



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb



Fast convergence in evolutionary equilibrium selection

Gabriel E. Kreindler^{a,*}, H. Peyton Young^b^a Nuffield College, New Road, Oxford OX1 1NF, United Kingdom^b Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, United Kingdom

ARTICLE INFO

Article history:

Received 3 July 2012

Available online 5 March 2013

JEL classification:

C72

C73

Keywords:

Stochastic stability

Logit learning

Markov chain

Convergence time

ABSTRACT

Stochastic best response models provide sharp predictions about equilibrium selection when the noise level is arbitrarily small. 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. Moreover, the expected waiting times are comparable to those in local interaction models.

© 2013 Elsevier Inc. All rights reserved.

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 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 myopic best response with a random component resulting from errors, unobserved utility shocks, or experiments. The long-run behavior of the resulting stochastic dynamical system 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 et al., 1993; Young, 1993).

A common criticism of this approach is that the waiting time to get close to the ergodic distribution may be exceptionally large. If the noise level is small and the system starts near an equilibrium that is *not* stochastically stable, it will remain close to this equilibrium for a very long period of time. Indeed, the time it takes to escape the initial (“wrong”) equilibrium grows exponentially in the population size when the noise is sufficiently small (Ellison, 1993; Sandholm, 2010). 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 could be that at intermediate levels of noise the selection process displays a fairly strong bias towards the stochastically stable states. If this is so the speed of convergence could be quite rapid.

* Correspondence to: G.E. Kreindler, Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, MA 02142, United States.

E-mail addresses: gek@mit.edu (G.E. Kreindler), peyton.young@economics.ox.ac.uk (H.P. Young).

A pioneering paper by Ellison (1993) shows that this is indeed the case when agents interact “locally” with small groups of neighbors. The paper deals with the specific case when agents are located at the nodes of a ring network and they are linked with agents within a specified distance. The form of interaction is a symmetric 2×2 coordination game. Ellison shows that the waiting time to get close to the stochastically stable equilibrium, say all- A , is bounded independently of population size, and that its absolute magnitude may be very small. The reason that this set-up leads to fast selection is local reinforcement. When the noise is sufficiently small (but not taken to zero), any small close-knit group of interconnected agents will not take long to adopt the action A . Once they have done this, they will continue to play A with high probability thereafter even when people outside the group are playing the alternative action B . Since this occurs in parallel across the entire network (independently of population size), it does not take long in expectation until almost all of the close-knit groups, and hence a large proportion of the population, have switched to A . In fact this argument is quite general and applies to a variety of “local” network structures and stochastic learning rules, as shown in Young (1998, 2011).

A separate line of work studies the conditions for equilibrium selection when agents (myopically) best respond to their available information. Sandholm (2001) considers the case when agents best respond to random samples of size d , and shows that any $1/d$ -dominant equilibrium is eventually reached with high probability from almost any initial condition, for sufficiently large population size.¹ In particular, when the sample size is two and interaction is given by a 2×2 symmetric coordination game with a unique risk-dominant equilibrium, say (A, A) , this equilibrium will be reached from any initial condition in which a positive fraction of the population plays A . The intuition for this result is that for almost any population action profile, the expected change in the proportion of A -players is strictly positive. The key idea is to work with the deterministic approximation of the process corresponding to an infinite population; in this setting the A -equilibrium is an almost global attractor. When the population is large but finite, the process is stochastic but well approximated by the deterministic dynamics, and convergence occurs in bounded time with high probability. (A similar deterministic approximation is used to analyze the diffusion of innovations in random networks; see López-Pintado (2006) and Jackson and Yariv (2007).) However, these papers do not attempt to characterize the expected waiting time to reach equilibrium.

The contribution of this paper is to provide an in-depth analysis of expected waiting times in evolutionary models with global interaction and stochastic best response dynamics. The model assumes a large, finite population of identical agents repeatedly interacting according to a symmetric 2×2 coordination game. Agents occasionally have the opportunity to revise their actions, and when they do they choose a perturbed best response to the distribution of actions in the population. We shall consider the case where agents know the actual distribution of actions and also the case where they only observe a finite random sample of actions. In the interest of analytical tractability we shall restrict attention to 2×2 coordination games and logit best response dynamics (Blume, 1993, 1995). The method of analysis applies to a much broader class of response functions, as we show in Section 7. For the sake of concreteness we shall call action A the *innovation* and action B the *status quo* and talk about the *adoption rate of the innovation* A starting from the status quo when everyone is playing B .

We can characterize the expected waiting time until most players have adopted A in terms of two easily interpretable parameters: the payoff gain α of the innovation relative to the status quo, and the noise level $1/\beta$, where β is the coefficient of the logit function. The analysis proceeds in four steps. First we characterize the deterministic (mean-field) dynamics for given levels of α and β , assuming an infinite population. Next, we characterize the combinations of α and β such that the deterministic dynamics have a unique rest point.² We then exploit the geometry of the aggregate response function to characterize the maximum length of time it takes for the process to reach a neighborhood of the unique rest point. Finally, we apply stochastic approximation theory (Benaïm and Weibull, 2003) to show that the estimations for the deterministic system are valid for the case of large but finite populations.

The main results are the following:

- (i) For a given payoff gain $\alpha > 0$ the dynamics exhibit a phase transition with respect to the noise level: there is a critical noise level below which the waiting time is exponential in the number of agents (selection is “slow”), and above which the waiting time is bounded irrespective of the number of agents (selection is “fast”).
- (ii) We provide an explicit estimate of the critical noise threshold, which turns out to be quite small. In particular, when $\frac{1}{\beta} > \frac{1}{2}$ selection is fast for *all* positive values of α . This threshold corresponds to an initial error rate – the probability of choosing A when everyone else is choosing B – of about 12%.
- (iii) For plausible parameter values the absolute magnitudes of the waiting times are very small. For example, when the innovation is 100% better than the status quo and the agents’ initial error rate is 5%, it takes less than 20 revisions per capita in expectation to reach an adoption rate of 99%.

These results can also be interpreted in terms of heterogeneity in individual payoffs rather than as noisy best responses. In particular, suppose that each individual has an idiosyncratic payoff from playing any given strategy, in addition to the payoff from the coordination game. Assume that these payoffs are fixed across time, and that they are drawn independently

¹ An equilibrium (A, A) is $1/d$ -dominant if A is a best response to any sample of size d that contains at least one player playing A .

² A number of authors have observed that when the noise in the logit response function is large enough, the deterministic (mean-field) dynamic has a unique global attractor (McKelvey and Palfrey, 1995; Blume and Durlauf, 2002; Sandholm, 2010, Example 6.2.2). However, they did not draw out the implications for convergence times.

across strategies and across agents from an extreme-value distribution with mean zero, parameter β and variance $\frac{\pi^2}{6\beta^2}$. In a large population the dynamics will behave as if the individuals were choosing via the logit process with parameter β . Consequently, the above results may be interpreted as saying that even quite low levels of variance in the payoffs can lead to fast selection.

The paper is organized 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 full 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 partial information case. In Section 7 we show that the method of analysis extends to many types of perturbed response functions in addition to the logit.

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 selection dynamics are deterministic (best response) or stochastic (noisy best response). 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:

	A	B	
A	$1 + \alpha, 1 + \alpha$	$0, 0$	$\alpha \geq 0$
B	$0, 0$	$1, 1$	

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 .

2.1. Deterministic best response, 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. A particularly interesting case occurs when all degrees in a connected network are at most d , and the payoff gain is larger than $d - 2$. In this case, any single initial adopter will cause the innovation to spread by contagion to the entire population. Morris does not address the issue of waiting times as such, rather, he identifies conditions that are necessary and sufficient for full adoption to occur under best response dynamics.

