

Evolution with Mutations Driven by Control Costs

ERIC VAN DAMME* AND JÖRGEN W. WEIBULL†

November 11, 1998

ABSTRACT. Bergin and Lipman (1996) show that the refinement effect from the random mutations in the adaptive dynamics in Kandori, Mailath and Rob (1993) and Young (1993) is due to restrictions on how these mutation rates vary across population states. We here model these mutation rates as endogenously determined mistake probabilities, by assuming that players at some cost or disutility can control their mistake probability, i.e., the probability of implementing another pure strategy than intended. This is shown to corroborate the result in Kandori-Mailath-Rob and Young that the risk-dominant equilibrium is selected in 2×2 -coordination games. Doc: *mute7.tex*.

1. INTRODUCTION

It has been shown by Kandori, Mailath and Rob (1993), henceforth "KMR," and Young (1993), henceforth "Young," that adding small noise to certain adaptive dynamics in games can lead to rejection of strict Nash equilibria. Specifically, in 2×2 coordination games these dynamics allow one to conclude that in the long run the risk-dominant equilibrium (Harsanyi and Selten, 1988) will result. This surprisingly strong result has recently been challenged by Bergin and Lipman (1996) who show that it depends on specific assumptions about the mutation process, namely that the mutation rate does not vary "too much" across the different states of the adaptive process. They show that, if mutation rates at different states are not taken to zero at the same rate, then many different outcomes are possible. Indeed, by choosing the mutation rates appropriately, any desired strict equilibrium may be selected in the long run.

Bergin and Lipman conclude from this lack of robustness that the nature of the mutation process must be scrutinized more carefully if one is to derive economically meaningful predictions, and they offer two suggestions for doing so. As suggested already in the original KMR and Young papers, there are two ways of interpreting the mutations in these models: Mutations may be thought of as arising from individuals' experiments or from their mistakes. In the first case, it is natural to expect the mutation rate to depend on the state - individuals may be expected to experiment less in states with higher payoffs. Also in the second case state-dependent mutation rates appear reasonable - exploring an idea proposed in Myerson (1978) one might argue that mistakes associated with larger payoff losses are less likely.

*Center for Economic Research, Tilburg University.

†Stockholm School of Economics and the Research Institute of Industrial Economics, Stockholm. The authors thank Maria Saez-Marti for helpful comments.

While Bergin and Lipman are right to point out that these considerations might lead to state-dependent mutation rates, they do not elaborate or formalize these ideas. Hence, it is not clear whether their concerns really matter for the conclusions drawn by KMR and Young. The aim of this study is to shed light on this issue by means of a straight-forward model of mutations as mistakes. The model is based on the assumption that players "rationally choose to make mistakes" because it is too costly to avoid them completely. It turns out that our model, when applied to 2×2 coordination games, produces mistake probabilities that do not vary "too much" with the state of the system. Hence, the concerns of Bergin and Lipman are irrelevant in this case.

Specifically, we build on the control-cost model of van Damme (1987, chapter 4). In essence, individuals are assumed to have a trembling hand, and by controlling it more carefully, which involves some disutility, the amount of trembles can be reduced. Since a rational player will try harder to avoid more serious mistakes, i.e. mistakes that lead to larger payoff losses, such mistakes will be less likely. However, they will still occur with positive probability since, by assumption, it is infinitely costly to avoid mistakes completely. Although mistake probabilities thus depend on the associated payoff losses, and therefore also on the state of the process, we show that, under mild regularity conditions, they all go to zero at the same rate when control costs become vanishingly small in comparison with the payoffs in the game. Consequently, the techniques developed by KMR and Young can be employed to show that, in 2×2 -games, only the risk-dominant equilibrium survives in the long run.

