focuses on the situation where agents use a logit response rule. The latter implies that the probability of making an error decreases as the payoff loss from making the error increases. In both cases, the main finding is that, when the network consists of small close-knit groups that interact mainly with each other rather than outsiders, then for intermediate levels of noise the process evolves quite rapidly to a state where most agents are playing A independently of the size of the population, and independently of the initial state.

Montanari and Saberi (2010) consider a similar situation: agents are located at the nodes of a fixed network and they update asynchronously using a logit response function. The authors characterize the expected waiting time to transit from all- B to all- A as a function of population size, network structure, and the size of the gain α . Like Ellison and Young, Montanari and Saberi show that local clusters tend to speed up the learning process, whereas overall well-connected networks tend to be slow. For example, learning is slow on random networks for small enough $\alpha > 0$, while learning on small-world networks – networks where agents are mostly connected locally but there also exist a few random 'distant' links – becomes slower as the proportion of distant links increases. These results stand in contrast with the dynamics of disease transmission, where contagion tends to be fast in well-connected networks and slow in localized networks (Anderson and May 1991, see also Watts and Dodds 2007).

The analytical framework of Montanari and Saberi differs from that of Ellison and Young in one crucial respect however: in the former the waiting time is characterized *as the noise is taken to zero*, whereas in the latter the noise is held fixed at a small but not arbitrarily small level. This difference has important implications for the magnitude of the expected waiting time: if the noise is extremely small, it takes an extremely long time in expectation for even one player to switch to A given that all his neighbors are playing B . Montanari and Saberi show that when the noise is vanishingly small, the expected waiting time to reach all- A is independent of the population size for some types of networks and not for others. However, their method of analyzing this issue requires that the *absolute* magnitude of the waiting time is very long in either case.

In summary, the previous literature has dealt mainly with four cases:

- i) Noisy learning dynamics and local interaction (Ellison, Young),
- ii) Deterministic dynamics and local interaction (Morris),
- iii) Low noise dynamics and both local and global interaction (Montanari and Saberi), and
- iv) Deterministic dynamics and global interaction (López-Pintado, Lelarge, Jackson and Yariv).

There remains the case of noisy dynamics with global interaction. Shah and Shin (2010) study a variant of logit learning where the rate at which players revise their actions changes over time. Focusing on a special class of potential games which does not include the case considered here, they prove that for intermediate values of noise the time to get close to equilibrium grows slowly (but is unbounded) in the population size.

The contribution of the present paper is to show that in the case of noisy dynamics with global interaction learning can be very rapid when the noise is small but not arbitrarily small. Furthermore the expected number of steps to reach the mostly- A state is similar in magnitude to the number of steps to reach such a state in local interaction models – on the order of 20-50 periods for an error rate around 5% – where each individual revises once per period in expectation. The conclusion is that fast learning occurs under global as well as local interaction for realistic (non-vanishing) noise levels.

3. The Model

Consider a large population of N agents. Each agent chooses one of two available actions, A and B . Interaction is given by a symmetric 2×2 coordination game with payoff matrix

	A	B
A	a, a	c, d
B	d, c	b, b

where $a > d$ and $b > c$. This game has the potential function

	A	B
A	$a - d$	0
B	0	$b - c$

Define the normalized potential gain associated with passing from the (B, B) equilibrium to the (A, A) equilibrium

$$(1) \quad \alpha = \frac{(a - d) - (b - c)}{b - c}$$

Without loss of generality assume that $a - d > b - c$, or equivalently $\alpha > 0$. This makes (A, A) the risk-dominant equilibrium; note that (A, A) need not be the same as the Pareto-dominant equilibrium. Standard results in evolutionary game theory say that the (A, A) equilibrium will be selected in the long run (Blume 2003; see also Kandori, Mailath and Rob 1993 and Young 1993).

A particular case of special interest occurs when the game is a *pure* coordination game with payoff matrix

$$(2) \quad \begin{array}{cc} & \begin{array}{c} A \\ B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{array}{cc} 1 + \alpha, 1 + \alpha & 0, 0 \\ 0, 0 & 1, 1 \end{array} \end{array}$$

We can think of B as the “status quo” and of A as the “innovation”, and in this case $\alpha > 0$ is also the *payoff gain* of the adopting the innovation relative to the status quo. The potential function in this case is proportional to the potential function in the general case, which implies that under logit learning and a suitable rescaling of the noise parameter, the two settings are equivalent. For the rest of this paper we will work with the game form in (2).

Agents revise their actions in the following manner. At times $t = \frac{k}{N}$ with $k \in \mathbb{N}$, and only at these times, one agent is randomly (independently over time) chosen to revise his action.¹ When revising, an agent gathers information about the current state of play. We consider two possible information structures. In the *full information* case, revising agents know the current proportion of adopters in the entire population. In the *partial information* case, revising agents randomly sample d other agents from the population (with replacement), and learn their current actions, where d is a positive integer that is independent of N .

After gathering information, a revising agent i calculates the fraction x of agents in his sample who are playing A , and chooses a noisy best response given by the logit model:

$$(3) \quad \Pr(i \text{ chooses } A \mid x) = f(x; \alpha, \beta) = \frac{e^{\beta(1+\alpha)x}}{e^{\beta(1+\alpha)x} + e^{\beta(1-x)}}$$

where $1/\beta$ is a measure of the noise in the revision process. For convenience we will sometimes drop the dependence of f on β and simply write $f(x; \alpha)$, or on both α and β and write $f(x)$. Denote $\varepsilon = \frac{1}{1+e^\beta}$ the associated error rate at zero adoption rate; given the bijective correspondence

¹ An alternative revision protocol runs as follows: time is continuous, and each agent has a Poisson clock that rings once per period in expectation. When an agent’s clock rings the agent revises his action. It is possible to show that results in this article remain unchanged under this alternative revision protocol.

between β and ε , we will use the two variables interchangeably to refer to the noise level in the system.

The logit model is one of the two models predominantly used in the literature. The other is the uniform error model (Kandori, Mailath and Rob 1993, Young 1993, Ellison 1993), which posits that agents make errors with a fixed probability. A characteristic feature of the logit model is that the probability of making an error is sensitive to the payoff difference between choices, making costly errors less probable; from an economic perspective, this feature is quite natural (Blume 1993, 1995, 2003). Another feature of logit is that it is a smooth response, whereas in the uniform error model an agent's decision changes abruptly around the indifference point. Finally, the logit model can also be viewed as a pure best-response to a noisy payoff observation. Specifically, if the payoff shocks ϵ_A and ϵ_B are independently distributed according to the extreme-value distribution given by $\Pr(\epsilon_i \geq z) = \exp(-\exp(\beta z))$, then this leads to the logit probabilities (Brock and Durlauf 2001, Anderson, Palma and Thisse 1992, McFadden 1976).

The revision process just described defines a stochastic process $\Gamma_N(\alpha, \beta)$ in the full information case and $\Gamma_N(\alpha, \beta, d)$ in the partial information case. The states of the process are the *adoption rates* $x_N(t) \in \left\{0, \frac{1}{N}, \dots, \frac{N-1}{N}, 1\right\}$, and by assumption the process starts in the *all-B* state, namely $x_N(0) = 0$.

We now turn to the issue of speed of convergence, measured in terms of the expected time until a large fraction of the population adopts action *A*. This measure is appropriate because the probability of being in the all-*A* state is extremely small. Formally, for any $p < 1$ define the random *hitting time*²

$$T_N(\alpha, \beta, p) = \min\{t : x_N(t) \geq p\}$$

Fast learning is defined as follows.

Definition 1. The family $\{\Gamma_N(\alpha, \beta) : N > 0\}$ has *fast learning* if there exists $S = S(\alpha, \beta)$ such that the expected waiting time until a majority of agents play *A* under process $\Gamma_N(\alpha, \beta)$ is at most S independently of N , or $ET_N\left(\alpha, \beta, \frac{1}{2}\right) < S$ for all N .

More generally, for any $p < 1$,

Definition 2. The family $\{\Gamma_N(\alpha, \beta) : N > 0\}$ has *fast learning to p* if there exists $S = S(\alpha, \beta, p)$ such that the expected waiting time until at least a fraction p of agents play *A* under process $\Gamma_N(\alpha, \beta)$ is at most S independently of N , or $ET_N(\alpha, \beta, p) < S$ for all N .

Note. When the above conditions are satisfied then we say, by a slight abuse of language, that $\Gamma_N(\alpha, \beta)$ has fast learning, or fast learning to p .

² The results in this paper continue to hold under the following stronger definition of waiting time. Given $p < 1$ let $\tilde{T}(\alpha, \beta, p)$ be the expected first time such that at least the proportion p has adopted and at all later times the probability is at least p that the proportion p has adopted.

4. Full information

The following theorem establishes how much better than the status quo an innovation needs to be in order for it to spread quickly in the population. Specifically, fast learning can occur for *any* noise level as long as the payoff gain exceeds a certain threshold; moreover this threshold is equal to zero for intermediate noise levels.

Theorem 1. If $\alpha > h(\beta)$ then $\Gamma_N(\alpha, \beta)$ displays fast learning, where

$$h(\beta) = \begin{cases} \frac{e^{\beta-1} + 4 - e}{\beta} - 2, & \beta > 2 \\ 0, & \beta \leq 2 \end{cases}$$

Moreover, when $\beta \geq 3$ (hence $\varepsilon < 5\%$) and $\alpha > h(\beta)$ then $\Gamma_N(\alpha, \beta)$ displays fast learning to 99%.

The main message of Theorem 1 is that fast learning holds in settings with *global* interaction. This result does not follow from previous results in models of *local* interaction (Ellison 1993, Young 1998, Montanari and Saberi 2010). Indeed, a key component of local interaction models is that agents interact only with a small, fixed group of neighbors, whereas here each agent observes the actions of the entire population. Theorem 1 is nevertheless reminiscent of results from models of local interaction. For example, Young 1998 shows that for certain families of local interaction graphs learning is fast for any positive payoff gain $\alpha > 0$. Theorem 1 shows that fast learning can occur for any positive payoff gain even when interaction is global, provided that $\beta \leq 2$, which is equivalent to an error rate larger than approximately 11.92%.

In the proof of Theorem 1 we show that for each noise level there exists a payoff gain threshold, denoted $\alpha^*(\beta)$, such that $\Gamma_N(\alpha, \beta)$ has fast learning for $\alpha > \alpha^*(\beta)$. Figure 1 shows the simulated payoff gain threshold $\alpha^*(\beta)$ (blue, dashed line) as well as the bound $h(\beta)$ (red, solid line). The x-axis displays the error rate $\varepsilon = \frac{1}{1+e^\beta}$, and the y-axis represents payoff gains α . Note that the difference between the two curves never exceeds about 15 percentage points.




Figure 1 - The critical payoff gain for fast learning as function of error rate. Simulation (blue, dashed) and upper bound (red, solid).

Theorem 1 shows that when the payoff gain is above a specific threshold, the time until a high proportion of players adopt A is bounded independently of the population size N . Simulations reveal that for realistic parameter values the expected waiting time can be very small.

Figure 2 shows a typical adoption path. It takes, on average, less than 20 revisions per capita until $p = 99\%$ of the population plays A , for a population size of $N = 1000$, with payoff gain $\alpha = 100\%$ and error rate $\varepsilon = 5\%$.




Figure 2 – Adoption path to 99%, $\alpha = 100\%$, $\varepsilon = 5\%$ ($\beta \approx 3$), $N = 1000$.

More generally, Table 1 shows how the expected waiting time depends on the population size N , the payoff gain α , the error rate ε , and on the target adoption level p . The main takeaway is that the absolute magnitude of the waiting times is small. We explore this effect in more detail in section 5.

Table 1 – Average waiting times (full information)

Average waiting time ($\varepsilon = 5\%$)			
	$N = 100$	$N = 1000$	$N = 10,000$
$\alpha = 70\%$ ^{a)}	33	101	> 8,000
$\alpha = 80\%$ ^{a)}	25	36	35
Average waiting time ($\varepsilon = 10\%$)			
	$N = 100$	$N = 1000$	$N = 10,000$
$\alpha = 4\%$ ^{b)}	38	190	> 8,000
$\alpha = 25\%$ ^{c)}	19	20	21

N = population size, α = innovation payoff gain, ε = error rate. The target adoption rate is ^{a)} $p = 99\%$, ^{b)} $p = 50\%$ and ^{c)} $p = 90\%$. Each result is based on 100 simulations.