2.2. Noisy best response, 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 in this section, the population is *finite* and the selection process is *stochastic*. The aim of the analysis is to characterize the ergodic distribution of the process rather than to approximate it by a 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 non-best response), 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 small but not vanishing 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 selection process, whereas overall well-connected networks tend to be slow. For example, selection is slow on random networks for small enough $\alpha > 0$, while selection 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 large in either case.

2.3. Deterministic best response, global interaction and sampling

A number of authors have considered the situation where agents interact globally and choose best responses to random samples drawn from the population. Some of these models analyze the resulting stochastic process for finite populations, while others focus exclusively on the mean-field dynamics with a continuum of agents. We have already mentioned the pioneering work of Sandholm (2001), who observes that when all samples have size d , then any $1/d$ -dominant equilibrium becomes a global attractor except from degenerate initial conditions. In the setting of a 2×2 pure coordination game, the $1/d$ -dominant condition is equivalent to $\alpha > d - 2$.³ Sandholm’s results have recently been generalized by Oyama et al. (2011) to iterated $1/d$ -dominant equilibria and settings where the sample size is itself random. They show that when samples of size at most d are sufficiently likely, a condition they call *d-goodness*, then the iterated $1/d$ -equilibrium becomes an almost global attractor.

López-Pintado (2006) 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 positive initial fraction of A -players (however small), the process evolves to a state in which a positive proportion of the population plays A forever. Jackson and Yariv (2007) use a similar mean-field approximation technique to analyze 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. Finally, Watts (2002) and Lelarge (2012) 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 .

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

- (i) Deterministic best response and local interaction (Morris),
- (ii) Noisy best response and local interaction (Ellison, Young),
- (iii) Low noise best response and both local and global interaction (Montanari and Saberi), and
- (iv) Deterministic best response and global interaction with sampling (Sandholm, López-Pintado, etc.).

There remains the case of noisy dynamics with global interaction. We know from prior results in the literature that when the noise in the logit best response process is large enough, a single equilibrium can survive as the global attractor of the deterministic (mean-field) dynamic (McKelvey and Palfrey, 1995; Blume and Durlauf, 2002; Sandholm, 2010, Example 6.2.2; Hommes and Ochea, 2012). It follows that convergence to any neighborhood of that equilibrium occurs in bounded time independently of the population size. However, this does not address the question of how long it takes to get close to the equilibrium.⁴ The contribution of the present paper is to derive explicit estimations of the waiting time as a function of the model parameters. These estimations show that selection occurs very rapidly when the noise is fairly 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 initial error rate

³ Morris (2000) finds the same threshold for contagion to occur in infinite d -regular trees.

⁴ One of the few papers to estimate waiting times explicitly is Shah and Shin (2010). They consider a special class of potential games which does not include the case considered here, and prove that for intermediate values of noise the time to get close to equilibrium grows slowly (but is unbounded) in the population size.

of about 5% – where each individual revises once per period in expectation. The conclusion is that selection is fast 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

$$\begin{array}{cc} & A & B \\ A & a, a & c, d \\ B & d, c & b, b \end{array}$$

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

$$\begin{array}{cc} & A & B \\ A & a - d & 0 \\ B & 0 & b - c \end{array}$$

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

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

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 et al., 1993 and Young, 1993).

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

$$\begin{array}{cc} & A & B \\ A & 1 + \alpha, 1 + \alpha & 0, 0 \\ B & 0, 0 & 1, 1 \end{array} \tag{2}$$

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 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.⁵ 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:

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

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 or the *initial error rate*; given the bijective correspondence 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 et al., 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

⁵ 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 her action. It is possible to show that results in this article remain unchanged under this alternative revision protocol.

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 et al., 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 $x_N(t) \in \{0, \frac{1}{N}, \dots, \frac{N-1}{N}, 1\}$, 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 time⁶

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

Fast selection is defined as follows.

Definition 1. The family $\{\Gamma_N(\alpha, \beta): N > 0\}$ displays *fast selection* 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 $ET_N(\alpha, \beta, \frac{1}{2}) < S$ for all N .

More generally, for any $p < 1$,

Definition 2. The family $\{\Gamma_N(\alpha, \beta): N > 0\}$ displays *fast selection 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 $\Gamma_N(\alpha, \beta)$ displays fast selection, or fast selection to p .

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 selection 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 selection, where

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

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

The main message of Theorem 1 is that fast selection 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 selection is fast for any positive payoff gain $\alpha > 0$. Theorem 1 shows that fast selection can occur for any positive payoff gain even when interaction is global, provided that $\beta \leq 2$, which is equivalent to an initial error rate larger than approximately 12%.

The key idea is that for any payoff gain $\alpha \geq 0$ there exists a critical noise level $\beta^*(\alpha)$ such that $\Gamma_N(\alpha, \beta)$ displays fast selection for $\beta < \beta^*(\alpha)$.⁷ Intuitively, fast selection occurs when the deterministic approximation of the process $\Gamma_N(\alpha, \beta)$ as the population grows large has a unique equilibrium. Equilibria of the deterministic process correspond to fixed points of the response function. For small noise levels the logit response function has three fixed points; as the noise level increases the first part of the response function “lifts up” and eventually the low and middle fixed points disappear. (See Fig. 2 for an illustration of this change.) Similarly, for a given noise level, the low and middle fixed points disappear for sufficiently large payoff gains.

This point has been noted by various authors, notably Sandholm (2001, 2010) and Blume and Durlauf (2002); the contribution of Theorem 1 is to provide a sharp estimate of the critical payoff gain. Later in the paper (in Section 5) we examine the ramifications of this phenomenon for the speed of convergence.

⁶ 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.

⁷ Also, for each noise parameter β there exists a payoff gain threshold, denoted $h^*(\beta)$, such that $\Gamma_N(\alpha, \beta)$ displays fast selection for $\alpha > h^*(\beta)$.




Fig. 1. The critical threshold for fast selection. Simulation (blue, dashed) and upper bound (red, solid). The x-axis is labeled with the initial error rate ε ; the corresponding noise level β appears in parentheses.




Fig. 2. The response functions for payoff gain $\alpha = 100\%$ and initial error rate $\varepsilon = 3\%$ (left panel) and $\varepsilon = 6\%$ (right panel). The two systems have three equilibria and a unique equilibrium, respectively. In each case, the vertical dashed line corresponds to the indifference point of the pure best response.

Fig. 1 shows the simulated noise threshold $\beta^*(\alpha)$ (blue, dashed line) as well as the bound provided by Theorem 1 (red, solid line).⁸ The x-axis displays the initial error rate $\varepsilon = \frac{1}{1+e^\beta}$, and the y-axis represents payoff gains α . Note that the difference between the error rates given by the two curves never exceeds about 0.5%.

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. Fig. 3 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 initial error rate $\varepsilon = 5\%$.

More generally, Table 1 shows how the expected waiting time depends on the population size N , the payoff gain α , the initial error rate ε , and on the target adoption level p .⁹ The last column shows the limit of the waiting time as the population size tends to infinity. 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 also suggests that the expected waiting time when the population is finite is generally less than the waiting time as the population tends to infinity. We conjecture that this is due to the increased volatility of the process when there are finitely many agents.¹⁰

⁸ The blue, dashed line in Fig. 1 also represents the payoff gain threshold $h^*(\beta)$, while the red, solid line represents the bound $h(\beta)$. Fig. 1 also shows that the function h is continuous and that $h(2) = h^*(2) = 0$.

⁹ Note that the process $I_N(\alpha, \beta)$ is a birth–death process, and the expected waiting times to achieve a specific adoption level can be computed explicitly using standard formulas (see Example 11A.5 in Sandholm, 2010).

¹⁰ We thank an anonymous referee for pointing this out.




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

Table 1

Expected waiting times (full information).