We do not know of any earlier work along the above lines. Blume (1993) studies strategic interaction between individuals who are located on a lattice and recurrently interact with their neighbors. These individuals play stochastically perturbed myopic best responses, in the sense that the choice probability for each pure strategy is an increasing positive function of its current payoff. As a consequence, more costly mistakes are assigned lower choice probabilities. Blume (1994) elaborates and extends this model to a more conventional random-matching setting in which choice probability ratios are an increasing function of the associated payoff differences. Maruta (1997) generalizes Blume's models by letting choice probability ratios be a function of both payoffs - not necessarily only of their difference. While these studies take random choice behavior as a starting point for the analysis, we here derive such behavior from an explicit decision-theoretic model in which individuals take account of their own mistake probabilities. Robles (1998) extends the KMR and Young models by letting mutation rates decline to zero over time, in one part of the study also allowing for state-dependent mutation rates. Also his work is complementary to ours in the sense that while he takes the state-dependence for given and analyzes implications thereof, we suggest a model that explains why and how mutation rates vary across population states.

The remainder of the text is organized as follows. In Section 2 we outline the models of KMR and Young. Section 3 discusses the Bergin-Lipman critique of these models and argues for the need to determine the probability of mutations/mistakes

endogenously. Section 4 provides such a model, in which mistakes arise out of rational deliberation. Section 5 analyzes this model, and proves that it selects the risk-dominant equilibrium in 2×2 -coordination games. Section 6 discusses a counterexample in the context of Young's model, and section 7 concludes.

2. ADAPTIVE DYNAMICS

Let an n -person game $G = \langle S_1, \dots, S_n, u_1, \dots, u_n \rangle$ be given and, for each player position $i = 1, \dots, n$, let C_i be a finite population of individuals. In the KMR and Young models, the game is played recurrently between individuals drawn from these populations. The so drawn individuals have some information about the past play of the game. Based on this information, the individuals form beliefs about how their opponents will play and choose a best response to these beliefs.¹ The chosen actions add to the history of play for the next round, and so on.

We restrict most of the subsequent analysis to 2×2 -coordination games with two strict equilibria, the reason being that Young gets his strongest results for this class, and that KMR restrict the main part of their analysis to symmetric such games. Formally, we then consider two-player games with payoff bi-matrices

$$\begin{array}{|c|c|c|}
 \hline
 & A & B \\
 \hline
 A & a_1, a_2 & b_1, b_2 \\
 \hline
 B & c_1, c_2 & d_1, d_2 \\
 \hline
 \end{array} \tag{1}$$

where $a_1 > c_1, a_2 > b_2, d_1 > b_1$ and $d_2 > c_2$. Hence, the pure Nash equilibria are strict and sit on the main diagonal of the bi-matrix.

In Young's model, a pair of individuals, one from population C_1 and one from population C_2 , are randomly drawn in each period to play the game. The state θ of the system in his model is a full description of the pure-strategy profiles played in the last m such rounds. Each individual drawn to play in position i of the game is assumed to make a statistically independent sample of k of these m profiles, and plays a best reply to the opponent population's empirical frequency of actions (pure strategies) in the sample.

KMR restrict their analysis to symmetric games, i.e. where $a_2 = a_1, b_2 = c_1, c_2 = b_1$ and $d_2 = d_1$. One may then represent the game by the payoff matrix of player 1,

$$\begin{array}{|c|c|c|}
 \hline
 & A & B \\
 \hline
 A & a & b \\
 \hline
 B & c & d \\
 \hline
 \end{array}, \tag{2}$$

where $a > c$ and $d > b$. Moreover, KMR assume that there is only one population, $C_1 = C_2$, and they assume that in every period each individual in this population

¹KMR allow individuals more freedom; they should just move, as an aggregate, to better responses than in the last round. Allowing for this here would complicate the notation without adding more insight.

plays against all other individuals. The state θ of the system is defined as the number of individuals who played the first action (pure strategy A) in the last round. All individuals are assumed to play a best reply to the current state.²

To obtain selection of an equilibrium, noise is added to these models. Basically, an individual in player position i plays a best response with probability $1 - \lambda_i \varepsilon$. A mistake or experiment occurs with positive probability $\lambda_i \varepsilon$, with mistakes (experiments) being statistically independent across time, states and individuals. Hence, all mistake probabilities are positive, and the ratios between mistake probabilities across states and player positions are constant as $\varepsilon \rightarrow 0$.