Note that Ellison 1993 obtains surprisingly similar simulation results in the case of local learning – most of the waiting times he presents lie between 10 and 50 (see figure 1, tables III and IV in that paper). Although the results are similar, the assumptions in the two models are very different. In Ellison’s model agents are located around a circle or at the nodes of a lattice and interact only with close neighbors. Also, he uses the uniform error model instead of logit learning. Finally, in his simulations the target adoption rate is $p = 75\%$, the payoff gain is $\alpha = 100\%$, and he presents results for error rates $\varepsilon = 1.25\%$, 2.5% and 5% .³

Proof of Theorem 1. The proof consists of two steps. First, we show that the results hold for a deterministic approximation of the stochastic process $\Gamma_N(\alpha, \beta)$. The second step is to show that fast learning is preserved by this approximation when N is large.

We begin by defining the deterministic approximation and the concepts of equilibrium, stability and fast learning in this setting. The deterministic process is denoted $\Gamma(\alpha, \beta)$ and has state variable $x(t)$. The process evolves in continuous time, and $x(t)$ is the adoption rate at time t . By assumption we take $x(0) = 0$.

In the process $\Gamma_N(\alpha, \beta)$, the probability that a revising agent chooses A when the population adoption rate is x is equal to $f(x; \alpha, \beta)$. Definition (3) can be rewritten as

$$(4) \quad f(x; \alpha, \beta) = \frac{1}{1 + e^{\beta(1-(\alpha+2)x)}}$$

This function depends on α and β , but it does not depend on N . For convenience, we shall sometimes omit the dependence of f on α and/or β in the notation. We define the deterministic dynamic by the differential equation

$$(5) \quad \dot{x} = f(x; \alpha, \beta) - x$$

where the dot above x denotes the time derivative.⁴

An *equilibrium* of this system is a rest point, that is an adoption rate x^* satisfying $\dot{x}^* = 0$, which is equivalent to $f(x^*) = x^*$. An equilibrium x^* is *stable* if after any small enough perturbation the process converges back to the same equilibrium. An equilibrium is *unstable* if the process never converges back to the same equilibrium after any (non-trivial) perturbation. Given that f is continuously differentiable, x^* is stable if and only if $f'(x^*)$ is strictly below 1. Similarly, x^* is unstable if and only if $f'(x^*)$ is strictly above 1. It is easy to see that there always exists a stable equilibrium.

The definitions of fast learning from the stochastic setting extend naturally to the deterministic case. The *hitting time* to reach an adoption rate p is

$$T(\alpha, \beta, p) = \min\{t : x(t) \geq p\}$$

³ These error rates correspond to randomization probabilities of 2.5%, 5% and 10% respectively.

⁴ Note that the expected change in the adoption rate in the stochastic process is given by $\frac{1}{N}(f(x) - x)$. More precisely, $\Delta_k = x_N\left(\frac{k+1}{N}\right) - x_N\left(\frac{k}{N}\right)$ takes values $\pm \frac{1}{N}$ with probabilities $(1-x)f(x)$ and $x(1-f(x))$ respectively, where $x = x_N\left(\frac{k}{N}\right)$. It follows that $E\Delta_k = \frac{1}{N}(f(x) - x)$.

Definition 3: The process $\Gamma(\alpha, \beta)$ displays *fast learning* if it reaches $x = \frac{1}{2}$ in finite time, that is $T(\alpha, \beta, \frac{1}{2}) < \infty$. Analogously, the process $\Gamma(\alpha, \beta)$ exhibits *fast learning to p* if $T(\alpha, \beta, p) < \infty$.

Remark: A necessary and sufficient condition for fast learning is that all equilibria lie strictly above $\frac{1}{2}$. Similarly, fast learning to p holds if and only all equilibria lie strictly above p . Indeed, clearly if $x^* \leq \frac{1}{2}$ ($x^* \leq p$) is an equilibrium, and $x(0) = 0$, then $x(t) < x^*$ for all t . Conversely, by uniform continuity the process always reaches $x = \frac{1}{2}$ ($x = p$) in finite time.

The following lemma shows that the deterministic process has at most three equilibria (rest points).

Lemma 1. For $\alpha > 0$ the process $\Gamma(\alpha, \beta)$ has a unique equilibrium x_H in the interval $(\frac{1}{2}, 1]$, and this equilibrium is stable (it is referred to as *the high equilibrium*). Furthermore, there exist at most two equilibria in the interval $[0, \frac{1}{2}]$.




Figure 3 – The adoption curves for payoff gain $\alpha = 20\%$ (left panel) and $\alpha = 120\%$ (right panel), error rate $\varepsilon = 5\%$ ($\beta \approx 3$). The two systems have three equilibria and a unique equilibrium, respectively.

Proof. For future reference, from (4) we find that for all x

$$(6) \quad f'(x) = \beta(\alpha + 2)f(x)(1 - f(x)),$$

$$(7) \quad f''(x) = \beta(\alpha + 2)f'(x)(1 - 2f(x)).$$

Note that $f(\frac{1}{2+\alpha}) = \frac{1}{2}$, hence identity (7) implies that

$$(8) \quad \text{The function } f \text{ is strictly convex below } \frac{1}{2+\alpha} \text{ and strictly concave above } \frac{1}{2+\alpha}.$$

To determine how many equilibria may exist in a given interval $[a, b]$ on which f is either strictly convex or strictly concave, we look at the signs of $f(a) - a$ and $f(b) - b$.

On the interval $\left[0, \frac{1}{2+\alpha}\right)$ the function f is strictly convex, and moreover

$$f(0) > 0 \quad \text{and} \quad f\left(\frac{1}{2+\alpha}\right) = \frac{1}{2} > \frac{1}{2+\alpha}.$$

It follows that there are at most two equilibria below $\frac{1}{2+\alpha}$.

Note that $f\left(\frac{1}{2+\alpha}\right) = \frac{1}{2}$, hence no point in the interval $\left(\frac{1}{2+\alpha}, \frac{1}{2}\right]$ can be an equilibrium, because f is increasing.

Finally, on the interval $\left[\frac{1}{2}, 1\right]$ the function f is strictly concave, and moreover

$$f\left(\frac{1}{2}\right) > \frac{1}{2} \quad \text{and} \quad f(1) < 1$$

It follows that there exists exactly one equilibrium above $\frac{1}{2}$. This equilibrium is stable because it corresponds to a down-crossing. ■

Consider the set

$$A(\beta) = \{\alpha : \Gamma(\alpha, \beta) \text{ displays fast learning}\}$$

We claim that if $\alpha \in A(\beta)$ and $\alpha < \alpha'$ then also $\alpha' \in A(\beta)$. First, from equation (4) we see that

$$(9) \quad \frac{\partial f}{\partial \alpha}(x; \alpha, \beta) = \beta x f(x)(1 - f(x))$$

This quantity is positive for all $x, \beta > 0$, so f is increasing in α . It follows that if $f(\cdot; \alpha, \beta)$ does not have any equilibria in the interval $\left[0, \frac{1}{2}\right]$, then neither does $f(\cdot; \alpha', \beta)$.

Define the *critical payoff gain (for fast learning)* as

$$\alpha^*(\beta) = \begin{cases} \inf A(\beta), & \text{if } A(\beta) \neq \emptyset \\ \infty, & \text{if } A(\beta) = \emptyset \end{cases}$$

The estimation of the critical payoff gain $\alpha^*(\beta)$ consists of two steps. First we show that the critical payoff gain is equal to zero if and only if $\beta \leq 2$ (high error rates). Secondly, we establish an upper bound for the critical payoff gain for $\beta > 2$ (low error rates).

Claim 1. When $\beta \leq 2$ then $\alpha^*(\beta) = 0$. When $\beta > 2$ then $\alpha^*(\beta) > 0$.

To establish this claim, we study $x_L(\alpha, \beta)$, the smallest equilibrium of f . We have already established that f is strictly increasing in α on $(0, 1)$, hence x_L is strictly increasing in α . Thus, to show that $\alpha^*(\beta) = 0$ it is sufficient to show that $x_L(0, \beta) = \frac{1}{2}$. Moreover, if $x_L(0, \beta) < \frac{1}{2}$ and $f'(x_L(0, \beta); 0, \beta) < 1$ it follows that for sufficiently small $\alpha > 0$ it still holds that $x_L(\alpha, \beta) < \frac{1}{2}$. This implies that $\alpha^*(\beta) > 0$.

Note that $f(\cdot; 0, \beta)$ is symmetric in the sense that $f(x; 0, \beta) + f(1 - x; 0, \beta) = 1$, and it is strictly convex on $(0, \frac{1}{2})$ and strictly concave on $(\frac{1}{2}, 1)$. Obviously $x = \frac{1}{2}$ is an equilibrium, so the system either has a single equilibrium or three equilibria, depending on whether $f'(\frac{1}{2}; 0, \beta)$ is below or above 1, respectively. Using (6) we have $f'(\frac{1}{2}; 0, \beta) = \frac{\beta}{2}$ so for $\beta \leq 2$ the system has a single equilibrium, and thus $x_L(0, \beta) = \frac{1}{2}$. For $\beta > 2$ the system has three equilibria, and the smallest corresponds to a down crossing. It follows that $x_L(0, \beta) < \frac{1}{2}$ and $f'(x_L(0, \beta); 0, \beta) < 1$.

Claim 2. Let $\beta > 2$ and

$$\alpha \geq h(\beta) = \frac{e^{\beta-1} + 4 - e}{\beta} - 2$$

Then $f(x; \alpha, \beta) - x > 0$ on the interval $[0, \frac{1}{2}]$. It follows that $\Gamma(\alpha, \beta)$ exhibits fast learning.

To establish this claim it will suffice to show that the minimum of $f(x; \alpha, \beta) - x$ on the interval $[0, \frac{1}{2}]$ is positive. Note first that $f(x; \alpha, \beta) - x$ is positive at the endpoints of this interval. The corresponding first order condition is

$$(10) \quad f'(x) = \beta(\alpha + 2)f(x)(1 - f(x)) = 1$$

If this equation does not have a solution in the interval $(0, \frac{1}{2})$ then we are done. Otherwise, denote $X = f(x_0)$ where $x_0 \in (0, \frac{1}{2})$ is a solution of (10), and denote $C = (\beta(\alpha + 2))^{-1}$. Equation (10) can be rewritten as

$$(11) \quad X - X^2 = C$$

It suffices to establish claim 2 for $\alpha = h(\beta)$. This condition translates to

$$(12) \quad C^{-1} = e^{\beta-1} + 4 - e$$

By assumption $\beta > 2$, hence $C \in (0, \frac{1}{4})$. Solving equation (11) we obtain

$$X = \frac{1 - \sqrt{1 - 4C}}{2} \in (0, \frac{1}{2})$$

We now show that $f(x_0) > x_0$, which is equivalent to $f(f(x_0)) > f(x_0)$ because f is strictly increasing. Using expression (4) the latter inequality becomes

$$\begin{aligned} \frac{1}{1 + e^{\beta(1-(\alpha+2)X)}} > X &\Leftrightarrow \frac{1}{X} > 1 + e^{\beta-\beta(\alpha+2)X} \\ &\Leftrightarrow \frac{1}{X} - 1 > e^{\beta-1}e^{1-\beta(\alpha+2)X} \end{aligned}$$

Using (12) to express $e^{\beta-1}$ in terms of C , and the fact that $\beta(\alpha + 2) = C^{-1}$, the inequality to prove becomes:

$$(13) \quad \frac{1 - X}{X} \geq (C^{-1} - 4 + e)e^{1 - \frac{X}{C}}$$

Denote $M = \frac{X}{1-X} \in (0,1)$, hence also $X^{-1} = \frac{1+M}{M}$. Equation (11) implies that $1 - \frac{X}{c} = -M$, and we can write

$$C^{-1} = \frac{1}{X - X^2} = \frac{X}{(1-X)} X^{-2} = \frac{(1+M)^2}{M}$$

Using these identities, inequality (13) becomes

$$\frac{1}{M} > \left(\frac{(1+M)^2}{M} - 4 + e \right) e^{-M} \Leftrightarrow e^M > M^2 + (e-2)M + 1$$

To establish this inequality, define

$$u(M) = e^M - M^2 - (e-2)M - 1$$

This function, depicted in Figure 4, is first increasing and then decreasing, and it is strictly positive on the interior of the interval $(0,1)$.