	Expected waiting time ($\varepsilon = 5\%$)			
	$N = 100$	$N = 1000$	$N = 10,000$	$N = \infty$
$\alpha = 70\%a$	32.65	102.88	33,691.57	∞
$\alpha = 80\%a$	24.44	33.59	35.29	35.19
	Expected waiting time ($\varepsilon = 10\%$)			
	$N = 100$	$N = 1000$	$N = 10,000$	$N = \infty$
$\alpha = 4\%b$	34.03	190.96	142,712.40	∞
$\alpha = 25\%c$	18.72	20.08	20.67	20.83

N = population size, α = innovation payoff gain, ε = initial error rate. The target adoption rate is ^a $p = 99\%$, ^b $p = 50\%$ and ^c $p = 90\%$.

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 Fig. 1, Tables 3 and 4 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% .¹¹

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 selection is preserved by this approximation when N is large.

We begin by defining the deterministic approximation and the concepts of equilibrium, stability and fast selection 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

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

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

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

where the dot above x denotes the time derivative.¹²

An equilibrium of this system is a rest point, that is an adoption rate x^* satisfying $\dot{x}^* = 0$, which is equivalent to $f(x^*) = x^*$. Note that this equilibrium of the process $\Gamma(\alpha, \beta)$ corresponds to the Logit Quantal Response Equilibrium defined in McKelvey and Palfrey (1995). 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. Fig. 4 plots the position of the stable and unstable rest points as the initial error rate ε varies, for $\alpha = 25\%$.

¹¹ These error rates correspond to randomization probabilities of 2.5%, 5% and 10% respectively.

¹² Note that the expected change in the adoption rate in the stochastic process is given by $\frac{1}{N}(f(x) - x)$.




Fig. 4. Rest points of the deterministic process ($\alpha = 25\%$): stable rest points (blue, solid line) and unstable rest points (red, dashed line). The x -axis is labeled with the initial error rate ε ; the corresponding noise level β appears in parentheses.

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

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

Definition 3. The process $\Gamma(\alpha, \beta)$ displays *fast selection* 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 selection to p* if $T(\alpha, \beta, p) < \infty$.

Remark. A necessary and sufficient condition for fast selection is that all equilibria lie strictly above $\frac{1}{2}$. Similarly, fast selection to p holds if and only if 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 $[0, \frac{1}{2}]$.

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

$$f'(x) = \beta(\alpha + 2)f(x)(1 - f(x)) \tag{6}$$

$$f''(x) = \beta(\alpha + 2)f'(x)(1 - 2f(x)) \tag{7}$$

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

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

The result now follows by inspecting the sign of $f(x) - x$ for $x = 0, \frac{1}{2+\alpha}$ and 1. \square

We now define and then estimate the critical payoff gain. Using the convention $\inf \emptyset = \infty$, let

$$h^*(\beta) = \inf\{\alpha: \Gamma(\alpha, \beta) \text{ displays fast selection}\}$$

Note that f is strictly increasing in α .¹³ Given $\alpha' > \alpha$, if $f(\cdot; \alpha, \beta)$ does not have any equilibria in the interval $[0, \frac{1}{2}]$, then neither does $f(\cdot; \alpha', \beta)$. It follows that $\Gamma(\alpha, \beta)$ displays fast selection for all $\alpha > h^*(\beta)$.

¹³ Differentiating Eq. (4) we obtain

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

This quantity is positive for all $x, \beta > 0$, so f is increasing in α .

The estimation of the critical payoff gain $h^*(\beta)$ 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 $h^*(\beta) = 0$. When $\beta > 2$ then $h^*(\beta) > 0$.

To establish this claim, we study $x_L(\alpha, \beta)$, the smallest equilibrium of f . The function f is strictly increasing in α on $(0, 1)$, hence x_L is strictly increasing in α . Thus, to show that $h^*(\beta) = 0$ it is sufficient to show that $x_L(0, \beta) = \frac{1}{2}$.

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 less than or equal to 1, or strictly greater than 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 strict down crossing. It follows that $x_L(0, \beta) < \frac{1}{2}$ and for sufficiently small $\alpha > 0$ it still holds that $x_L(\alpha, \beta) < \frac{1}{2}$. This implies that $h^*(\beta) > 0$.

When $\beta > 2$ the critical payoff gain threshold is non-trivial, namely for small payoff gains the system has two equilibria smaller than $\frac{1}{2}$. At the critical payoff gain $h^*(\beta)$, there is a unique equilibrium smaller than $\frac{1}{2}$, namely the tangency point x^* between the function f and the 45-degree line. For given β , the equilibrium x^* and the critical payoff gain $h^* = h^*(\beta)$ solve the equations

$$f'(x^*; h^*, \beta) = 1, \quad \text{and} \tag{9}$$

$$f(x^*; h^*, \beta) = x^* \tag{10}$$

Using identity (6) and then applying (10), Eq. (9) becomes

$$\beta(h^* + 2)f(x^*) = \frac{1}{1 - f(x^*)} \Rightarrow \beta(h^* + 2)x^* = \frac{1}{1 - x^*} \tag{11}$$

Writing Eq. (10) explicitly and using (11) yields:

$$\frac{1}{1 + e^{\beta - (h^* + 2)x^*}} = x^* \Leftrightarrow \frac{1}{1 + e^{\beta - \frac{1}{1 - x^*}}} = x^*$$

The last equation uniquely identifies x^* as a function of β . Eq. (11) can be rewritten as a quadratic equation:

$$x^* - (x^*)^2 = \frac{1}{\beta(h^* + 2)} \tag{12}$$

There is a unique solution smaller than $\frac{1}{2}$ given by

$$x^* = \frac{1}{2} \left(1 - \sqrt{1 - \frac{4}{\beta(h^* + 2)}} \right)$$

This equation implies that h^* is uniquely determined given x^* , hence it is uniquely determined given β .

We now present an informal estimation of $h^*(\beta)$ based on taking β to be very large. Claim 2 establishes an upper bound for $h^*(\beta)$ for all $\beta > 2$. For large β the equilibrium x^* tends to zero, hence the quadratic term in (12) is negligible. This means that

$$x^* \approx \frac{1}{\beta(h^* + 2)} \tag{13}$$

Using this approximation twice, Eq. (10) becomes

$$\begin{aligned} \frac{1}{1 + e^{\beta - (h^* + 2)x^*}} = x^* &\Rightarrow \frac{1}{1 + e^{\beta - 1}} \approx x^* \\ &\Rightarrow \frac{1}{1 + e^{\beta - 1}} \approx \frac{1}{\beta(h^* + 2)} \end{aligned}$$

Rearranging terms yields

$$h^*(\beta) \approx \frac{e^{\beta - 1} + 1}{\beta} - 2 \tag{14}$$

This approximation is accurate for large β . However, for small β the term $(x^*)^2$ in (12) is not negligible and the approximation breaks down. Moreover, the expression in (14) is negative when evaluated at $\beta = 2$; recall that $h^*(2) = 0$.

The next claim shows that if the expression in (14) is adjusted to equal zero at $\beta = 2$, it becomes a *global* upper bound for $h^*(\beta)$.

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 selection.

Note that $h(\beta)$ differs from the expression in (14) only by a constant ($4 - e \approx 1.28$) in the numerator, and that $h(2) = h^*(2) = 0$. The proof of Claim 2 is deferred to Appendix A.