With these mistakes as part of the process, each state of the system is reachable with positive probability from every other state. Hence, the full process is an irreducible Markov chain on the finite state space Θ , where $\Theta \subset (S_1 \times S_2)^m$ in Young's model and $\Theta = \{0, \dots, |C_1|\}$ in the KMR-model. Consequently, there exists a unique stationary distribution $\mu \varepsilon$ for each $\varepsilon > 0$. KMR and Young establish the existence of the limit distribution $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu \varepsilon$, and study its properties. They call an equilibrium of the underlying game G *stochastically stable* if μ^* places positive probability weight on the state in which this equilibrium is played (in the last period in KMR, in the last m periods in Young). The main result obtained is that, for generic 2×2 -games with two strict equilibria, the risk-dominant equilibrium is the unique stochastically stable equilibrium.

Risk dominance (Harsanyi and Selten, 1988) in a game with payoff bi-matrix (1) may be characterized as follows. Let p_i be the probability that player i assigns to his first pure strategy, A , in the unique mixed-strategy equilibrium of the game. Then (A, A) , i.e. each player choosing his first action, is the risk-dominant equilibrium if and only if

$$p_1 + p_2 < 1 \tag{3}$$

If the reversed strict inequality holds, then (B, B) - each player choosing his second action - is the risk-dominant equilibrium.

In the symmetric version (2) of the game, we have $p_1 = p_2$, and condition (3) is equivalent to the condition that player 1 would strictly prefer pure strategy A if his opponent were to play both pure strategies with the same probability. Hence, (A, A) is the risk dominant equilibrium of the symmetric game (2) if and only if

$$a + b > c + d . \tag{4}$$

The intuition for the main result can most easily be seen in the KMR-model. If (4) is satisfied and if the system is in state "A", i.e., all individuals in the population played pure strategy A in the last round, then the state is upset only if more than half the population simultaneously make a mistake (experiment). In contrast, state "B" is upset if less than half the population simultaneously make a mistake. Since, in

²See footnote 1.

the limit, the second possibility is infinitely more likely than the first, the process will spend (virtually) all time in state "A", i.e. at the risk-dominant equilibrium (A, A) .³

3. ROBUSTNESS AND ENDOGENOUS MUTATION RATES

The reasoning in the preceding paragraph relies on the assumption that the mistake/mutation probability is the same in all states. More generally, KMR and Young derive their results under the assumption that the ratio between any pair of mutation probabilities is kept constant as $\varepsilon \rightarrow 0$. Bergin and Lipman (1996) note that their results continue to hold even with state-dependent mutation probabilities if the ratio between any pair of mutation probabilities, across all population states and player positions, has a non-zero limit when $\varepsilon \rightarrow 0$. However, Bergin and Lipman (1996) also show that if the mutation probabilities in different states are allowed to go to zero at different rates, then any stationary distribution in the mutation-free process can be turned into the unique limiting distribution $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu\varepsilon$ of the process with mutations.

For example, they show that the limiting distribution μ^* in the KMR model places unit probability on the risk-dominated equilibrium (B, B) in the game (2) with payoffs $a = 6, b = 4, c = 0$ and $d = 8$, if the mutation rate in state "B" is $\varepsilon\beta$, and the mutation rate in state "A" is ε , for any $\beta > 1.5$.⁴ As ε goes to zero, mutations in state "B" become infinitely rarer than mutations in "A". Consequently, it is more difficult for the population state to get out of the basin of attraction of "B", although "A" has a larger basin of attraction (in the sense of containing more population states). They conclude: "In other words, any refinement effect from adding mutations is solely due to restrictions of how mutation rates vary across states." (Bergin and Lipman, p. 944).