Figure 4 – The function $u(M)$

This concludes the proof that if $\beta > 2$ and $\alpha \geq h(\beta)$ then $\Gamma(\alpha, \beta)$ exhibits fast learning.

We now show that when $\beta \geq 3$ and $\alpha > h(\beta)$ the process $\Gamma(\alpha, \beta)$ has fast learning to 99%. We claim that the high equilibrium is increasing in both α and β . Indeed, identity (9) implies that $\frac{\partial f}{\partial \alpha}$ is positive for all $x > 0$. We also have

$$\frac{\partial f}{\partial \beta}(x; \alpha, \beta) = ((\alpha + 2)x - 1)f(x)(1 - f(x))$$

By definition $x_H > \frac{1}{2}$ and thus $\frac{\partial f}{\partial \beta}(x_H; \alpha, \beta) > 0$ as claimed.

It is thus sufficient to show that when $\beta = 3$ and $\alpha = h(3) = \frac{e^2+4-e}{3} - 2 > 89\%$ the high equilibrium is above 99%. An explicit calculation shows that

$$f(0.99; h(3), 3) > \frac{1}{1 + e^{3(1-(.89+2) \cdot (0.99))}} \approx 0.9963 > 0.99.$$

It follows that $x_H > 0.99$.

The final part of the proof is to show that the deterministic process is well approximated by the stochastic process for sufficiently large population size N .

Let α and β be such that the process $\Gamma(\alpha, \beta)$ has fast learning, namely there exists a unique equilibrium x_H , and this equilibrium is strictly above $\frac{1}{2}$. Given a small *precision level* $\delta > 0$, recall that $T(\alpha, \beta, x_H - \delta)$ is the time until the deterministic process comes closer than δ to the equilibrium x_H . Similarly, $T_N(\alpha, \beta, x_H - \delta)$ is the time until the stochastic process with population size N comes closer than δ to the equilibrium x_H .

Lemma 2. If the deterministic process $\Gamma(\alpha, \beta)$ exhibits fast learning, then $\Gamma_N(\alpha, \beta)$ also exhibits fast learning. More precisely, for any $\delta > 0$ we have

$$(14) \quad \lim_{N \rightarrow \infty} ET_N(\alpha, \beta, x_H - \delta) = T(\alpha, \beta, x_H - \delta)$$

Proof. The key result is Lemma 1 in Benaïm and Weibull 2003, which bounds the maximal deviation of the finite process from the deterministic approximation on a bounded time interval, as the population size goes to infinity (see also Kurtz 1970). Before stating the result, we introduce some notation. Denote by $x_N(k\tau)$ the random variable describing the adoption rate in the process $\Gamma_N(\alpha, \beta)$, where $\tau = \frac{1}{N}$ and $k \in \mathbb{N}$. To extend the process x_N to a continuous time process, define the *step process* \bar{x}_N and the *interpolated process* \hat{x}_N as

$$\begin{aligned} \bar{x}_N(t) &= x_N(k\tau), \text{ and} \\ \hat{x}_N(t) &= x_N(k\tau) + \frac{t - k\tau}{\tau} \left(x_N((k+1)\tau) - x_N(k\tau) \right), \end{aligned}$$

for any $t \in [k\tau, (k+1)\tau)$.

Lemma 3. (adapted from Lemma 1 in Benaïm and Weibull 2003) For any $T > 0$ there exists a constant $c = c(T) > 0$ such that for any $\mu > 0$ and N sufficiently large:

$$\Pr(|x(T) - \hat{x}_N(T)| \geq \mu) \leq 2e^{-\mu^2 cN}$$

The proof is relegated to the appendix.

For convenience, we omit the dependence of T and T_N on α and β . Assuming that $T(x_H - \delta) < \infty$, it is now easy to prove equality (14). Consider a small $\epsilon > 0$, take $\mu = \frac{\epsilon}{2}$ and denote $T_\epsilon = T(x_H - \delta + \epsilon)$. Lemma 3 implies

$$\Pr(x_N(T_\epsilon) > x_H - \delta + \epsilon - \mu) \geq 1 - 2e^{-\frac{\epsilon^2 cN}{4}}$$

It follows that

$$\Pr\left(T_N\left(x_H - \delta + \frac{\epsilon}{2}\right) \leq T_\epsilon\right) \geq 1 - 2e^{-\epsilon^2 cN/4}$$

We claim that $\Pr(T_N < T) \geq q$ implies that $E(T_N) < \frac{T}{q}$. The proof relies on the following simple argument. With probability q we have $T_N < T$. With the remaining probability $1 - q$, for $t \geq T$ we know that $x(t)$ is lower bounded by a process y satisfying the same differential equation $\dot{y} = P(y) - y$ and with a different starting point, namely $y(T) = 0$. This implies that with probability at least $(1 - q)q$ we have $T_N < 2T$. By iteration we obtain

$$E(T_N) \leq (q + 2(1 - q)q + 3(1 - q)^2q + \dots)T = q \frac{1}{q^2} T = \frac{T}{q}$$

It follows that

$$ET_N(x_H - \delta) < ET_N(x_H - \delta + \epsilon/2) \leq \frac{T_\epsilon}{1 - 2e^{-\epsilon^2 cN/4}}$$

Taking limits in N on both sides we get that $\lim_{N \rightarrow \infty} ET_N(x_H - \delta) \leq T_\epsilon$ for any $\epsilon > 0$, hence $\lim_{N \rightarrow \infty} ET_N(x_H - \delta) \leq T$ by taking $\epsilon \rightarrow 0$. A similar argument show that $\lim_{N \rightarrow \infty} ET_N(x_H - \delta) \geq T$. ■

This concludes the proof of theorem 1. ■

5. Theoretical Bounds on Waiting times

We now embark on a systematic analysis of the magnitude of the time it takes to get close to the high equilibrium, starting from the all- B state. Figure 5 shows simulations of the waiting times until the adoption rate is within $\delta = 1\%$ of the high equilibrium. The blue, dashed line represents the critical payoff gain $\alpha^*(\beta)$, such that for $\alpha > \alpha^*(\beta)$ learning is fast. Pairs (ϵ, α) on the red line correspond to waiting times $T(\alpha, \beta, x_H - \delta)$ of 40 revisions per capita, while on the orange and green lines the corresponding times are 20 and 10 revisions per capita, respectively.




Figure 5 – Waiting times to within $\delta = 1\%$ of the high equilibrium; 40 (red), 20 (orange) and 10 (green) revisions per capita.

The following result provides a characterization of the expected waiting time as a function of the payoff gain.

Theorem 2. For any noise level $\beta > 2$ and precision level $\delta > 0$ there exists a constant $S = S(\beta, \delta)$ such that for every $\alpha > \alpha^*(\beta)$ and all sufficiently large N given α, β and δ , the expected waiting time $T_N = T_N(\alpha, \beta, x_H - \delta)$ until the adoption rate is within δ of the high equilibrium satisfies

$$(15) \quad E(T_N) \leq \frac{S}{\sqrt{\frac{\alpha + 2}{\alpha^*(\beta) + 2} - 1}} + \log(\delta^{-1}) + \delta$$

To understand Theorem 2, note that as the payoff gain α approaches the threshold $\alpha^*(\beta)$, a “bottleneck” appears for intermediate adoption rates, which slows down the process. Figure 6 illustrates this phenomenon by highlighting the distance between the updating function f and the identity function (recall that the speed of the process at adoption rate x is given by $f(x) - x$). The first term on the right hand side of inequality (15) tends to infinity as α tends to $\alpha^*(\beta)$, and the proof of Theorem 2 shows that inequality (14) holds for the following explicit value of the constant S :

$$(16) \quad S_0 = \frac{4\sqrt{2}}{\sqrt{1 - \frac{4}{\beta(\alpha^*(\beta) + 2)}}$$

When the payoff gain is large, the main constraining factor is the precision level δ , namely how close we want the process to approach the high equilibrium. The last two terms on the right hand side of inequality (15) take care of this possibility.




Figure 6 – An informal illustration of the evolution of the process

Figure 7 plots, in log-log format, the expected time to get within $\delta = 1\%$ of the high equilibrium as a function of $\alpha - \alpha^*$, for $\varepsilon = 1\%$ (left panel) and $\varepsilon = 5\%$ (right panel). These error rate correspond to

$\beta \approx 4.59$ and $\beta \approx 2.95$ respectively. The constant S_0 takes the values $S_0(1\%) \approx 6$ and $S_0(5\%) \approx 8$. The red, solid line represents the estimated upper-bound (using constant S_0), while the blue, dashed line shows the simulated deterministic time $T(\alpha, \beta, x_H - \delta)$.

Finally, we give a concrete example of the estimated waiting time. As above, fix the error rate at $\varepsilon = 5\%$ and the precision level at $\delta = 1\%$. Then the critical payoff gain is $\alpha^*(\beta) \approx 74\%$. Our estimate for the waiting time when $\alpha = 100\%$ is $T = 30.5$, while the simulated average waiting time is 16.95 (for $N = 1000$).




Figure 7 – Waiting time until the process is within $\delta = 1\%$ of the high equilibrium, as a function of the payoff gain differential $\alpha - \alpha^*$. Both axes have logarithmic scales. Simulation (blue, dashed) and upper bound (red, solid). Error rates $\varepsilon = 1\%$ (left panel) and $\varepsilon = 5\%$ (right panel).

Proof of Theorem 2. We prove the statement for the deterministic approximation $\Gamma(\alpha, \beta)$ and the associated waiting time $T = T(\alpha, \beta, x_H - \delta)$. It then follows, using Lemma 2, that Theorem 2 holds for all sufficiently large N given α and β .

The proof unfolds as follows. The first step is to show that as the payoff gain α approaches the critical payoff gain $\alpha^* = \alpha^*(\beta)$ from above, the waiting time is controlled by the bottleneck that appears around the low equilibrium $x_L^* = x_L(\alpha^*)$. More specifically, the waiting time scales with the inverse square root of the height of the bottleneck, called the *gap*. In the second step we show that for α close to α^* the gap is approximately linear in the payoff gain difference $(\alpha^* - \alpha)$, and we derive the constant S_0 . The third step consists of showing that the limit waiting time as the payoff gain α tends to infinity is smaller than the second term on the right hand-side of (15). The final step wraps up the proof by showing that the case of intermediate α is covered by an appropriate choice of the constant S .

We first consider the case when the payoff gain α is close to $\alpha^*(\beta)$. We begin by showing that the low equilibrium disappears for $\alpha > \alpha^*(\beta)$. Fix $\beta > 2$ and omit the explicit dependence of f and α^* on β . By claim 1, the condition $\beta > 2$ implies that $\alpha^* > 0$. The low equilibrium satisfies $x_L^* \equiv x_L(\alpha^*) < \frac{1}{2}$, because $f\left(\frac{1}{2}\right) > \frac{1}{2}$. By definition, $x_L(\alpha) > \frac{1}{2}$ for all $\alpha > \alpha^*$, hence

$$x_L^* < \frac{1}{2} \leq \lim_{\alpha \searrow \alpha^*} x_L(\alpha)$$

Lemma 1 now implies that at $\alpha = \alpha^*$ there are two equilibria, namely x_L^* and x_H^* , while for $\alpha > \alpha^*$ there exists a single equilibrium x_H , and it satisfies $x_H > \frac{1}{2}$. As the payoff gain α approaches the threshold α^* from above, a bottleneck appears around x_L^* , which slows down the evolution of the process. Concretely, the minimal speed of the process is controlled by the minimal value of $f(x; \alpha) - x$ on $[0, x_H - \delta]$, and this quantity approaches zero as α approaches α^* . We calculate the rate at which the time to pass through the bottleneck grows to infinity as $\alpha \searrow \alpha^*$. This rate depends on the width of the bottleneck, namely the minimal value of $f(x; \alpha) - x$, as well as on its length, which is controlled by the curvature $\frac{1}{f''(x_L^*; \alpha^*)}$.