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

We now show that when $\beta \geq 3$ and $\alpha > h(\beta)$ the process $\Gamma(\alpha, \beta)$ exhibits fast selection to 99%. We claim that the high equilibrium is increasing in both α and β . Indeed, identity (*) in footnote 13 implies that $\partial f / \partial \alpha$ 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 $\alpha = h(3) = (e^2 + 4 - e)/3 - 2 > 89\%$, the high equilibrium is above 99%. An explicit calculation shows that

$$f(0.99; h(3), 3) > \frac{1}{1 + e^{3(1-(0.89+2) \cdot (0.99))}} \approx 0.9963 > 0.99$$

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 $\Gamma(\alpha, \beta)$ displays fast selection, 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 selection, then $\Gamma_N(\alpha, \beta)$ also exhibits fast selection. More precisely, for any $\delta > 0$ we have

$$\lim_{N \rightarrow \infty} ET_N(\alpha, \beta, x_H - \delta) = T(\alpha, \beta, x_H - \delta) \tag{15}$$

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

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

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:

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

The proof is relegated to Appendix A.

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 (15). Consider a small $\epsilon > 0$, take $\mu = \frac{\epsilon}{2}$ and denote $T_\epsilon = T(x_H - \delta + \epsilon)$. Lemma 3 implies

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




Fig. 5. Waiting times to within $\delta = 1\%$ of the high equilibrium; 40 (red, solid), 20 (orange, dots) and 10 (green, dash-dot) revisions per capita. The x-axis is labeled with the initial error rate ε ; the corresponding noise level β appears in parentheses.

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 $E(T_N) < \frac{T}{q}$.¹⁴ 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 $\limsup_{N \rightarrow \infty} ET_N(x_H - \delta) \leq T_\epsilon$ for any $\epsilon > 0$, hence $\limsup_{N \rightarrow \infty} ET_N(x_H - \delta) \leq T$ by taking $\epsilon \rightarrow 0$. A similar argument shows that $\liminf_{N \rightarrow \infty} ET_N(x_H - \delta) \geq T$. \square

This concludes the proof of Theorem 1. \square

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. Fig. 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 $h^*(\beta)$, such that for $\alpha > h^*(\beta)$ selection is fast. Pairs (ε, α) on the red, solid line correspond to waiting times $T(\alpha, \beta, x_H - \delta)$ of 40 revisions per capita, while on the orange, dotted and green, dash-dotted lines the corresponding times are 20 and 10 revisions per capita, respectively.

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 $\alpha > h^*(\beta)$ 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

$$E(T_N) \leq \frac{S}{\sqrt{\frac{\alpha+2}{h^*(\beta)+2} - 1}} + \log(\delta^{-1}) + \delta \tag{16}$$

¹⁴ 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 $\dot{y} = f(y) - y$ 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

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




Fig. 6. An informal illustration of the evolution of the process.




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

To understand Theorem 2, note that as the payoff gain α approaches the threshold $h^*(\beta)$, a “bottleneck” appears for intermediate adoption rates, which slows down the process. Fig. 6 illustrates this phenomenon by highlighting the distance between the updating function f and the identity function (recall that the speed of the change of the process at adoption rate x is given by $f(x) - x$). The first term on the right hand side of inequality (16) tends to infinity as α tends to $h^*(\beta)$, and the proof of Theorem 2 shows that inequality (16) holds for the following explicit value of the constant S :

$$S_0 = \frac{4\sqrt{2}}{\sqrt{1 - \frac{4}{\beta(h^*(\beta)+2)}}} \tag{17}$$

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 (16) take care of this possibility.

Fig. 7 plots, in log–log format, the expected time to get within $\delta = 1\%$ of the high equilibrium as a function of $\alpha - h^*$, for $\varepsilon = 1\%$ (left panel) and $\varepsilon = 5\%$ (right panel). These initial error rates 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 $T(\alpha, \beta, x_H - \delta)$. Note that in both panels the two lines are parallel for small values of $\alpha - h^*$; this shows that the rate of convergence of the waiting time presented in Theorem 2 (inverse square root) is the “correct” rate.

Finally, we give a concrete example of the estimated waiting time. As above, fix the initial error rate at $\varepsilon = 5\%$ and the precision level at $\delta = 1\%$. Then the critical payoff gain is $h^*(\beta) \approx 74\%$. Our estimate for the waiting time when $\alpha = 100\%$ is $T = 30.5$, while the expected waiting time for $N = 1000$ is $T_{1000} = 17.5$.

Proof of Theorem 2. We shall show that Theorem 2 holds for the deterministic approximation $\Gamma(\alpha, \beta)$. It then follows, using Lemma 2, that Theorem 2 holds for all sufficiently large N given α, β and δ . The proof involves studying the behavior of the waiting time $T = T(\alpha, \beta, x_H - \delta)$ at the extremes, namely for α close to $h^* = h^*(\beta)$ and for α large. The case of intermediate α is covered by an appropriate choice of the constant S . Detailed proofs are relegated to Appendix B.




Fig. 8. The function $v(x, \alpha)$ (red, solid) is a lower bound for $f(x; \alpha)$ (green, dots).

The essential ideas are as follows. When α approaches h^* from above, a bottleneck forms close to the position of the low equilibrium x_L^* . The process slows down as it approaches the bottleneck, and the time it takes to traverse it increases as the bottleneck width $g(\alpha)$ decreases. (The term $g(\alpha)$ will be defined formally in the proof of the next lemma.) The length of the bottleneck is controlled by the curvature of the response function at the (tangential) low equilibrium x_L^* . The following lemma shows that as α approaches h^* the waiting time grows at a rate equal to the inverse square root of the width and length of the bottleneck. (For convenience, we drop the dependence of f on β , which is held fixed in the proof.)

Lemma 4. For any $\beta > 2$ we have

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

The idea of the proof of this lemma is depicted in Fig. 8. The waiting time is asymptotically the same when the response function $f(x; \alpha, \beta)$ is replaced by the function $v(x; \alpha, \beta)$, which is the upper envelope of the functions $f(x; h^*, \beta)$ and $x + g(\alpha)$. The bottleneck is the interval $[x_1, x_2]$ on which $v(x; \alpha, \beta) = x + g(\alpha)$. It can be shown that the time to approach the bottleneck, the time to traverse it, and the time to get away from it all grow at the same rate $(f''(x_L^*; h^*)g(\alpha))^{-1/2}$.

Our next result gives a linear estimate for the gap $g(\alpha)$ for α close to h^* .

Lemma 5.

$$\liminf_{\alpha \searrow h^*} \left(\frac{h^* + 2}{\alpha - h^*} \right) f''(x_L^*; h^*)g(\alpha) \geq 1 - \frac{4}{\beta(h^* + 2)}$$

Combining the results in Lemmas 4 and 5 we conclude that the behavior of the waiting time as α approaches h^* is given by:

$$\limsup_{\alpha \searrow h^*} T(\alpha, \beta, x_H - \delta) \sqrt{\frac{\alpha + 2}{h^* + 2} - 1} \leq \frac{4\sqrt{2}}{\sqrt{1 - \frac{4}{\beta(h^* + 2)}}} \tag{18}$$

Note that the right hand side of (18) is exactly S_0 as defined in (17).

When α is large, the process slows down when it approaches the high equilibrium. The behavior in this case is controlled by the precision level δ , as described in the following lemma.

Lemma 6. There exists α_0 such that for all $\alpha > \alpha_0$

$$T(\alpha; \beta, x_H - \delta) < \log(\delta^{-1}) + \frac{\delta}{2}$$

Theorem 2 follows by putting together these results. Note that $T(\alpha, \beta, x_H - \delta)$ is continuous in α , and Eq. (18) implies that there exists a constant $S > 0$ such that for all $\alpha \in (h^*, \alpha_0]$

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

Together with Lemma 6, we find that for any $\alpha > h^*$

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

The exact result in Theorem 2 follows by applying Lemma 2. \square

6. Partial information

We now turn to the case where agents have a limited, finite capacity to gather information. Assume that each player samples d other players before revising, where $d \geq 3$ is independent of N . As we shall see, 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 h^* 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 selection 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 $T_N(\alpha, \beta, d, p)$ and so forth.