What is missing in these analyses, and this is Bergin's and Lipman's main message, is a theory of why and how mutations occur. For instance, their counter-example cannot be discarded if mutations indeed are (at least half) an order of magnitude rarer in state "B" than in state "A". One reason why this may be the case, suggested by Bergin and Lipman, is that mutation rates might be lower in high-payoff states than in low-payoff states, which might be expected if mutations are due to individuals' experimentation (see Bergin and Lipman, 1996, pp. 944, 945 and 947). Another reason why mutation probabilities may differ across population states, also suggested by Bergin and Lipman (pp. 945 and 955), is that mutations leading to larger payoff losses might have lower probabilities than mutations leading to smaller payoff losses, for reasons similar to those given in Myerson's (1978) motivation of the concept of proper equilibrium. Bergin and Lipman do not investigate the consequences of either of these two ideas. We now follow up on their latter suggestion and show that it leads to a confirmation of the earlier results of KMR and Young.

³For a complete analysis one also needs to consider mutations in intermediate states.

⁴More exactly, the best-reply correspondence divides the state space into two basins of attraction, one for state "A" and one for state "B". Bergin and Lipman assume that the mutation rate is $\varepsilon\beta$ in all population states in B's basin, and ε in all population states in A's basin.

Indeed, in the case of a symmetric game, it is straightforward to see why such a result should come about in the KMR model. Namely, condition (4) is equivalent to $a - c > d - b$, which says that mistakes at (A, A) are more serious - involve larger payoff losses - than mistakes at (B, B) . Hence, any "reasonable" theory of endogenous mistakes should imply that mistakes at (A, A) are less likely than mistakes at (B, B) . In other words, the basin of attraction of state "A" should not only be "larger" than that of "B", it should also be "deeper", thus making it even more difficult to upset this equilibrium.⁵ In the next section we formally demonstrate this result for symmetric games played by a single population, in a model where mistakes arise from control-cost considerations. However, we develop the model in a more general setting that includes asymmetric games played by two populations. In such settings, the intuitive argument above is not available: In the risk-dominant equilibrium one of the two deviation losses may be quite small, thus inducing relatively large mistake probabilities in one of the player positions at that equilibrium.

4. ADAPTIVE DYNAMICS WITH CONTROL COSTS

Van Damme (1987, chapter 4) develops a model where mistakes arise in implementing pure strategies in games. The basic idea is that players make mistakes because it is too costly to prevent these completely. Each player has a trembling hand, and by controlling it more carefully (which involves higher costs) the amount of trembles can be reduced. It is assumed to be infinitely costly to eliminate trembles completely.

For a general n -person normal-form game, with pure strategy sets S_1, \dots, S_n , mixed-strategy profiles $\sigma = (\sigma_1, \dots, \sigma_n)$, and payoffs $\pi_i(\sigma)$, these ideas are formalized as follows. If σ is played, player i 's payoff is not just $\pi_i(\sigma)$; in addition to this "regular" payoff, i incurs costs (disutility) to control his trembling hand and to actually implement σ_i . As a result, the player's payoff is given by $u_i(\sigma) = \pi_i(\sigma) - \delta v_i(\sigma_i)$, where $\delta > 0$ is a scaling parameter that measures the importance of control costs (disutility of control effort) relative to the original payoffs, and $v_i(\sigma_i)$ is the cost or disutility player i incurs in order to implement the mixed strategy $\sigma_i \in \Delta(S_i)$. The function $v_i : \text{int}[\Delta(S_i)] \rightarrow \mathbb{R}_+$ is called the *control-cost function*. This function is assumed to be strictly convex, symmetric, and twice differentiable with $\lim_{\sigma_i^k \rightarrow 0} v_i(\sigma_i) = +\infty$ for every pure strategy $k \in S_i$. Symmetry here means that $v_i(\tau_i) = v_i(\sigma_i)$ for every mixed strategy σ_i and permutation τ_i of the components of σ_i .⁶ These assumptions imply that it is costly to keep down the probability that a given pure strategy is played, that the marginal cost of reducing this probability is increasing, and that the cost of reducing such a probability to zero is prohibitive.