We now define the width of the bottleneck formally by studying the minimum value of $f(x; \alpha) - x$ on $[0, x_H - \delta]$. Note that $f(x; \alpha^*) \geq x$ on $[0, x_H - \delta]$, with equality if and only if $x = x_L^*$. This implies that $f'(x_L^*; \alpha^*) = 1$. By (8) we know that f is first convex and then concave, hence $f''(x_L^*; \alpha^*) > 0$. The implicit function theorem then implies that there exists a function $t(\alpha)$, defined in a neighborhood (α_0, α_1) containing α^* , that satisfies $f'(t(\alpha); \alpha) = 1$. By an appropriate choice of the neighborhood (α_0, α_1) we can ensure that $t(\alpha)$ is also the unique solution to the problem $\min_{x \in [0, x_H - \delta]} f(x; \alpha) - x$. For any $\alpha \in [\alpha^*, \alpha_1)$ define the *gap* as

$$(17) \quad g(\alpha) = f(t(\alpha); \alpha) - t(\alpha)$$

The gap is the smallest distance between the function f and the identity function, on an interval bounded away from the high equilibrium. By definition we have that

$$f(x; \alpha) - x \geq g(\alpha) \text{ for all } x \in [0, x_H - \delta].$$

Note that the gap is non-negative, because f is increasing in α . In particular, we have $t(\alpha^*) = x_L^*$ and $g(\alpha^*) = 0$.

To estimate the hitting time $T(\alpha, \beta, x_H - \delta)$ we shall construct a lower bound for $f(x; \alpha)$. The idea is that in a neighborhood around x_L^* we bound $f(x; \alpha)$ by $x + g(\alpha)$, while everywhere else on $[0, \frac{1}{2}]$ we bound it by $f(x; \alpha^*)$. Figure 8 illustrates the approximation.




Figure 8 – The function $h(x, \alpha)$ (red) is a lower bound for $f(x; \alpha, \beta)$ (green)

To formalize this construction, consider the second order Taylor's expansion of $f(x; \alpha^*) - x$ around $x = x_L^*$

$$(18) \quad \begin{aligned} f(x; \alpha^*) - x &= (f(x_L^*; \alpha^*) - x_L^*) + (f'(x_L^*; \alpha^*) - 1)(x - x_L^*) + \\ &\quad + \frac{f''(x_L^*; \alpha^*)}{2}(x - x_L^*)^2 + \mathcal{O}((x - x_L^*)^3) \\ &= \frac{f''(x_L^*; \alpha^*)}{2}(x - x_L^*)^2 + \mathcal{O}((x - x_L^*)^3) \end{aligned}$$

The later equality holds because $f(x_L^*; \alpha^*) = x_L^*$ and $f'(x_L^*; \alpha^*) = 1$.

Fix some $\theta < 1$ that is close to 1. There exist bounds $\underline{x} = \underline{x}(\theta) < x_L^*$ and $\bar{x} = \bar{x}(\theta) > x_L^*$ such that for all $x \in [\underline{x}, \bar{x}]$ we have

$$(19) \quad \theta \frac{f''(x_L^*; \alpha^*)}{2}(x - x_L^*)^2 < f(x; \alpha^*) - x < \frac{1}{\theta} \frac{f''(x_L^*; \alpha^*)}{2}(x - x_L^*)^2$$

We now consider the equation (in x) $f(x; \alpha^*) - x = g(\alpha)$. Given θ there exists $\alpha_2 \in (\alpha^*, \alpha_1]$ such that for any $\alpha \in (\alpha^*, \alpha_2)$ the equation has exactly two solutions $x_1 = x_1(\alpha)$ and $x_2 = x_2(\alpha)$ that satisfy $\underline{x} < x_1(\alpha) < x_L^*$ and $x_L^* < x_2(\alpha) < \bar{x}$. Using (19) we obtain

$$(20) \quad \sqrt{\theta \frac{2g(\alpha)}{f''(x_L^*; \alpha^*)}} < |x_L^* - x_i(\alpha)| < \sqrt{\frac{1}{\theta} \frac{2g(\alpha)}{f''(x_L^*; \alpha^*)}} \text{ for } i = 1, 2$$

We now decompose the waiting time

$$T(\alpha, \beta, x_H - \delta) = T\left(0 \rightarrow \frac{1}{2}\right) + T\left(\frac{1}{2} \rightarrow x_H - \delta\right),$$

where $T(x_0 \rightarrow x_1)$ denotes the waiting time until the process $\Gamma(\alpha, \beta)$ with initial condition $x(0) = x_0$ reaches x_1 .

We claim that the second term is lower-bounded by some constant K_1 independently of α . Indeed, this waiting time is continuous in α , it is finite for $\alpha = \alpha^*$ and it is bounded as $\alpha \rightarrow \infty$. The last claim follows from the fact that $f'(x_H; \alpha) - 1$ is negative and bounded away from zero independently of $\alpha \geq \alpha^*$.

We now turn to an estimate of the term $T\left(0 \rightarrow \frac{1}{2}\right)$. The idea is to consider a process that is slower than $\Gamma(\alpha, \beta)$ and that is analytically easier to work with. For every $\alpha \in [\alpha^*, \alpha_2)$ define $v(x; \alpha)$ as follows

$$v(x; \alpha) = \begin{cases} f(x; \alpha^*), & x \in [0, x_1(\alpha)) \\ x + g(\alpha), & x \in [x_1(\alpha), x_2(\alpha)] \\ f(x; \alpha^*), & x \in (x_2(\alpha), 1]. \end{cases}$$

Obviously, this satisfies $v(x; \alpha) \leq f(x; \alpha)$ for all $x \in [0, \frac{1}{2}]$ and $\alpha \in [\alpha^*, \alpha_2)$. Let $V(\alpha, \beta)$ be the process with state variable $y(t)$ and dynamics given by $\dot{y} = v(y; \alpha) - y$. Let $T_V(x_0 \rightarrow x_1)$ be the time until the process $V(\alpha, \beta)$ with initial condition $y(0) = x_0$ reaches x_1 .

Clearly $T_V(x_0 \rightarrow x_1) \geq T(x_0 \rightarrow x_1)$ for all $x_0, x_1 \in [0, \frac{1}{2}]$. We further decompose $T_V(0 \rightarrow \frac{1}{2})$ as

$$(21) \quad T_V\left(0 \rightarrow \frac{1}{2}\right) = T_V(0 \rightarrow \underline{x}) + T_V(\underline{x} \rightarrow x_1(\alpha)) + \\ + T_V(x_1(\alpha) \rightarrow x_2(\alpha)) + T_V(x_2(\alpha) \rightarrow \bar{x}) + T_V(\bar{x} \rightarrow \frac{1}{2})$$

The first and last terms are lower bounded by constants K_2 and K_3 independently of α . (The reason is that v does not depend on α in the intervals $[0, \underline{x}]$ and $[\bar{x}, \frac{1}{2}]$). Using (20), the middle term satisfies

$$(22) \quad T_V(x_1(\alpha) \rightarrow x_2(\alpha)) = \frac{x_2(\alpha) - x_1(\alpha)}{g(\alpha)} < 2 \sqrt{\frac{1 - 2g(\alpha)}{\theta f''(x_L^*; \alpha^*)} \frac{1}{g(\alpha)}} = \frac{1}{\sqrt{\theta}} \frac{2\sqrt{2}}{\sqrt{f''(x_L^*; \alpha^*)g(\alpha)}}$$

We now look at $T_V(\underline{x} \rightarrow x_1(\alpha))$. The idea is to use approximation (19) to lower bound the process $V(\alpha, \beta)$ with initial condition $y(0) = \underline{x}$. Consider the process $z(t)$ with dynamics given by

$$(23) \quad \dot{z} = \theta \frac{f''(x_L^*; \alpha^*)}{2} (x_L^* - z)^2 \quad \text{and initial condition } z(0) = \underline{x}$$

By (19) we know that $z(t) \leq y(t)$ for all $t \geq 0$. The solution to equation (23) is

$$(24) \quad z(t) = x_L^* - \frac{1}{\theta} \frac{2}{(t + t_0)f''(x_L^*; \alpha^*)}$$

The term t_0 solves the equation $z(0) = \underline{x}$, and is thus positive.

Solving for T in $z(T) = x_1(\alpha)$ we obtain

$$(25) \quad T_V(\underline{x} \rightarrow x_1(\alpha)) < T = \frac{1}{\theta} \frac{2}{f''(x_L^*; \alpha^*)(x_L^* - x_1(\alpha))} - t_0 < \frac{1}{\sqrt{\theta^3}} \frac{\sqrt{2}}{\sqrt{f''(x_L^*; \alpha^*)g(\alpha)}}$$

By a similar argument we obtain $T_V(x_2(\alpha) \rightarrow \bar{x}) < \frac{1}{\sqrt{\theta^3}} \frac{\sqrt{2}}{\sqrt{f''(x_L^*; \alpha^*)g(\alpha)}}$.

Putting results (21), (22) and (25) together we obtain that for $\alpha \in [\alpha^*, \alpha_2)$ we have

$$(26) \quad T(\alpha, \beta, x_H - \delta) \leq K_4 + \frac{1}{\sqrt{\theta^3}} \frac{4\sqrt{2}}{\sqrt{f''(x_L^*; \alpha^*)g(\alpha)}}$$

where $K_4 = K_1 + K_2 + K_3$ is independent of α . Taking the limit as α tends to α^* we obtain

$$\limsup_{\alpha \searrow \alpha^*} T(\alpha, \beta, x_H - \delta) \sqrt{f''(x_L^*; \alpha^*)g(\alpha)} \leq \frac{4\sqrt{2}}{\sqrt{\theta^3}}$$

This inequality holds for all $\theta < 1$, hence

$$(27) \quad \limsup_{\alpha \searrow \alpha^*} T(\alpha, \beta, x_H - \delta) \sqrt{f''(x_L^*; \alpha^*)g(\alpha)} \leq 4\sqrt{2}.$$

Continuing to the second step of the proof, we now lower bound the quantity $f''(x_L^*; \alpha^*)g(\alpha)$ for small $\alpha > \alpha^*$. Firstly, using (7) we obtain

$$(28) \quad \begin{aligned} f''(x_L^*; \alpha^*) &= \beta(\alpha^* + 2)f'(x_L^*; \alpha^*)(1 - 2f(x_L^*; \alpha^*)) \\ &= \beta(\alpha^* + 2)(1 - 2x_L^*). \end{aligned}$$

Consider the first order Taylor's expansion of $g(\alpha)$ around $\alpha = \alpha^*$

$$\begin{aligned} g(\alpha) &= g(\alpha^*) + g'(\alpha^*)(\alpha - \alpha^*) + \mathcal{O}((\alpha - \alpha^*)^2) \\ &= g'(\alpha^*)(\alpha - \alpha^*) + \mathcal{O}((\alpha - \alpha^*)^2). \end{aligned}$$

The last equality holds because $g(\alpha^*) = 0$. Fix some $\theta < 1$ that is close to 1. Then there exists $\alpha_3 \in (\alpha^*, \alpha_1]$ such that for any $\alpha \in (\alpha^*, \alpha_3)$ we have

$$(29) \quad g(\alpha) > \theta g'(\alpha^*)(\alpha - \alpha^*).$$

From the definition of the gap we have that for all $\alpha \in [\alpha^*, \alpha_1)$

$$(30) \quad g'(\alpha) = f'(t(\alpha); \alpha)t'(\alpha) + \frac{df}{d\alpha}(t(\alpha); \alpha) - t'(\alpha) = \frac{df}{d\alpha}(t(\alpha); \alpha).$$

Recall that in general

$$(31) \quad \frac{df}{d\alpha}(x; \alpha) = \beta x f(x; \alpha)(1 - f(x; \alpha)) = \frac{x}{\alpha + 2} f'(x; \alpha)$$

Putting (29), (30) and (31) together, and using $f'(x_L^*; \alpha^*) = 1$, we obtain that for any $\alpha \in (\alpha^*, \alpha_3)$

$$g(\alpha) > \theta \frac{(\alpha - \alpha^*)x_L^*}{\alpha^* + 2}$$

Combining this inequality with identity (28) we obtain that for any $\alpha \in (\alpha^*, \alpha_3)$

$$(32) \quad f''(x_L^*; \alpha^*)g(\alpha) > \theta(\alpha^* + 2)\beta(x_L^* - 2(x_L^*)^2) \frac{\alpha - \alpha^*}{\alpha^* + 2}.$$

The equations $f(x_L^*; \alpha^*) = x_L^*$ and $f'(x_L^*; \alpha^*) = 1$ imply that (see equation (11))

$$x_L^* - (x_L^*)^2 = C \equiv \frac{1}{(\alpha^* + 2)\beta}.$$

Note that the solution verifies $x_L^* < C + 4C^2$.⁵ We can thus write

$$(\alpha^* + 2)\beta(x_L^* - 2(x_L^*)^2) = \frac{x_L^* - 2(x_L^* - C)}{C} = \frac{2C - x_L^*}{C} > \frac{2C - C - 4C^2}{C} = 1 - 4C$$

⁵ The exact solution is $x_L^* = \frac{1 - \sqrt{1 - 4C}}{2}$ and the desired inequality follows by direct computation.