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

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

Proof. 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. Detailed proofs are found in Appendix C; here we shall outline the main steps in the argument.

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 $f(\frac{k}{d})$. Hence the agent chooses action A with probability

$$f_d(x) = \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} f\left(\frac{k}{d}\right) \tag{19}$$

The definition of the deterministic approximation of the stochastic process $\Gamma_N(\alpha, \beta, d)$ is analogous to 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.

The following two lemmas state that the shape of the response function and the selection property established for full information continue to hold with partial information.

Lemma 7. The function f_d is first strictly convex then strictly concave, and the inflection point is at most $\frac{1}{2}$.

Lemma 8. 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 $[0, \frac{1}{2}]$.

The definition of the critical payoff gain $h^*(\beta, d)$ is analogous to the definition of $h^*(\beta)$ in the proof of Theorem 1. With the convention $\inf \emptyset = \infty$, let

$$h^*(\beta, d) = \inf\{\alpha: \Gamma(\alpha, \beta, d) \text{ displays fast selection}\}$$




Fig. 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, solid), 20 (orange, dots) and 10 (green, dash-dot) revisions per capita. The x-axis is labeled with the initial error rate ε ; the corresponding noise level β appears in parentheses.

We claim that f_d is strictly increasing in α , which implies that if $\alpha' > \alpha$ and $f_d(\cdot; \alpha, \beta)$ does not have any equilibria in the interval $[0, \frac{1}{2}]$, then neither does $f_d(\cdot; \alpha', \beta)$. It then follows that $\Gamma(\alpha, \beta, d)$ displays fast selection for all $\alpha > h^*(\beta, d)$.

To establish the claim, we differentiate expression (19) with respect to α and obtain

$$\frac{\partial f_d}{\partial \alpha}(x; \alpha, \beta) = \sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{\partial f}{\partial \alpha} \left(\frac{k}{d}; \alpha, \beta \right) \tag{20}$$

Using the expression for $\frac{\partial f}{\partial \alpha}$ available in (*) in footnote 13 we find that $\frac{\partial f_d}{\partial \alpha}(x; \alpha, \beta)$ is positive for all $x, \beta > 0$, which means that f_d is increasing in α .

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

Lemma 9. *The threshold $h^*(\beta, d)$ verifies $h^*(\beta, d) < h(\beta)$ and $h^*(\beta, d) < d - 2$.*

The intuition for these results is as follows. Identity (19) shows that f_d is a convex combination of the function f evaluated at points k/d , for $0 \leq k \leq d$. For low adoption rates, most of the weight falls on points k/d in the range where f is convex, and thus f_d lies above f . It follows that the low and middle equilibria disappear at lower values of α in the finite sampling case. For the second result, note that when $\alpha > d - 2$ it is a best response to choose A when the sample contains at least one adopter, and this holds for any value of β . This implies that $f_d(x) > x$ for $x \leq 1/2$, hence the process $\Gamma(\alpha, \beta, d)$ exhibits fast selection.

Theorem 3 now follows directly. \square

Partial information also leads to fast selection in an absolute sense. Fig. 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 $h^*(\beta, d)$, such that for any $\alpha > h^*(\beta, d)$ selection is fast. Pairs (ε, α) on the red, solid line correspond to waiting times $T_N(\alpha, \beta, d, x_H - \delta)$ of 40 revisions per capita, while on the orange, dotted and green, dash-dotted lines the corresponding waiting times are 20 and 10 revisions per capita, respectively. For example, we see that if the sample size is $d = 10$, the initial error rate is $\varepsilon = 5\%$, and the payoff gain is $\alpha = 60\%$, then the waiting time is below 20 revisions per capita.

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

Theorem 4. *For any $d \geq 3$, for any noise level β satisfying $h^*(\beta, d) > 0$ and for any precision level $\delta > 0$, there exists a constant $S = S(\beta, \delta, d)$ such that for all $\alpha > h^*(\beta, d)$ and all sufficiently large N given α, β and δ , 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*

$$E(T_N) \leq \frac{S}{\sqrt{\frac{\alpha+2}{h^*(\beta, d)+2} - 1}} + \log(\delta^{-1} \varepsilon^{-\frac{1}{d-1}}) \tag{21}$$

The intuition behind Theorem 4 is the same as for Theorem 2. As the payoff gain α approaches the threshold $h^*(\beta, d)$, a “bottleneck” appears for intermediate adoption rates, which slows down the process. This effect is captured by the first




Fig. 10. Waiting time until the process is within $\delta = 1\%$ of the high equilibrium, as a function of the payoff gain difference $\alpha - h^*$. Both axes have logarithmic scales. Simulation (blue, dashed) and upper bound (red, solid). Information parameters and initial error rates, and values for constant S : $d = 4, \epsilon = 1\%, S = 6.3$ (upper left panel), $d = 4, \epsilon = 5\%, S = 8.6$ (upper right panel), $d = 10, \epsilon = 1\%, S = 4.9$ (lower left panel), $d = 10, \epsilon = 5\%, S = 5.9$ (lower right panel).

term on the right hand side of (21). 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 (21) takes care of this aspect.

Fig. 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 $\alpha - h^*$. The initial error rates are $\epsilon = 1\%$ (left panels) and $\epsilon = 5\%$ (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 indicated values for the constant S . These figures show that Theorem 4 captures the “correct” rate of convergence of the waiting time as α tends to $h^*(\beta, d)$.

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 $h^* = h^*(\beta, d)$ 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 h^* the gap is approximately linear in the payoff gain difference $(\alpha - h^*)$. 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 (21). 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 (21) 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 (21).

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

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

Plugging this into identity (19) we find that for all $x \in [0, 1]$

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} f_d(x; \alpha, \beta) &= (1-x)^d \epsilon + \sum_{k=1}^d \binom{d}{k} x^k (1-x)^{d-k} \\ &= 1 - (1-\epsilon)(1-x)^d \end{aligned} \tag{22}$$

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 (22) 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

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

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

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

We conclude that

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

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 (21) holds for all $\alpha > h^*(\beta, d)$. \square

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 selection 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 selection 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 selection 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 α 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.

Fig. 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 selection, 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 at the high equilibrium. Just as in the case of the logit, one can estimate the critical payoff threshold that leads to fast selection for each given level of noise.

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 could be applied to larger games that have more than two equilibria. However, recent work by Hommes and Ochea (2012) shows that larger games may exhibit more complex dynamics and hence precise estimates of the waiting times may be more difficult to attain.




Fig. 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).

Acknowledgments

We thank Sam Bowles and Bill Sandholm for helpful comments on an earlier draft. This research was sponsored by the Office of Naval Research, grant N00014-09-1-0751.

Appendix A. Proofs in Theorem 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 selection.

Proof. 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

$$f'(x) = \beta(\alpha + 2)f(x)(1 - f(x)) = 1 \tag{23}$$

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 (23), and denote $C = (\beta(\alpha + 2))^{-1}$. Eq. (23) can be rewritten as

$$X - X^2 = C \tag{24}$$

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

$$C^{-1} = e^{\beta-1} + 4 - e \tag{25}$$

By assumption $\beta > 2$, hence $C \in (0, \frac{1}{4})$. Solving Eq. (24) we obtain

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

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

$$\frac{1 - X}{X} \geq (C^{-1} - 4 + e)e^{1-\frac{X}{C}} \tag{26}$$




Fig. 12. The function $u(M)$.

Denote $M = \frac{X}{1-X} \in (0, 1)$, hence also $X^{-1} = \frac{1+M}{M}$. Eq. (24) 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 (26) 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 Fig. 12, is first increasing and then decreasing, and it is strictly positive on the interior of the interval $(0, 1)$. \square

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. 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

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

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

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

Hence the difference between the deterministic and the stochastic processes is

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

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 $\Psi(T) = \max_{0 \leq t \leq T} |\int_0^t U(s) ds|$. 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 $\Pr(\Psi(T) e^{\lambda T} \geq \frac{\mu}{2})$. The following lemma will be useful.