We now apply these ideas to the KMR and Young models. Let θ denote the current state of the system, as described above, and let $\iota_i(\theta)$ be the information that an individual drawn to play in player position i has about this state. In Young's

⁵See footnote 3.

⁶Note that symmetry implies that v_i achieves its minimum value when all pure strategies are assigned the same probability. In van Damme (1987), the cost functions v_i are given in the symmetric additive form $v_i(\sigma_i) = \sum_{k \in S_i} f_i(\sigma_i^k)$, for f_i strictly convex and twice differentiable.

model, $\iota_i(\theta) \in (S_1 \times S_2)^k$, while in the KMR model $\iota_1(\theta) = \iota_2(\theta) = \theta \in \{0, \dots, |C_1|\}$. In both models this information determines a probabilistic belief for each of the two individuals about the action to be taken by the opponent. Since the state space Θ is finite, the set Ω of possible beliefs in these models is finite (see section 2). Let $\pi_i(\sigma_i, \omega)$ denote the expected payoff to an individual who plays mixed strategy σ_i in player position $i = 1, 2$, when his belief is ω .⁷ (We write $\pi_i(A, \omega)$ when he plays pure strategy A etc.) In the unperturbed process, each individual chooses a mixed strategy $\sigma_i \in \Delta(S_i)$ in order to maximize $\pi_i(\sigma_i, \omega)$. Taking control costs into account, each individual chooses $\sigma_i \in \text{int}[\Delta(S_i)]$ to maximize

$$u_i(\sigma_i, \omega) = \pi_i(\sigma_i, \omega) - \delta v_i(\sigma_i) . \quad (5)$$

The resulting stochastic process is ergodic and has a unique stationary distribution $\mu\delta$. The limiting case $\delta = 0$ represents a situation in which all individuals are perfectly rational in the sense of being able to perfectly control their actions at no cost or effort. We are interested in the limit as $\delta \rightarrow 0$, i.e. when control costs become insignificant in comparison with the payoffs in the game.

5. THE RESULT

Focusing on 2×2 -games with payoff bi-matrices (1), we first introduce some simplifying notation. For any belief $\omega \in \Omega$ of an individual in player position $i = 1, 2$, let

$$b_i(\omega) = \max \{ \pi_i(A, \omega), \pi_i(B, \omega) \} , \quad w_i(\omega) = \min \{ \pi_i(A, \omega), \pi_i(B, \omega) \} , \quad (6)$$

and $l_i(\omega) = b_i(\omega) - w_i(\omega)$. Given the finiteness of Ω , we may, without loss of generality, assume that the payoff loss $l_i(\omega)$ in case of a mistake is positive for both player positions and all beliefs.⁸

Now consider the game with payoff functions specified by (5). We assume that the control costs are the same in both player positions, and simplify notation by writing $v(p)$ for $v_i(p, 1-p)$ ($= v_i(1-p, p)$ by symmetry). Writing ε for the probability with which an individual in player position i by mistake plays the non-optimal action, we can write the game as a game on the open unit square, where the individual in each player position i chooses a mistake probability $\varepsilon \in (0, 1)$ in order to maximize his or her expected payoff, $b_i(\omega) - \varepsilon l_i(\omega) - \delta v(\varepsilon)$. The optimal mistake probability at ω , $\varepsilon_i(\delta, \omega)$, is thus determined by the first-order condition

$$-v'(\varepsilon) = l_i(\omega)/\delta . \quad (7)$$

Here $-v'(\varepsilon)$ is the marginal cost of reducing the mistake probability from ε . Our assumptions imply that the function $-v'$ is decreasing on $(0, \frac{1}{2})$ from plus infinity to

⁷In the KMR model, ω is one of the finitely many mixed strategies that correspond to a population distribution over the set of pure strategies in the last round. In Young's model, ω is one of the finitely many empirical frequency distributions over the opponent player position's set of pure strategies that correspond to a sample of size k from the opponent population's past m plays.