Together with inequality (32), we have that for any $\alpha \in (\alpha^*, \alpha_3)$

$$f''(x_L^*; \alpha^*)g(\alpha) > \theta \left(1 - \frac{4}{\beta(\alpha^* + 2)}\right) \left(\frac{\alpha + 2}{\alpha^* + 2} - 1\right)$$

Combining this inequality with inequality (27) we obtain that

$$\limsup_{\alpha \searrow \alpha^*} T(\alpha, \beta, x_H - \delta) \sqrt{\left(1 - \frac{4}{\beta(\alpha^* + 2)}\right) \left(\frac{\alpha + 2}{\alpha^* + 2} - 1\right)} \leq \frac{4\sqrt{2}}{\sqrt{\theta}}$$

Using that the inequality holds for all $\theta < 1$ and rearranging terms, we obtain

$$(33) \quad \limsup_{\alpha \searrow \alpha^*} T(\alpha, \beta, x_H - \delta) \sqrt{\frac{\alpha + 2}{\alpha^* + 2} - 1} \leq \frac{4\sqrt{2}}{\sqrt{1 - \frac{4}{\beta(\alpha^* + 2)}}}$$

Note that the right hand side of (33) is exactly S_0 as defined in (16).

Moving on to the third step of the proof, we now study the waiting time when the payoff gain α is large. Denote by x_α the solution to the ordinary differential equation

$$\dot{x} = f(x; \alpha, \beta) - x \quad \text{with initial condition } x(0) = 0.$$

Denote by x_∞ the solution to the ordinary differential equation

$$\dot{x} = 1 - x \quad \text{with initial condition } x(0) = 0.$$

We show that x_α converges to x_∞ pointwise, in the sense that for any $t \geq 0$ we have $\lim_{\alpha \rightarrow \infty} x_\alpha(t) = x_\infty(t)$. This follows from the fact that f converges to 1 as α tends to infinity, in the sense that $\lim_{\alpha \rightarrow \infty} f(x; \alpha, \beta) = 1$ for all $x \in (0, 1]$. A more detailed argument is necessary, however, as this convergence is not uniform in x , because $f(0; \alpha, \beta) = \varepsilon$ for all α .

Formally, for any small enough $\mu > 0$ there exists $\alpha(\mu)$ such that for all $\alpha > \alpha(\mu)$ and all $x \geq \mu$ we have $f(x; \alpha, \beta) > 1 - \mu$. We also have $f(x; \alpha, \beta) \geq \varepsilon$ for all x , hence x_α is lower bounded by the solution to the equation

$$\dot{y} = \begin{cases} \varepsilon - y, & x < \mu \\ 1 - \mu - y, & x \geq \mu, \end{cases} \quad \text{with initial condition } y(0) = 0.$$

A simple yet somewhat involved calculation shows that $\lim_{\mu \rightarrow 0} y(t) = x_\infty(t)$.⁶

The solution x_∞ is given by $x_\infty(t) = 1 - e^{-t}$. The waiting time satisfying $x_\infty(T) = 1 - \delta$ is $T = \log(\delta^{-1})$, which implies that

⁶ Explicitly, we have

$$y(t) = \begin{cases} \varepsilon - e^{\log(\varepsilon)-t}, & \text{for } t \leq \log\left(\frac{\varepsilon}{\varepsilon - \mu}\right) \\ 1 - \mu - e^{\log\left(\frac{\varepsilon(1-2\mu)}{\varepsilon - \mu}\right)-t}, & \text{for } t \geq \log\left(\frac{\varepsilon}{\varepsilon - \mu}\right). \end{cases}$$

$$(34) \quad \lim_{\alpha \rightarrow \infty} T(\alpha, \beta, x_H - \delta) = \log(\delta^{-1})$$

In the last step of the proof, we put the previous steps together to prove Theorem 2. Note first that $T(\alpha) = T(\alpha, \beta, x_H - \delta)$ is continuous in $\alpha \in (\alpha^*, \infty)$.

Equation (34) implies that there exists an α_0 such that for all $\alpha > \alpha_0$ we have

$$(35) \quad T(\alpha) < \log(\delta^{-1}) + \frac{\delta}{2}$$

On the other hand, equation (33) implies that the quantity $T(\alpha) \sqrt{\frac{\alpha+2}{\alpha^*+2}} - 1$ is upper bounded on the interval $(\alpha^*, \bar{\alpha}]$. Let the constant S be such that for all $\alpha \in (\alpha^*, \bar{\alpha}]$

$$(36) \quad T(\alpha) \leq \frac{S}{\sqrt{\frac{\alpha+2}{\alpha^*+2}} - 1}$$

Putting inequalities (35) and (36) together, we find that the deterministic process $\Gamma(\alpha, \beta)$ verifies, for any $\alpha > \alpha^*$,

$$T(\alpha) \leq \frac{S}{\sqrt{\frac{\alpha+2}{\alpha^*+2}} - 1} + \log(\delta^{-1}) + \frac{\delta}{2}.$$

The exact result in Theorem 2 follows by applying Lemma 2. ■

6. Partial information

Now consider the case where agents have a limited, finite capacity to gather information, namely each player samples exactly d other players before revising with $d \geq 3$ independent of N . It turns out that partial information facilitates the spread of the innovation, in the sense that the critical payoff gain is lower than in the full information case. Intuitively, for low adoption rates the effect of sample variability is asymmetric: the increased probability of adoption when adopters are over-represented in the sample outweighs the decreased probability of adoption when adopters are under-represented. In particular, the coarseness of a finite sample implies that the threshold α^* is no longer unbounded as the noise tends to zero. Indeed, for $\frac{1}{\alpha+2} < \frac{1}{d}$, or equivalently $\alpha > d - 2$, the existence of a single adopter in the sample makes it a best response to adopt the innovation. This implies that the process displays fast learning for *any* noise level. This argument is formalized in Theorem 3 below. (Here and for the remainder of the section, we modify the previous notation by adding d as a parameter. For example, the process with payoff gain α , noise parameter β , population size N and sampling size d is denoted $\Gamma_N(\alpha, \beta, d)$; the waiting time to adoption level p is denoted $T_N(\alpha, \beta, p, d)$ and so forth.)

Theorem 3. Consider $3 \leq d < \infty$. If $\alpha > h_d(\beta)$ then $\Gamma_N(\alpha, \beta, d)$ has fast learning, where

$$h_d(\beta) = \min(h(\beta), d - 2).$$

Proof of Theorem 3. The proof follows the same logic as the proof of Theorem 1. Specifically, we shall show that the result holds for the deterministic approximation of the finite process, which implies that it also holds for sufficiently large population size. We begin by characterizing the response function in the partial information case. The next step is to show that, as in the case of full information, the high equilibrium of the deterministic process is unique. We then show that the threshold for partial information is below the threshold for full information, and also that $d - 2$ forms an upper bound on the threshold for partial information.

When agents have access to partial information only, the response function denoted $f_d(x; \alpha, \beta)$ depends on the population adoption rate x as well as on the sample size d . For notational convenience, fix the dependence of f and f_d on α and β , and write $f(x)$ and $f_d(x)$ instead of $f(x; \alpha, \beta)$ and $f_d(x; \alpha, \beta)$. The probability that exactly k players in a randomly selected sample of size d are playing A is $\binom{d}{k} x^k (1-x)^{d-k}$, for any $k = 0, 1, \dots, d$. In this case the agent chooses action A with probability $f\left(\frac{k}{d}\right)$. Hence the agent chooses action A with probability

$$(37) \quad f_d(x) = \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} f\left(\frac{k}{d}\right).$$

To calculate the derivative of f_d , let $\Delta_d f(k)$ be the discrete derivative of f with step $1/d$, evaluated at $x = \frac{k}{d}$, namely

$$\Delta_d f(k) = \frac{f\left(\frac{k+1}{d}\right) - f\left(\frac{k}{d}\right)}{1/d}.$$

Differentiating (37) and using the above notation we obtain

$$(38) \quad f'_d(x) = \sum_{k=0}^{d-1} \binom{d-1}{k} x^k (1-x)^{d-1-k} \Delta_d f(k).$$

The definition of the deterministic approximation of the stochastic process $\Gamma_N(\alpha, \beta, d)$ is analogous to the corresponding definition in the perfect information case. The continuous-time process $\Gamma(\alpha, \beta, d)$ has state variable $x(t)$ that evolves according to the ordinary differential equation

$$\dot{x} = f_d(x; \alpha, \beta) - x, \quad \text{with } x(0) = 0.$$

The stochastic process $\Gamma_N(\alpha, \beta, d)$ is well-approximated by the process $\Gamma(\alpha, \beta, d)$ for large population size N . Indeed, the statement and proof of Lemma 2 apply without change to the partial information case.

We now proceed to proving the result in Theorem 3 for the deterministic process $\Gamma(\alpha, \beta, d)$. The following lemmas will be useful throughout the proof.

Lemma 4. For any $\alpha \geq 0$ and any $x \in [0, 1]$ we have $f(x) + f(1-x) \geq 1$. Equality occurs if and only if $\alpha = 0$.

Proof. By explicit calculation using (4) the claim is equivalent to

$$(39) \quad 1 \geq e^{\beta(1-(2+\alpha)x)} e^{\beta(1-(2+\alpha)(1-x))} = e^{-\beta\alpha}.$$

Note that by assumption $\beta > 0$, so inequality (39) is true for all $\alpha \geq 0$, and equality occurs if and only if $\alpha = 0$.

Lemma 5. For any $\alpha > 0$ we have $f_d\left(\frac{1}{2}\right) > \frac{1}{2}$.

Proof. Using Lemma 4, we have that for any k

$$(40) \quad \begin{aligned} \binom{d}{k} f\left(\frac{k}{d}\right) + \binom{d}{d-k} f\left(1 - \frac{k}{d}\right) &= \binom{d}{k} \left(f\left(\frac{k}{d}\right) + f\left(1 - \frac{k}{d}\right) \right) \\ &> \binom{d}{k} \\ &= \binom{d}{k} \frac{1}{2} + \binom{d}{d-k} \frac{1}{2}. \end{aligned}$$

Adding up inequality (40) for $k = 1, \dots, \left\lfloor \frac{d}{2} \right\rfloor$ we obtain⁷

$$f_d\left(\frac{1}{2}\right) = \frac{1}{2^d} \sum_{k=0}^d \binom{d}{k} f\left(\frac{k}{d}\right) > \frac{1}{2^d} \sum_{k=0}^d \binom{d}{k} \frac{1}{2} = \frac{1}{2}. \quad \blacksquare$$

Lemma 6. The function f_d is first strictly convex then strictly concave, and the inflection point is at most $\frac{1}{2}$.

The straightforward but somewhat involved proof is deferred to the appendix.