Lemma 10. Let \mathcal{F}_k denote the σ -algebra generated by $\{x_N(t) : t \leq k\tau\}$. For any $\theta \in \mathbb{R}$

$$E(e^{\theta U_k} | \mathcal{F}_k) \leq e^{2|\theta|^2} \tag{27}$$

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 $g(1) \leq 2|\theta|^2$. \square

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 10, 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:

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

(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 $c = \frac{e^{-\lambda T}}{32T}$. \square

Appendix B. Proofs in Theorem 2

Lemma 4. For any $\beta > 2$ we have

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

Proof. We begin by defining the *breadth of the bottleneck* or *gap*. For α in a neighborhood of h^* let

$$g(\alpha) = \min_{x \in [0, x_H - \delta]} f(x; \alpha) - x$$

Note that for $\beta > 2$ the process $\Gamma(h^*, \beta)$ has exactly two equilibria $x_L^* < \frac{1}{2} < x_H^*$. Moreover, x_L^* is a tangential equilibrium, namely $f'(x_L^*; h^*) = 1$ and $f(x; h^*) \geq x$ for $x \in [0, x_H^* - \delta]$. It follows that $g(h^*) = 0$, and by the implicit function theorem there exists a function $t(\alpha)$ defined in a neighborhood of h^* such that $f'(t(\alpha); \alpha) = 1$ and

$$g(\alpha) = f(t(\alpha); \alpha) - t(\alpha) \tag{28}$$

To estimate the hitting time $T(\alpha, \beta, x_H - \delta)$ we shall construct a lower bound for $f(x; \alpha)$ that is easier to work with. For α in a neighborhood of h^* let

$$v(x; \alpha) = \max\{x + g(\alpha), f(x; h^*)\}$$

Recall that f is increasing in α , hence $v(x; \alpha) \leq f(x; \alpha)$ for all $x \in [0, \frac{1}{2}]$. Let $x_1(\alpha) < x_2(\alpha)$ denote the two solutions of the equation $g(\alpha) = f(x; h^*) - x$ in a neighborhood of x_L^* . Fig. 8 illustrates the construction of the function v and points x_1 and x_2 .

We now introduce a second order approximation of $f(x; h^*) - x$ around x_L^* . This allows us to estimate $x_1(\alpha)$ and $x_2(\alpha)$, as well as to approximate the waiting time to get close to the bottleneck. For x in a neighborhood of x_L^* we have

$$f(x; h^*) - x \approx (f(x_L^*; h^*) - x_L^*) + (f'(x_L^*; h^*) - 1)(x - x_L^*) + \frac{f''(x_L^*; h^*)}{2}(x - x_L^*)^2$$

The first two terms are zero because x_L^* is a tangential equilibrium. To ease notation, denote $\Omega = f''(x_L^*; h^*, \beta)$. We obtain

$$f(x; h^*) - x \approx \frac{\Omega}{2}(x - x_L^*)^2 \tag{29}$$

Solving $g(\alpha) = f(x; h^*, \beta) - x$ yields approximate solutions (for α close to h^*)

$$|x_L^* - x_i(\alpha)| \approx \sqrt{\frac{2g(\alpha)}{\Omega}} \tag{30}$$

We now decompose the waiting time

$$T(\alpha, \beta, x_H - \delta) = T\left(0 \rightarrow \frac{1}{2}\right) + T\left(\frac{1}{2} \rightarrow x_H - \delta\right) \tag{31}$$

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 $\alpha = h^*$ 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 $\alpha \geq h^*$.

We now turn to an estimate of the term $T(0 \rightarrow \frac{1}{2})$. 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 $x_0, x_1 \in [0, \frac{1}{2}]$. The waiting time $T_V(0 \rightarrow \frac{1}{2})$ can be further decomposed as

$$T_V\left(0 \rightarrow \frac{1}{2}\right) = T_V\left(0 \rightarrow x_1(\alpha)\right) + T_V\left(x_1(\alpha) \rightarrow x_2(\alpha)\right) + T_V\left(x_2(\alpha) \rightarrow \frac{1}{2}\right) \tag{32}$$

Using (30), for α close to h^* the middle term is approximately

$$T_V\left(x_1(\alpha) \rightarrow x_2(\alpha)\right) = \frac{x_2(\alpha) - x_1(\alpha)}{g(\alpha)} \approx \frac{2}{g(\alpha)} \sqrt{\frac{2g(\alpha)}{\Omega}} = \frac{2\sqrt{2}}{\sqrt{\Omega g(\alpha)}} \tag{33}$$

We now look at $T_V(0 \rightarrow x_1(\alpha))$. For α close to h^* this waiting time is controlled by the shape of $f(x; h^*)$ for x close to x_L^* . In particular, consider α close to h^* and y close to $x_1(\alpha)$, $y < x_1(\alpha)$. Using (29) we obtain

$$\begin{aligned} \dot{y} &= f(y; h^*, \beta) - y \\ &\approx \frac{\Omega}{2}(x_L^* - y)^2 \end{aligned}$$

Consider the process $z(t)$ with dynamics given by

$$\dot{z} = \frac{\Omega}{2}(x_L^* - z)^2 \quad \text{and} \quad \text{initial condition } z(0) = 0$$

Both waiting times $T_V(0 \rightarrow x_1(\alpha))$ and $T_z(0 \rightarrow x_1(\alpha))$ diverge as α approaches h^* ; because $y(t)$ and $z(t)$ follow asymptotically identical differential equations close to $x_1(\alpha)$ as α approaches h^* , it follows that the ratio of the two waiting times tends to 1.

The process $z(t)$ is given by

$$z(t) = x_L^* - \frac{2}{t\Omega}$$

Solving for T in $z(T) = x_1(\alpha)$ and using approximation (30) we obtain

$$T_V(0 \rightarrow x_1(\alpha)) \approx T = \frac{2}{\Omega(x_L^* - x_1(\alpha))} \approx \frac{\sqrt{2}}{\sqrt{\Omega g(\alpha)}} \tag{34}$$

By a similar argument we obtain

$$T_V\left(x_2(\alpha) \rightarrow \frac{1}{2}\right) \approx \frac{\sqrt{2}}{\sqrt{\Omega g(\alpha)}}$$

Combining inequalities (32), (33) and (34) yields

$$T_V\left(0 \rightarrow \frac{1}{2}\right) \approx \frac{4\sqrt{2}}{\sqrt{\Omega g(\alpha)}}$$

Inequality (31) gives

$$T(\alpha, \beta, x_H - \delta) < T_V\left(0 \rightarrow \frac{1}{2}\right) + K_1$$

where K_1 is independent of α . Multiplying both sides by $\sqrt{\Omega g(\alpha)}$ and taking the limit $\alpha \searrow h^*$ and replacing $\Omega = f''(x_L^*; h^*, \beta)$, we obtain

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

Lemma 5.

$$\liminf_{\alpha \searrow h^*} \left(\frac{h^* + 2}{\alpha - h^*}\right) f''(x_L^*; h^*) g(\alpha) \geq 1 - \frac{4}{\beta(h^* + 2)}$$

Proof. Using (7) we obtain

$$f''(x_L^*; h^*) = \beta(h^* + 2) f'(x_L^*; h^*) (1 - 2f(x_L^*; h^*))$$

Given that $f(x_L^*; h^*) = x_L^*$ and $f'(x_L^*; h^*) = 1$ this becomes

$$f''(x_L^*; h^*) = \beta(h^* + 2)(1 - 2x_L^*) \tag{35}$$

Turning to $g(\alpha)$, Taylor’s expansion around h^* implies that $g(\alpha) \approx g'(h^*)(\alpha - h^*)$, in the sense that $\liminf_{\alpha \searrow h^*} g(\alpha)(\alpha - h^*)^{-1} = g'(h^*)$.