⁸More precisely, this is true for any finite population sizes in the Young and KMR models, for generic payoffs (1) or (2), respectively.

zero. Figures 1 and 2 show the graphs of the control-cost function $v(p) = -\log p - \log(1-p)$ and the associated marginal-cost function.

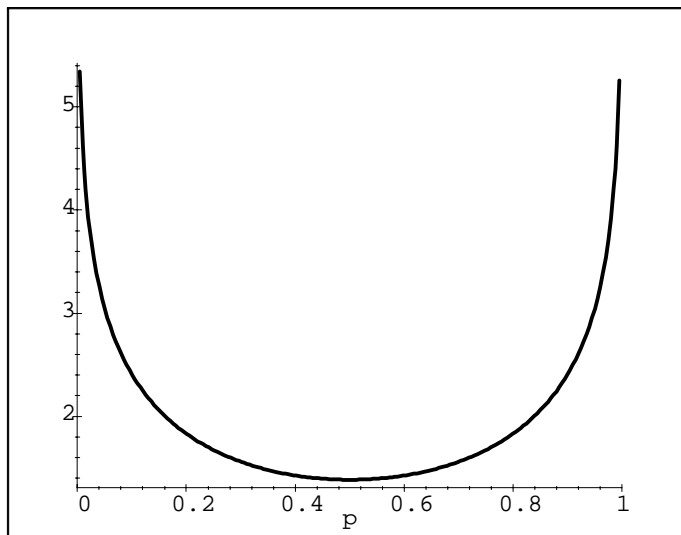


Figure 1: The control-cost function $v(p) = -\log [p(1-p)]$.

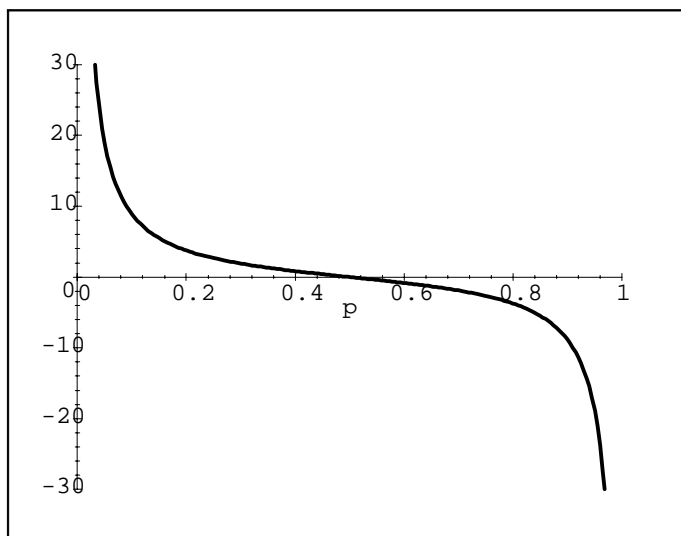


Figure 2: The marginal-cost function $-v'$ corresponding to the control-cost function v in Figure 1.

Since the marginal-cost function $-v'$ is decreasing, larger potential payoff losses are accompanied by lower optimal mistake probabilities. For all beliefs ω and ω' :

$$l_i(\omega) < l_i(\omega') \quad \Rightarrow \quad \varepsilon_i(\delta, \omega') < \varepsilon_i(\delta, \omega) . \quad (8)$$