The selection property established for full information continues to hold with partial information.

Lemma 7. For any $\alpha > 0$ there exists a unique equilibrium $x_H > \frac{1}{2}$ (the *high equilibrium*), and it is stable. Furthermore, there exist at most two equilibria in the interval $\left[0, \frac{1}{2}\right]$.

Proof. Lemma 5 showed that $f_d\left(\frac{1}{2}\right) > \frac{1}{2}$. Since f is continuous and always less than 1, there exists a down crossing point in the interval $\left(\frac{1}{2}, 1\right)$. This corresponds a high equilibrium. Furthermore, the strict concavity of f_d on $\left(\frac{1}{2}, 1\right)$, established in Lemma 6, guarantees that the high equilibrium is unique.

By Lemmas 6 and 5 f_d is convex and then concave on $\left[0, \frac{1}{2}\right]$, and $f_d\left(\frac{1}{2}\right) > \frac{1}{2}$. This implies that there exist at most two equilibria in this interval. To see this, let x_I be f_d 's inflection point, and consider two cases. If $f_d(x_I) \geq \frac{1}{2}$ there are at most two equilibria in the interval $[0, x_I]$, because f_d is strictly

⁷ For even values of d we multiply the inequality corresponding to $k = \frac{d}{2}$ by one-half to avoid double counting.

convex on this interval. Moreover, there are no equilibria in $(x_I, \frac{1}{2}]$, because f_d is strictly concave on this interval. If $f_d(x_I) < \frac{1}{2}$ there exists exactly one equilibrium in the interval $[0, x_I]$, and also exactly one equilibrium in the interval $(x_I, \frac{1}{2}]$. ■

The definition of the critical payoff gain $\alpha^*(\beta, d)$ is analogous to the definition of $\alpha^*(\beta)$ in the proof of Theorem 1. Define the set

$$A(\beta, d) = \{\alpha : \Gamma(\alpha, \beta, d) \text{ displays fast learning}\}.$$

Let

$$\alpha^*(\beta, d) = \begin{cases} \inf A(\beta, d), & \text{if } A(\beta, d) \neq \emptyset \\ \infty, & \text{if } A(\beta, d) = \emptyset. \end{cases}$$

We claim that if $\alpha \in A(\beta, d)$ and $\alpha < \alpha'$ then also $\alpha' \in A(\beta, d)$. First, differentiating expression (37) with respect to α we obtain

$$(41) \quad \frac{\partial f_d}{\partial \alpha}(x; \alpha, \beta) = \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{\partial f}{\partial \alpha}\left(\frac{k}{d}; \alpha, \beta\right)$$

Using the expression for $\frac{\partial f}{\partial \alpha}$ available in (9) we find that $\frac{\partial f_d}{\partial \alpha}(x; \alpha, \beta)$ is positive for all $x, \beta > 0$, so f_d is increasing in α . It follows that if $f_d(\cdot; \alpha, \beta)$ does not have any equilibria in the interval $[0, \frac{1}{2}]$, then neither does $f_d(\cdot; \alpha', \beta)$.

The estimation of $\alpha^*(\beta, d)$ consists of two steps. First, we show that $h(\beta)$ is an upper bound on $\alpha^*(\beta, d)$. Secondly, we show that $d - 2$ is also an upper bound on $\alpha^*(\beta, d)$. Theorem 3 then follows directly.

To prove that $\alpha^*(\beta, d) \leq \alpha^*(\beta)$, we show that if for given $\alpha > 0$ and β the process $\Gamma(\alpha, \beta)$ exhibits fast learning, then so does $\Gamma(\alpha, \beta, d)$, for any $d \geq 3$.

In other words, assume that the function f does not have any equilibrium in the interval $[0, \frac{1}{2}]$. We will show that neither does the function f_d .

Let m be the largest integer such that $f(1 - \frac{m}{d}) \leq 1 - \frac{m}{d}$. By assumption that $\Gamma(\alpha, \beta)$ exhibits fast learning $m \leq \frac{d}{2}$. By definition of m we have that

$$(42) \quad f\left(\frac{k}{d}\right) > \frac{k}{d} \quad \text{for any } k \text{ such that } m < k < d - m.$$

We also use the identity

$$(43) \quad \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{k}{d} = x.$$

Fix x in the interval $\left[0, \frac{1}{2}\right]$, and rewrite expression (37) as

$$f_d(x) = \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{k}{d} + \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \left(f\left(\frac{k}{d}\right) - \frac{k}{d} \right)$$

Using identity (43) and inequality (42), and then rearranging terms, we obtain

$$(44) \quad \begin{aligned} f_d(x) &\geq x + \sum_{\substack{0 \leq k \leq m, \\ d-m \leq k \leq d}} \binom{d}{k} x^k (1-x)^{d-k} \left(f\left(\frac{k}{d}\right) - \frac{k}{d} \right) \\ &= x + \sum_{k=0}^m \binom{d}{k} x^{d-k} (1-x)^k \left[\left(\frac{1-x}{x}\right)^{d-2k} \left(f\left(\frac{k}{d}\right) - \frac{k}{d} \right) + \left(f\left(\frac{d-k}{d}\right) - \frac{d-k}{d} \right) \right] \end{aligned}$$

Fix $k \in \{0, 1, \dots, m\}$. Then $\left(\frac{1-x}{x}\right)^{d-2k} \geq 1$, and $f\left(\frac{k}{d}\right) - \frac{k}{d}$ is positive, so the term in square brackets is at least

$$f\left(\frac{k}{d}\right) - \frac{k}{d} + f\left(\frac{d-k}{d}\right) - \frac{d-k}{d} = f\left(\frac{k}{d}\right) + f\left(1 - \frac{k}{d}\right) - 1.$$

The last term is strictly positive by Lemma 4.

Using (44) we have now established that $f_d(x) > x$ for all $x \leq \frac{1}{2}$. It follows that $\alpha^*(\beta, d) \leq \alpha^*(\beta)$. In particular, $\alpha^*(\beta, d) \leq h(\beta)$.

The second part of the estimation of the critical payoff gain is to show that $\alpha^*(\beta, d) < d - 2$. We show that when $\alpha \geq d - 2$ then $f_d(x) > x$ for all $x \leq \frac{1}{2}$.

Note that $f\left(\frac{1}{d}\right) \geq f\left(\frac{1}{\alpha+2}\right) = \frac{1}{2}$. Hence for all $1 \leq k \leq \frac{d}{2}$

$$f\left(\frac{k}{d}\right) \geq f\left(\frac{1}{d}\right) \geq \frac{1}{2} \geq \frac{k}{d}$$

In addition $f_d(0) > 0$. It follows that

$$(45) \quad f\left(\frac{k}{d}\right) \geq \frac{k}{d} \text{ for all } 0 \leq k \leq \frac{d}{2}.$$

Lemma 8. For any $x \in \left(0, \frac{1}{2}\right]$ and $0 \leq k \leq \frac{d}{2}$ we have

$$(46) \quad \begin{aligned} x^k (1-x)^{d-k} f\left(\frac{k}{d}\right) + x^{d-k} (1-x)^k f\left(\frac{d-k}{d}\right) \\ > x^k (1-x)^{d-k} \frac{k}{d} + x^{d-k} (1-x)^k \left(1 - \frac{k}{d}\right). \end{aligned}$$

Proof. Denote $V_k = x^k (1-x)^{d-k}$ and $W_k = x^{d-k} (1-x)^k$. The assumption in the Lemma imply that $V_k \geq W_k$. Using inequality (45) we obtain that

$$(47) \quad (V_k - W_k)f\left(\frac{k}{d}\right) \geq (V_k - W_k)\frac{k}{d}.$$

Note that $W_k > 0$, so lemma 4 implies that

$$(48) \quad W_k \left(f\left(\frac{k}{d}\right) + f\left(\frac{d-k}{d}\right) \right) > W_k \left(\frac{k}{d} + 1 - \frac{k}{d} \right).$$

Lemma 8 follows by adding up inequalities (47) and (48). ■

Weighing inequality (46) by $\binom{d}{k} = \binom{d}{d-k}$ and summing up for $k = 0, 1, \dots, \lfloor \frac{d}{2} \rfloor$ we obtain

$$f_d(x) > \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{k}{d} = x.$$

We have now established the two upper bounds on the critical payoff gain in the case of partial information, namely $h(\beta)$ and $d - 2$. This concludes the proof of Theorem 3. ■

Partial information also leads to fast learning in an absolute sense. Figure 9 illustrates the waiting times until the adoption rate is within $\delta = 1\%$ of the high equilibrium for $d = 4$ (left panel), and $d = 10$ (right panel). The blue, dashed line represents the critical payoff gain $\alpha^*(\beta, d)$, such that for any $\alpha > \alpha^*(\beta, d)$ learning is fast. Pairs (ε, α) on the red line correspond to waiting times $T_N(\alpha, \beta, d, x_H - \delta)$ of 40 revisions per capita, while on the orange and green lines the corresponding waiting times are 20 and 10 revisions per capita, respectively.




Figure 9 – Waiting times to within $\delta = 1\%$ of the high equilibrium, for $d = 4$ (left panel) and $d = 10$ (right panel). Waiting times on solid lines are 40 (red), 20 (orange) and 10 (green) revisions per capita.

The following result provides a characterization of the expected waiting time in terms of the payoff gain, error rate and precision level.

Theorem 4. For any $d \geq 3$, for any noise level β satisfying $\alpha^*(\beta, d) > 0$ and for any precision level $\delta > 0$, there exists a constant $S = S(\beta, \delta, d)$ such that for all $\alpha > \alpha^*(\beta, d)$ and all sufficiently large N given α , the expected waiting time $T_N = T_N(\alpha, \beta, d, x_H - \delta)$ until the adoption rate is within δ of the high equilibrium satisfies

$$(49) \quad E(T_N) \leq \frac{S}{\sqrt{\frac{\alpha + 2}{\alpha^*(\beta, d) + 2} - 1}} + \log\left(\delta^{-1} \varepsilon^{-\frac{1}{d-1}}\right)$$

The intuition behind Theorem 4 is the same as for Theorem 2. As the payoff gain α approaches the threshold $\alpha^*(\beta, d)$, a “bottleneck” appears for intermediate adoption rates, which slows down the process. This effect is captured by the first term on the right hand side of (49). When the payoff gain is large, the main constraining factor is the precision level δ , namely how close we want the process to be to the high equilibrium. The second term on the right hand side of (49) takes care of this aspect.

Figure 10 plots, in log-log format, the expected time to get within $\delta = 1\%$ of the high equilibrium as a function of the payoff gain difference $\alpha - \alpha^*$. The error rates are $\varepsilon = 1\%$ (left panels) and $\varepsilon = 5\%$ (right panels), and the information parameters are $d = 4$ (top panels) and $d = 10$ (bottom panels). The blue, dashed line shows the simulated deterministic time $T(\alpha, \beta, d, x_H - \delta)$, while the red, solid line represents the estimated upper-bound using the following values for the constant S :




Figure 10 – Waiting time until the process is within $\delta = 1\%$ of the high equilibrium, as a function of the payoff gain difference $\alpha - \alpha^*$. Both axes have logarithmic scales. Simulation (blue, dashed) and upper bound (red, solid). Information parameters and error rates, and values for constant S :

$d = 4, \varepsilon = 1\%, S = 6.3$ (upper left panel), $d = 4, \varepsilon = 5\%, S = 8.6$ (upper right panel)
 $d = 10, \varepsilon = 1\%, S = 4.9$ (lower left panel), $d = 10, \varepsilon = 5\%, S = 5.9$ (lower right panel)

Proof of Theorem 4. The proof of Theorem 4 is essentially identical to the proof of Theorem 2. Here we outline the main steps in the argument, and refer the reader to the proof of Theorem 2 for the details.

The first step is to show that as the payoff gain α approaches the critical payoff gain $\alpha^* = \alpha^*(\beta, d)$ from above, the waiting time scales with the inverse square root of the gap, i.e. the height of the “bottleneck.” The second step is to show that for α close to α^* the gap is approximately linear in the payoff gain difference ($\alpha^* - \alpha$). These arguments are very similar to those in the proof of Theorem 2, hence we omit the details.