Differentiating the definition of the gap in Eq. (28) with respect to α , we obtain that

$$\begin{aligned} g'(\alpha) &= f'(t(\alpha); \alpha) t'(\alpha) + \frac{\partial f}{\partial \alpha}(t(\alpha); \alpha) - t'(\alpha) \\ &= \frac{\partial f}{\partial \alpha}(t(\alpha); \alpha) \end{aligned}$$

Recall that in general

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

Noting that $f'(x_L^*; h^*) = 1$ we obtain $g'(h^*) = \frac{x_L^*}{h^* + 2}$. Hence, for α close to h^*

$$g(\alpha) \approx \frac{(\alpha - h^*)x_L^*}{h^* + 2} \quad (36)$$

Putting (35) and (36) together yields

$$\begin{aligned} \left(\frac{h^* + 2}{\alpha - h^*}\right) f''(x_L^*; h^*) g(\alpha) &\approx \beta(h^* + 2)(1 - 2x_L^*)x_L^* \\ &= \frac{x_L^* - 2(x_L^*)^2}{C} \end{aligned}$$

where we have used the notation $C \equiv \frac{1}{(h^* + 2)\beta}$. In the proof of Theorem 1 it is shown that $x_L^* - (x_L^*)^2 = C$ (see Eq. (24)). Thus we can write

$$\left(\frac{h^* + 2}{\alpha - h^*}\right) f''(x_L^*; h^*) g(\alpha) \approx 1 - \frac{(x_L^*)^2}{C}$$

It is easy to check that $x_L^* < 2C$ hence $\frac{(x_L^*)^2}{C} \leq 4C$. It follows that

$$\liminf_{\alpha \searrow h^*} \left(\frac{h^* + 2}{\alpha - h^*}\right) f''(x_L^*; h^*) g(\alpha) \geq 1 - \frac{4}{\beta(h^* + 2)} \quad \square$$

Lemma 6. *There exists α_0 such that for all $\alpha > \alpha_0$*

$$T(\alpha; \beta, x_H - \delta) < \log(\delta^{-1}) + \frac{\delta}{2}$$

Proof. 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 $x \in (0, 1]$. This convergence is not uniform, because $f(0; \alpha, \beta) = \varepsilon$ for all α , so a more detailed argument is necessary. 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 $\lim_{\mu \rightarrow 0} y(t) = x_\infty(t)$.

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

$$\lim_{\alpha \rightarrow \infty} T(\alpha, \beta, x_H - \delta) = \log(\delta^{-1}) \quad (37)$$

The statement of Lemma 6 follows easily. \square

Appendix C. Proofs in Theorem 3

The following two lemmas are useful throughout the proof.

Lemma 11. 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

$$1 \geq e^{\beta(1-(2+\alpha)x)} e^{\beta(1-(2+\alpha)(1-x))} = e^{-\beta\alpha} \tag{38}$$

Note that by assumption $\beta > 0$, so inequality (38) is true for all $\alpha \geq 0$, and equality occurs if and only if $\alpha = 0$. \square

Lemma 12. For any $\alpha > 0$ we have $f_d(\frac{1}{2}) > \frac{1}{2}$.

Proof. Using Lemma 11, we have that for any k

$$\begin{aligned} \binom{d}{k} f\left(\frac{k}{d}\right) + \binom{d}{d-k} f\left(1 - \frac{k}{d}\right) &= \binom{d}{k} \left(f\left(\frac{k}{d}\right) + f\left(1 - \frac{k}{d}\right) \right) \\ &> \binom{d}{k} \\ &= \binom{d}{k} \frac{1}{2} + \binom{d}{d-k} \frac{1}{2} \end{aligned} \tag{39}$$

Adding up inequality (39) for $k = 1, \dots, \lfloor \frac{d}{2} \rfloor$ we obtain¹⁵

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

Lemma 7. The function f_d is first strictly convex then strictly concave, and the inflection point is at most $\frac{1}{2}$.

Proof. 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 $f_d''(\frac{1}{2}) \leq 0$, which implies that inflection point of f_d is at most $\frac{1}{2}$.

We begin by introducing some notation. Let $\Delta_d f(k)$ be the discrete derivative of f with step $1/d$, evaluated at $x = \frac{k}{d}$, namely

$$\Delta_d f(k) = \frac{f\left(\frac{k+1}{d}\right) - f\left(\frac{k}{d}\right)}{1/d} \tag{40}$$

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

$$\Delta_d f(k_l) \geq \frac{f(x_l) - f(\frac{k_l}{d})}{x_l - \frac{k_l}{d}} > \Delta_d f(k_l - 1)$$

The last inequality follows because f is strictly convex on $[0, x_l]$.

¹⁵ For even values of d we multiply the inequality corresponding to $k = \frac{d}{2}$ by one-half to avoid double counting.

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 $k_0 \in \{k_l - 1, k_l\}$ 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

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

We now express the second derivative of f_d in terms of s_k and h_k . First, by differentiating (19) and using the notation in (40) the derivative of f_d can be written as

$$f'_d(x) = \sum_{k=0}^{d-1} \binom{d-1}{k} x^k (1-x)^{d-1-k} \Delta_d f(k) \tag{41}$$

Differentiating and rearranging terms we obtain

$$\begin{aligned} f''_d(x) &= \sum_{k=0}^{d-2} \binom{d-2}{k} x^k (1-x)^{d-2-k} \frac{\Delta_d f(k+1) - \Delta_d f(k)}{1/(d-1)} \\ &= \sum_{k=0}^{d-2} s_k h_k(x) \end{aligned} \tag{42}$$

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 13. *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

$$s_k h_k(y) < \lambda s_k h_k(x) \tag{43}$$

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

$$s_k h_k(y) < \lambda s_k h_k(x) \tag{44}$$

Finally, for $k = k_0$ we have

$$s_k h_k(y) = \lambda s_k h_k(x) \tag{45}$$

Adding up expressions (43), (44) and (45) for $k = 0, 1, \dots, d - 2$ we obtain that $f''_d(y) < \lambda f''_d(x) \leq 0$. \square

We shall now show that $f''_d(\frac{1}{2}) \leq 0$ by direct computation. Rearranging the terms in (42) 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

$$\sum_{k=0}^d ((2k-d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq 0 \tag{46}$$

In outline, the remainder of the proof runs as follows. We shall first establish that Q is increasing on the interval $[0, \frac{1}{2}]$. Next, we shall show that as k increases from 0 to $\lfloor \frac{d}{2} \rfloor$ the coefficient $((2k - d)^2 - d)$ is first positive and then negative. These two facts imply that in order to prove inequality (46), it is sufficient to prove the inequality after dropping the Q term. The last part of the proof will establish inequality (46) after dropping the Q term.

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

To establish this claim, fix $x \in [0, \frac{1}{2}]$ 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 11.

Note that as k increases from 0 to $\lfloor \frac{d}{2} \rfloor$ 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 $Q(\frac{k}{d}) \geq Q(\frac{k^*}{d})$. Hence

$$((2k - d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq ((2k - d)^2 - d) \binom{d}{k} Q\left(\frac{k^*}{d}\right) \tag{47}$$

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

$$((2k - d)^2 - d) \binom{d}{k} Q\left(\frac{k}{d}\right) \leq ((2k - d)^2 - d) \binom{d}{k} Q\left(\frac{k^*}{d}\right) \tag{48}$$

From inequalities (47) and (48) 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

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

This follows from the identities

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

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

Lemma 8. 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 $[0, \frac{1}{2}]$.