It follows that our model selects the risk-dominant equilibrium in the KMR-setting. For, as already noted, equilibrium (A, A) being risk-dominant in game (2) is

equivalent to a mistake in state "A" resulting in a smaller payoff loss than a mistake in state "B". More precisely, let ω_A^p denote the belief that one's opponent will play pure strategy A with probability p , and let ω_B^p denote the belief that one's opponent will play pure strategy B with probability p . Such beliefs correspond in the KMR model to situations where the population share p played pure strategy A (B) in the last round. Thus ω_A^p corresponds to a population state at "distance" $1 - p$ from state "A", and ω_B^p to a population state at "distance" $1 - p$ from state "B". For each p sufficiently close to one, the payoff loss from a mistake at the first state is larger than at the second state. Hence, by (8), the optimal mistake probability at the first state is smaller than at the second. In fact:

$$p > \frac{a - c}{a - c + d - b} \quad \Rightarrow \quad l(\omega_A^p) > l(\omega_B^p) \quad \Rightarrow \quad \varepsilon_i(\delta, \omega_A^p) < \varepsilon_i(\delta, \omega_B^p), \quad (9)$$

showing that the basin of attraction of state "A" is not only "larger" but also (point-wise) "deeper" than that of "B".

Using this observation and following the steps of the proof in KMR one establishes:

Proposition 1. *The risk-dominant equilibrium is the unique stochastically stable equilibrium in the KMR-model.*

Below we establish a similar result for Young's model. We do this by way of identifying a mild regularity condition under which the mistake probabilities at all information states are of the same order of magnitude when δ is taken to zero in that model. For this purpose, let $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the inverse of $-v'$. Before stating the regularity condition, however, we note that for no control-cost function v and scalar $\lambda > 1$ is it the case that the ratio $\varphi(\lambda x)/\varphi(x)$ converges to zero as x goes to plus infinity:

Proposition 2. $\limsup_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x) > 0$ for all $\lambda > 1$.

Proof: Suppose the contrary, i.e., $\limsup_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x) = 0$ for some $\lambda > 1$. Then $\lim_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x) = 0$ for this λ , since φ is positive. Thus $\varphi(\lambda x)/\varphi(x) \leq \frac{1}{2}\lambda^{-1}$ for all x sufficiently large, say $x > x_0$. Hence,

$$\int_{x_0}^{\infty} \varphi(x) dx \leq (\lambda - 1)x_0\varphi(x_0) \sum_{t=0}^{\infty} 2^{-t} < \infty,$$

and thus $\int_0^{\infty} \varphi(x) dx < \infty$ since φ is decreasing with $\varphi(0) = \frac{1}{2}$. But we also have

$$\int_0^{\infty} \varphi(x) dx = - \int_0^{\frac{1}{2}} v'(p) dp = \lim_{p \rightarrow 0} v(p) - v\left(\frac{1}{2}\right),$$

which contradicts the assumption that $\lim_{p \rightarrow 0} v(p) = +\infty$. \square

We define v to be *nice* if $\liminf_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x) > 0$ for some $\lambda > 1$. It follows from proposition 2 that a control-cost function is nice if $\lim_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x)$ exists for some $\lambda > 1$. Moreover, it is easily seen that if v is nice then $\liminf_{x \rightarrow \infty} \varphi(\eta x)/\varphi(x) > 0$ for all $\eta \geq 1$.⁹ An example of a nice control-cost function is the one used in Figures 1 and 2. That function v is clearly strictly convex, symmetric and twice differentiable with $\lim_{p \rightarrow 0} v(p) = +\infty$, and it is easily verified that $\lim_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x) = 1/\lambda$ for any $\lambda > 0$.¹⁰ See Figures 3 and 4.