The third step consists of showing that the limit waiting time as the payoff gain tends to infinity is smaller than the second term on the right hand-side of (49). The details of this argument are presented below. This step is necessary because for α large relative to δ the constraining factor on the waiting time is the precision level δ .

The first two steps deal with inequality (49) for low values of α , while the third step takes care of high values of α . The intermediate values of α are covered by an appropriate choice of S in the first term of (49).

We now find the limit of the waiting time as the payoff gain tends to infinity. Identity (4) readily implies that

$$\lim_{\alpha \rightarrow \infty} f(x; \alpha, \beta) = \begin{cases} 1, & x > 0, \\ \varepsilon, & x = 0. \end{cases}$$

Plugging this into identity (37) we find that for all $x \in [0, 1]$

$$(50) \quad \begin{aligned} \lim_{\alpha \rightarrow \infty} f_d(x; \alpha, \beta) &= (1-x)^d \varepsilon + \sum_{k=1}^d \binom{d}{k} x^k (1-x)^{d-k} \\ &= 1 - (1-\varepsilon)(1-x)^d \end{aligned}$$

Denote by x_α the solution to the ordinary differential equation

$$\dot{x} = f_d(x; \alpha, \beta) - x \quad \text{with initial condition } x(0) = 0.$$

Denote by x_∞ the solution to the ordinary differential equation

$$\dot{x} = 1 - (1-\varepsilon)(1-x)^d - x \quad \text{with initial condition } x(0) = 0.$$

Note that convergence in (50) is uniform in $x \in [0, 1]$. This implies that x_α converges to x_∞ pointwise, in the sense that for any $t \geq 0$ we have $\lim_{\alpha \rightarrow \infty} x_\alpha(t) = x_\infty(t)$.

It can be checked that

$$x_\infty(t) = 1 - \left(1 - \varepsilon + \varepsilon e^{(d-1)t}\right)^{-\frac{1}{d-1}}$$

The limit waiting time T satisfies $x_\infty(T) = 1 - \delta$, which yields

$$T = \frac{1}{d-1} \log\left(\frac{\delta^{-(d-1)} - (1-\varepsilon)}{\varepsilon}\right) < \log\left(\delta^{-1} \varepsilon^{-\frac{1}{d-1}}\right)$$

We conclude that

$$\lim_{\alpha \rightarrow \infty} T(\alpha, \beta, d, x_H - \delta) < \log\left(\delta^{-1} \varepsilon^{-\frac{1}{d-1}}\right)$$

This establishes an upper bound for the waiting time to get close to the high equilibrium for sufficiently large payoff gain α . Together with the first two steps and an appropriate choice for the constant S , inequality (49) holds for all $\alpha > \alpha^*(\beta, d)$. ■

7. Extensions

This paper has examined the long-run behavior of an evolutionary model of equilibrium selection when the noise is bounded away from zero. The key finding is that the stochastically stable equilibrium is established quite rapidly when the noise is small but not extremely small. The boundary between the fast and slow learning regimes is sharp. If the potential gain between the two equilibria is below a critical threshold (for a given noise level), the expected waiting time to get close to the stochastically stable equilibrium is unbounded in the population size. However, if the potential gain is above the critical threshold the expected waiting time is bounded. Furthermore there is a tradeoff between the level of noise and the critical potential gain: a higher noise level allows fast learning to occur for lower potential gains.

The waiting times can be quite short – on the order of twenty to fifty revision opportunities per player. These numbers are comparable to those obtained by Ellison (1993) in a model of local interaction, but here the mechanism that produces fast learning is quite different. In local interaction models there are typically many equilibria of the associated deterministic dynamic; convergence occurs quickly because local groups can establish the stochastically stable equilibrium independently of the rest of the population. Hence equilibration occurs in parallel across the population at a rate that is more or less independent of the population size. Under global interaction, by contrast, there are at most three equilibria of the associated deterministic dynamic. The key observation is that when the noise is large enough, two of these equilibria disappear (the low and the middle); hence the process does not get stuck at a low adoption level. The expected motion towards a neighborhood of the high equilibrium is positive and bounded away from zero, and this remains true when the population is finite and large.

A natural question is whether our results hinge on the specific features of the logit response function. It turns out that this is not the case. The same line of argument applies to a large family of noisy response functions that are qualitatively similar to the logit, i.e., that are increasing in the payoff gain and approach full adoption as the payoff gain increases. In fact similar results hold even when the pure response function involves a constant rate of error (independent of the payoff gain) and agents have partial information based on finite samples. The reason is that sampling combined with error produces a smooth response function that is qualitatively similar to the logit, and that makes it easier for the system to escape from the low equilibrium.

Figure 11 illustrates this phenomenon. The blue, dashed line shows the response function for error rate $\varepsilon = 5\%$, sample size $d = 5$, and payoff gain $\alpha = 40\%$. This situation leads to slow learning, because the process gets stuck in the low equilibrium where the response function first crosses the

45-degree line. However, if the gain is somewhat larger ($\alpha = 50\%$) we obtain the green, solid line, which first crosses the 45-degree line near the high equilibrium. Just as in the case of the logit, one can estimate the critical payoff threshold that leads to fast learning for each given level of noise.




Figure 11 – Response functions for the uniform error model with $\varepsilon = 5\%$, sample size $d = 5$, and $\alpha = 40\%$ (blue, dashed line), $\alpha = 50\%$ (green, solid line).

Like many other papers in the evolutionary literature we have focused on 2×2 games in the interest of analytical tractability. This restriction allows us to obtain precise estimates of the critical values that separate fast from slow convergence. In principle a similar approach can be applied to larger games that have more than two equilibria, though the critical values will depend in a more complex way on the payoffs. As in the 2×2 case, the basic logic is that for a given noise level, the stochastically stable equilibrium will be established quickly if there is a sufficiently large gap in stochastic potential between it and the other equilibria.

Appendix

Lemma 3. (adapted from Lemma 1 in Benaïm and Weibull 2003) For any $T > 0$ there exists a constant $c = c(T) > 0$ such that for any $\mu > 0$ and N sufficiently large

$$\Pr(|x(T) - \hat{x}_N(T)| \geq \mu) \leq 2e^{-\mu^2 c N}.$$

Proof of Lemma 3. The deterministic adoption rate x satisfies the differential equation

$$\dot{x} = F(x) = f(x) - x.$$

The function F is Lipschitz continuous; denote its Lipschitz constant by λ .

Note that F also describes the expected change in the stochastic adoption rate x_N , so define

$$U_k = \frac{1}{\tau} \left(x_N((k+1)\tau) - x_N(k\tau) \right) - F(x_N(k\tau)).$$

By the previous observation $EU_k = 0$. The step extension of U_k is defined as $U(t) = U_k$ for all $t \in [k\tau, (k+1)\tau)$.

The deterministic adoption rate satisfies the integral equation

$$x(t) = \int_0^t F(x(s)) ds.$$

The stochastic adoption rate satisfies

$$\begin{aligned} \hat{x}_N(t) &= \int_0^t \tau (\bar{x}_N(t+1) - \bar{x}_N(t)) ds \\ &= \int_0^t \left(F(\bar{x}_N(s)) + U(s) \right) ds \\ &= \int_0^t \left(F(\hat{x}_N(s)) + \left(F(\bar{x}_N(s)) - F(\hat{x}_N(s)) \right) + U(s) \right) ds. \end{aligned}$$

Hence the difference between the deterministic and the stochastic processes is

$$\begin{aligned} |x(t) - \hat{x}_N(t)| &= \left| \int_0^t \left(\left(F(x(s)) - F(\hat{x}_N(s)) \right) + \left(F(\bar{x}_N(s)) - F(\hat{x}_N(s)) \right) + U(s) \right) ds \right| \\ &\leq \lambda \int_0^t |x(s) - \hat{x}_N(s)| ds + \lambda \tau T + \left| \int_0^t U(s) ds \right|, \end{aligned}$$

where the inequality uses that F is Lipschitz with constant λ and that $|\bar{x}_N(t) - \hat{x}_N(t)| < \tau$ for all t . Denote $\Psi(T) = \max_{0 \leq t \leq T} \left| \int_0^t U(s) ds \right|$. The above inequality shows that $|x(t) - \hat{x}_N(t)|$ grows at most exponentially in t . Specifically, Grönwall's inequality says that

$$\begin{aligned} |x(T) - \hat{x}_N(T)| &\leq (\lambda \tau T + \Psi(T)) e^{\lambda T} \\ &= \lambda \tau T e^{\lambda T} + \Psi(T) e^{\lambda T}. \end{aligned}$$

To bound the first term we take N sufficiently large, specifically $N \geq \frac{2\lambda T e^{\lambda T}}{\mu}$. This implies that $\tau \leq \frac{\mu}{2\lambda T e^{\lambda T}}$ and hence $\lambda\tau T e^{\lambda T} \leq \frac{\mu}{2}$. The remainder of the proof will be concerned with bounding $\Pr\left(\Psi(T)e^{\lambda T} \geq \frac{\mu}{2}\right)$. The following lemma will be useful.

Lemma 9. Let \mathcal{F}_k denote the σ -algebra generated by $\{x_N(t) : t \leq k\tau\}$. For any $\theta \in \mathbb{R}$

$$(51) \quad E\left(e^{\theta U_k} \mid \mathcal{F}_k\right) \leq e^{2|\theta|^2}.$$

Proof. Make the transformation $a = \theta U_k$ and note that $E(a) = 0$. The function $g(t) = \log E(e^{ta})$ satisfies $g' = \frac{E(ae^{ta})}{E(e^{ta})}$ and $g'' = \frac{E(a^2 e^{ta})E(e^{ta}) - E(ae^{ta})^2}{E(e^{ta})^2}$. It follows that $g(0) = g'(0) = 0$ and $g'' > 0$ by virtue of the Cauchy-Schwartz inequality. Moreover, $g''(t) \leq \frac{E(a^2 e^{ta})}{E(e^{ta})} \leq 4|\theta|^2$ because $|U_k| \leq 2$. It follows that $g(t) \leq \frac{4|\theta|^2}{2}$ hence $g(1) \leq 2|\theta|^2$. ■

To estimate $\Pr(\Psi(T) \geq \beta)$, define

$$Z_k = \exp\left(\sum_{i=0}^{k-1} \tau\theta U_i - 2k\tau^2|\theta|^2\right)$$

By Lemma 9, Z_k is a supermartingale, that is, $E(Z_k | Z_{k-1}, \dots, Z_1) \leq Z_{k-1}$. One can inductively define a martingale Y_k such that almost surely $Z_k \leq Y_k$, for all k . We have:

$$\begin{aligned} \Pr\left(\max_{0 \leq t \leq T} \int_0^t \theta U(s) \geq \gamma\right) &= \Pr\left(\max_{0 \leq k \leq T/\tau} \exp\left(\sum_{i=0}^{k-1} \tau\theta U_i\right) \geq \exp(\gamma)\right) \\ &\leq \Pr\left(\max_{0 \leq k \leq T/\tau} Z_k \geq \exp(\gamma - 2(T/\tau)\tau^2|\theta|^2)\right) \\ &\leq \Pr\left(\max_{0 \leq k \leq T/\tau} Y_k \geq \exp(\gamma - 2(T/\tau)\tau^2|\theta|^2)\right) \\ &\leq \exp\left(2\left(\frac{T}{\tau}\right)\tau^2|\theta|^2 - \gamma\right) \\ &= \exp(2T\tau|\theta|^2 - \gamma). \end{aligned}$$

(Here we have used Doob's martingale inequality to pass from line 3 to line 4.) Setting $\theta = \pm \frac{2\gamma}{\mu e^{\lambda T}}$ and adding up the probabilities we obtain:

$$\Pr\left(\max_{0 \leq t \leq T} \left|\int_0^t U(s)\right| \geq \frac{\mu}{2} e^{-\lambda T}\right) \leq 2 \exp\left(2T\tau \frac{4\gamma^2}{\mu^2 e^{-2\lambda T}} - \gamma\right).$$

It is optimal to choose $\gamma = \frac{\mu^2 e^{-\lambda T}}{16T\tau}$, in which case

$$\Pr\left(\Psi(T) \geq \frac{\mu}{2} e^{-\lambda T}\right) \leq 2 \exp\left(-\frac{\mu^2 e^{-\lambda T}}{32T\tau}\right).$$

Finally, noting that $\tau = \frac{1}{N}$, we obtain the desired inequality with $c = \frac{e^{-\lambda T}}{32T}$. ■

Lemma 6. The function f_d is first strictly convex then strictly concave, and the inflection point is at most $\frac{1}{2}$.