Proof. Lemma 12 showed that $f_d(\frac{1}{2}) > \frac{1}{2}$. Since f is continuous and always less than 1, there exists a down crossing point in the interval $(\frac{1}{2}, 1)$. This corresponds to a high equilibrium. Furthermore, the strict concavity of f_d on $(\frac{1}{2}, 1)$, established in Lemma 7, guarantees that the high equilibrium is unique.

By Lemmas 7 and 12 f_d is convex and then concave on $[0, \frac{1}{2}]$, and $f_d(\frac{1}{2}) > \frac{1}{2}$. 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 $f_d(x_I) \geq x_I$ there are at most two equilibria in the interval $[0, x_I]$, because f_d is strictly 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 $f_d(x_I) < x_I$ there exists exactly one equilibrium in the interval $[0, x_I]$, and also exactly one equilibrium in the interval $(x_I, \frac{1}{2}]$. \square

Lemma 9. The threshold $h^*(\beta, d)$ verifies $h^*(\beta, d) < h(\beta)$ and $h^*(\beta, d) < d - 2$.

Proof. To prove that $h^*(\beta, d) \leq h^*(\beta)$, we show that if for given $\alpha > 0$ and β the process $\Gamma(\alpha, \beta)$ exhibits fast selection, 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 selection $m \leq \frac{d}{2}$. By definition of m we have that

$$f\left(\frac{k}{d}\right) > \frac{k}{d} \quad \text{for any } k \text{ such that } m < k < d - m \quad (49)$$

We also use the identity

$$\sum_{k=0}^d \binom{d}{k} x^k (1-x)^{d-k} \frac{k}{d} = x \quad (50)$$

Fix x in the interval $[0, \frac{1}{2}]$, and rewrite expression (19) as

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

The second line follows from identity (50). Using inequality (49) for all k such that $m < k < d - m$ we obtain

$$f_d(x) \geq x + \sum_{\substack{0 \leq k \leq m, \\ d-m \leq k \leq d}} \binom{d}{k} x^k (1-x)^{d-k} \left(f\left(\frac{k}{d}\right) - \frac{k}{d} \right)$$

Rearranging terms in the summation yields:

$$f_d(x) \geq x + \sum_{k=0}^m \binom{d}{k} x^{d-k} (1-x)^k \left[\left(\frac{1-x}{x} \right)^{d-2k} \left(f\left(\frac{k}{d}\right) - \frac{k}{d} \right) + \left(f\left(\frac{d-k}{d}\right) - \frac{d-k}{d} \right) \right] \quad (51)$$

Fix $k \in \{0, 1, \dots, m\}$. Then $(\frac{1-x}{x})^{d-2k} \geq 1$, and $f(\frac{k}{d}) - \frac{k}{d}$ 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 11.

Using (51) we have now established that $f_d(x) > x$ for all $x \leq \frac{1}{2}$. It follows that $h^*(\beta, d) \leq h^*(\beta)$. In particular, $h^*(\beta, d) \leq h(\beta)$.

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

Note that $f(\frac{1}{d}) \geq f(\frac{1}{\alpha+2}) = \frac{1}{2}$. 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

$$f\left(\frac{k}{d}\right) \geq \frac{k}{d} \quad \text{for all } 0 \leq k \leq \frac{d}{2} \quad (52)$$

Lemma 14. For any $x \in (0, \frac{1}{2}]$ and $0 \leq k \leq \frac{d}{2}$ we have

$$x^k (1-x)^{d-k} f\left(\frac{k}{d}\right) + x^{d-k} (1-x)^k f\left(\frac{d-k}{d}\right) > x^k (1-x)^{d-k} \frac{k}{d} + x^{d-k} (1-x)^k \left(1 - \frac{k}{d}\right) \quad (53)$$

Proof. Denote $V_k = x^k(1-x)^{d-k}$ and $W_k = x^{d-k}(1-x)^k$. The assumptions in the lemma imply that $V_k \geq W_k$. Using inequality (52) we obtain that

$$(V_k - W_k) f\left(\frac{k}{d}\right) \geq (V_k - W_k) \frac{k}{d} \quad (54)$$

Note that $W_k > 0$, so Lemma 11 implies that

$$W_k \left(f\left(\frac{k}{d}\right) + f\left(\frac{d-k}{d}\right) \right) > W_k \left(\frac{k}{d} + 1 - \frac{k}{d} \right) \quad (55)$$

Lemma 14 follows by adding up inequalities (54) and (55). \square

Weighing inequality (53) 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 Lemma 9. \square

References

- Anderson, R.M., May, R., 1991. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press.
- Anderson, S., De Palma, A., Thisse, J., 1992. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
- Benaïm, M., 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 Econ. Behav.* 5 (3), 387–424.
- Blume, L.E., 1995. The statistical mechanics of best-response strategy revision. *Games Econ. Behav.* 11 (2), 111–145.
- Blume, L.E., 2003. How noise matters. *Games Econ. Behav.* 44 (2), 251–271.
- Blume, L.E., Durlauf, S., 2002. Equilibrium concepts for social interaction models. Working paper 7. Department of Economics, University of Wisconsin–Madison.
- Brock, W.A., Durlauf, S.N., 2001. Interactions-based models. In: Heckman, J.J., Leamer, E.E. (Eds.), *Handbook of Econometrics*, vol. 5, pp. 3297–3380 (Chapter 54).
- Ellison, G., 1993. Learning, local interaction, and coordination. *Econometrica* 61 (5), 1047–1071.
- Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* 38, 219–232.
- Freidlin, M., Wentzell, A., 1984. *Random Perturbations of Dynamical Systems*. Springer-Verlag, Berlin.
- Hommes, C.H., Ochea, M.I., 2012. Multiple equilibria and limit cycles in evolutionary games with logit dynamics. *Games Econ. Behav.* 74 (1), 434–441.
- Jackson, M.O., Yariv, L., 2007. Diffusion of behavior and equilibrium properties in network games. *Amer. Econ. Rev.* 97 (2), 92–98.
- Kandori, M., Mailath, G.J., 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. *J. Appl. Probab.* 7, 49–58.
- Lelarge, M., 2012. Diffusion and cascading behavior in random networks. *Games Econ. Behav.* 75 (2), 752–775.
- López-Pintado, D., 2006. Contagion and coordination in random networks. *Int. J. Game Theory* 34 (3), 371–381.
- McFadden, D.L., 1976. Quantal choice analysis: a survey. In: *NBER Chapters. Ann. Econ. Soc. Meas.* 5 (4), 8–66.
- McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. *Games Econ. Behav.* 10, 6–38.
- Montanari, A., Saberi, A., 2010. The spread of innovations in social networks. *Proc. Natl. Acad. Sci.* 107 (47), 20196–20201.
- Morris, S., 2000. Contagion. *Rev. Econ. Stud.* 67 (1), 57–78.
- Oyama, D., Sandholm, W., Tercieux, O., 2011. Sampling best response dynamics and deterministic equilibrium selection. Mimeo, October 2011.
- Sandholm, W.H., 2001. Almost global convergence to p-dominant equilibrium. *Int. J. Game Theory* 30, 107–116.
- Sandholm, W.H., 2010. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA.
- Shah, D., Shin, J., 2010. Dynamics in congestion games. In: *Proceedings of the ACM SIGMETRICS Conference*.
- Watts, D.J., 2002. A simple model of global cascades on random networks. *Proc. Natl. Acad. Sci.* 99 (9), 5766–5771.
- Watts, D.J., Dodds, 2007. Influentials, networks, and public opinion formation. *J. Cons. Res.* 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 University Press, Princeton, NJ.
- Young, H.P., 2011. The dynamics of social innovation. *Proc. Natl. Acad. Sci.* 108 (Suppl. 4), 21285–21291.