Proof of Lemma 6. The proof has two steps. Firstly, we shall use a monotone likelihood ratio argument to show that the second derivative of f_d is initially strictly positive and then strictly negative. Secondly, we shall prove that $f_d''\left(\frac{1}{2}\right) \leq 0$, which implies that the inflection point of f_d is at most $\frac{1}{2}$.

We begin by introducing some notation. Recall that $\Delta_d f(k)$ denotes the discrete derivative of f , namely,

$$\Delta_d f(k) = \frac{f\left(\frac{k+1}{d}\right) - f\left(\frac{k}{d}\right)}{1/d}.$$

For each $k = 0, \dots, d-2$ let

$$s_k = \binom{d-2}{k} \frac{(\Delta_d f(k+1) - \Delta_d f(k))}{1/(d-1)}.$$

For every $x \in [0,1]$ let

$$h_k(x) = x^k(1-x)^{d-2-k}.$$

We claim that $\Delta_d f(k)$ is single-peaked in k . To see this, let x_I be the inflection point of f , and let k_I be the integer that satisfies $\frac{k_I}{d} \leq x_I < \frac{k_I+1}{d}$. Clearly $\Delta_d f(k)$ is increasing for $k \leq k_I - 1$, and decreasing for $k \geq k_I + 1$. We are left to prove that $\Delta_d f(k_I)$ is larger than either $\Delta_d f(k_I - 1)$ or $\Delta_d f(k_I + 1)$. This follows according to whether the point $(x_I, f(x_I))$ lies below or above the line uniting the points $\left(\frac{k}{d}, f\left(\frac{k}{d}\right)\right)$ and $\left(\frac{k+1}{d}, f\left(\frac{k+1}{d}\right)\right)$. Assume we are in the former case (the other case is similar), then we can write successively

$$\begin{aligned} \Delta_d f(k_I) &\geq \frac{f(x_I) - f\left(\frac{k_I}{d}\right)}{x_I - \frac{k_I}{d}} \\ &> \Delta_d f(k_I - 1). \end{aligned}$$

The last inequality follows because f is strictly convex on $[0, x_I)$.

The result in the last paragraph implies that the sequence s_0, \dots, s_{d-2} satisfies single-crossing, in the sense that there exists $k_0 \in \{k_I - 1, k_I\}$ such that $s_k > 0$ for $k < k_0$ and $s_k < 0$ for $k > k_0$.

The sequence of functions h_0, h_1, \dots, h_{d-2} has a monotone likelihood ratio in x , that is for any $k = 0, 1, \dots, d-3$ the ratio $\frac{h_{k+1}(x)}{h_k(x)}$ is increasing in x . To see this take $0 < x < y < 1$ and note that

$$\frac{h_{k+1}(x)}{h_k(x)} < \frac{h_{k+1}(y)}{h_k(y)} \Leftrightarrow \frac{x}{1-x} < \frac{y}{1-y}.$$

Differentiating the expression for $f'_d(x)$ in (38) and rearranging terms we obtain

$$(52) \quad \begin{aligned} f''_d(x) &= \sum_{k=0}^{d-2} \binom{d-2}{k} x^k (1-x)^{d-2-k} \frac{\Delta_d f(k+1) - \Delta_d f(k)}{1/(d-1)} \\ &= \sum_{k=0}^{d-2} s_k h_k(x). \end{aligned}$$

The following result shows that f''_d is first strictly positive and then strictly negative. It follows that f_d is first strictly convex and then strictly concave.

Lemma 10. With the above notation, the function f''_d satisfies single-crossing from positive to negative, in the sense that there exists $x_0 \in [0,1]$ such that $f''_d(x) > 0$ for $x < x_0$ and $f''_d(x) < 0$ for $x > x_0$.

Proof. We shall prove that whenever $0 < x < y < 1$ and $f''_d(x) \leq 0$ then $f''_d(y) < 0$. Let $\lambda = \frac{h_{k_0}(y)}{h_{k_0}(x)} > 0$.

For every $k < k_0$ we have $h_k(y) < \lambda h_k(x)$ and $s_k > 0$, hence

$$(53) \quad s_k h_k(y) < \lambda s_k h_k(x).$$

For every $k > k_0$ we have $h_k(y) > \lambda h_k(x)$ and $s_k < 0$, hence

$$(54) \quad s_k h_k(y) < \lambda s_k h_k(x).$$

Finally, for $k = k_0$ we have

$$(55) \quad s_k h_k(y) = \lambda s_k h_k(x).$$

Adding up expressions (53), (54) and (55) for $k = 0, 1, \dots, d-2$ we obtain that $f''_d(y) < \lambda f''_d(x) \leq 0$. ■

We shall now show that $f''_d\left(\frac{1}{2}\right) \leq 0$ by direct computation. Rearranging the terms in (52) yields

$$\begin{aligned} f''_d\left(\frac{1}{2}\right) &= \frac{1}{2^{d-2}} \sum_{k=0}^{d-2} \binom{d-2}{k} \frac{\Delta_d f(k+1) - \Delta_d f(k)}{1/(d-1)} = \frac{1}{2^{d-2}} \sum_{k=0}^{d-1} (2k - (d-1)) \binom{d-1}{k} \Delta_d f(k) \\ &= \frac{1}{2^{d-2}} \sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} f\left(\frac{k}{d}\right). \end{aligned}$$

It will be useful to cast the last expression into a symmetric form. Let $Q(x) = f(x) + f(1-x)$, then the inequality to prove becomes

$$(56) \quad \sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq 0.$$

In outline, the remainder of the proof runs as follows. We shall first establish that Q is increasing on the interval $\left[0, \frac{1}{2}\right]$. Next, we shall show that as k increases from 0 to $\left[\frac{d}{2}\right]$ the coefficient $((2k-d)^2 - d)$ is first positive and then negative. These two facts imply that in order to prove inequality (56), it is

sufficient to prove the inequality after dropping the Q term. The last part of the proof will establish inequality (56) after dropping the Q term.

Claim 3. Q is increasing on the interval $\left[0, \frac{1}{2}\right]$.

To establish this claim, fix $x \in \left[0, \frac{1}{2}\right]$ and let $X = f(x)$ and $Y = f(1-x)$. Then

$$\begin{aligned} Q'(x) &= f'(x) - f'(1-x) \\ &= \beta(\alpha+2)(X - X^2 - Y + Y^2) \\ &= \beta(\alpha+2)(X - Y)(1 - (X + Y)) \end{aligned}$$

The last expression is positive because $X \leq Y$ and $1 \leq X + Y$ by Lemma 4.

Note that as k increases from 0 to $\left\lfloor \frac{d}{2} \right\rfloor$ the coefficient $((2k-d)^2 - d)$ is first positive and then negative. Let k^* be the smallest k such the coefficient $((2k-d)^2 - d)$ is negative. In other words, the coefficient is negative when $k^* \leq k \leq d - k^*$, and non-negative otherwise.

For every k such that $k^* \leq k \leq d - k^*$, we have that $((2k-d)^2 - d) < 0$ and $Q\left(\frac{k}{d}\right) \geq Q\left(\frac{k^*}{d}\right)$. Hence

$$(57) \quad ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k^*}{d}\right).$$

For every k such that $k < k^*$ or $d - k^* < k$, we have that $((2k-d)^2 - d) \geq 0$ and $Q\left(\frac{k}{d}\right) \leq Q\left(\frac{k^*}{d}\right)$. Hence

$$(58) \quad ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k^*}{d}\right).$$

From inequalities (57) and (58) it follows that

$$\sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq Q\left(\frac{k^*}{d}\right) \left(\sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} \right).$$

The final step is to establish that

$$\sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} = 0.$$

We use the identities

$$\sum_{k=0}^d \binom{d}{k} k = d2^{d-1} \quad \text{and} \quad \sum_{k=0}^d \binom{d}{k} k^2 = d(d+1)2^{d-2}.$$

We now write

$$\begin{aligned}
\sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} &= 4 \sum_{k=0}^d \binom{d}{k} k^2 - 4d \sum_{k=0}^d \binom{d}{k} k + (d^2 - d) \sum_{k=0}^d \binom{d}{k} \\
&= 4d(d+1)2^{d-2} - 4d^2 2^{d-1} + d(d-1)2^d \\
&= 0.
\end{aligned}$$

Therefore the inflection point of f_d is at most $\frac{1}{2}$. ■

References

- Anderson, R.M. and May, R. (1991): Infectious diseases of humans: dynamics and control, Oxford University Press.
- Anderson, S., De Palma, A. and Thisse, J. (1992): Discrete Choice Theory of Product Differentiation, Cambridge, U.S.A.: MIT Press.
- Benaïm, M., and Weibull, J. (2003): Deterministic Approximation of Stochastic Evolution In Games, *Econometrica*, 71(3), 873-903.
- Blume, L.E. (1993): The Statistical Mechanics of Strategic Interaction, *Games and Economic Behavior*, 5(3), 387–424.
- Blume, L.E. (1995): The Statistical Mechanics of Best-Response Strategy Revision, *Games and Economic Behavior*, 11(2), 111–145.
- Blume, L.E. (2003): How noise matters, *Games and Economic Behavior*, 44(2), 251–271.
- Brock, W. A. and Durlauf, S. N. (2001): Interactions-based models, *Handbook of Econometrics*, in: Heckman, J.J. and Leamer, E.E. (eds.), *Handbook of Econometrics*, 1(5), chapter 54, 3297–3380.
- Ellison, G. (1993): Learning, Local Interaction, and Coordination, *Econometrica*, 61(5), 1047-71.
- Foster, D. and Young, H.P. (1990): Stochastic Evolutionary Game Dynamics, *Theoretical Population Biology*, 38:219-232.
- Freidlin, M. and Wentzell, A. (1984): *Random Perturbations of Dynamical Systems*. Berlin: Springer-Verlag.
- Jackson, M.O. and Yariv, L. (2007): Diffusion of Behavior and Equilibrium Properties in Network Games, *American Economic Review*, 97(2), 92-98.
- Kandori, M., Mailath, G.J. and Rob, R. (1993): Learning, Mutation, and Long Run Equilibria in Games, *Econometrica*, 61(1), 29-56.

- Kurtz, T.G. (1970): Solutions of ordinary differential equations as limits of pure jump Markov processes, *Journal of Applied Probability*, 7, 49-58.
- Lelarge, M. (2010): Diffusion and Cascading Behavior in Random Networks, eprint arXiv:1012.2062.
- López-Pintado, D. (2006): Contagion and coordination in random networks, *International Journal of Game Theory*, 34(3), 371-381.
- McFadden, D.L. (1976): Quantal Choice Analysis: A Survey, NBER Chapters, in: *Annals of Economic and Social Measurement*, 5(4), 8-66.
- Montanari, A. and Saberi, A. (2010): The spread of innovations in social networks, *Proceedings of the National Academy of Sciences*, 107(47), 20196-20201.
- Morris, S. (2000): Contagion, *Review of Economic Studies*, 67(1), 57-78.
- Shah, D. and Shin, J. (2010): Dynamics in Congestion Games, *Proceedings of the ACM SIGMETRICS Conference*.
- Watts, D.J. (2002): A simple model of global cascades on random networks, *Proceedings of the National Academy of Sciences*, 99(9), 5766–5771.
- Watts and Dodds (2007): Influentials, Networks, and Public Opinion Formation, *Journal of Consumer Research*, 34(4), 441-445.
- Young, H.P. (1993): The Evolution of Conventions, *Econometrica* 61:57-84.
- Young, H.P. (1998): *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton, NJ: Princeton University Press.
- Young, H.P. (2011): The dynamics of social innovation, *Proceedings of the National Academy of Sciences*, forthcoming.