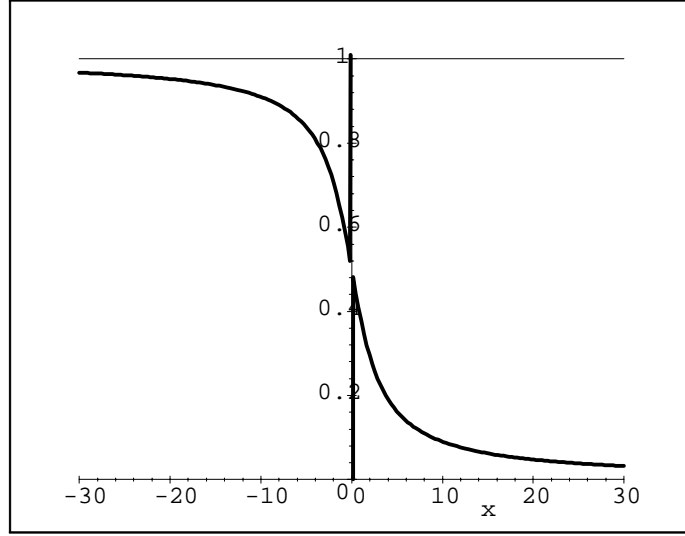


Figure 3: The inverse φ of the marginal-cost function $-v'$ in Figure 2.

Our main result is that if the control-cost function v is nice, then all mistake probabilities are of the same order of magnitude at all states:

Proposition 3. *If the control-cost function v is nice and $l_i(\omega) > l_i(\omega')$, then*

$$0 < \liminf_{\delta \rightarrow 0} \varepsilon_i(\delta, \omega)/\varepsilon_i(\delta, \omega') \leq \limsup_{\delta \rightarrow 0} \varepsilon_i(\delta, \omega)/\varepsilon_i(\delta, \omega') \leq 1$$

Proof: Rewriting (7), we have $\varepsilon_i(\delta, \omega) = \varphi[l_i(\omega)/\delta]$, and the first inequality follows from the observation after the definition of niceness. The second inequality follows from (8). \square

⁹By choosing h such that $\lambda^h \geq \eta$, one obtains $\liminf_{x \rightarrow \infty} \varphi(\eta x)/\varphi(x) \geq (\liminf_{x \rightarrow \infty} \varphi(\lambda x)/\varphi(x))^h > 0$.

¹⁰Here $-v'(p) = 1/p - 1/(1-p)$, and hence $\varphi(x) = 1/2 + 1/x - \text{sign}(x)\sqrt{1/4 + 1/x^2}$ for $x \neq 0$, and $\varphi(0) = 1/2$. An application of l'Hopital's rule gives the sought limit ratio.

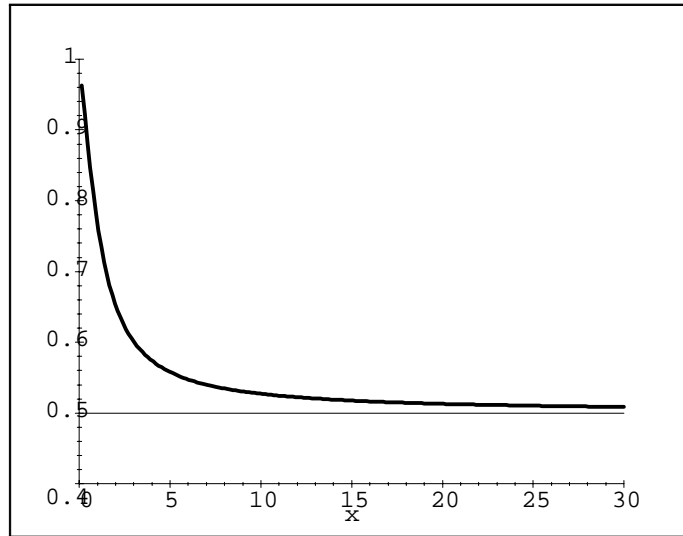


Figure 4: The ratio $\varphi(2x)/\varphi(x)$ for the marginal-cost function $-v'$ in Figure 2.

Corollary 1. *If the control-cost function v is nice, then the risk-dominant equilibrium is the unique stochastically stable equilibrium also in Young's model.*

Proof: This follows from proposition 3 combined with the arguments used in Young. While Young establishes the result formally only for the case where all mistakes are a positive multiple of ε (at each state θ player i makes a mistake with probability $\lambda_i\varepsilon$), it is easily verified that his arguments remain valid in the slightly more general case considered here, where the mistake probabilities are of the same order of magnitude in all states. The interested reader may consult Samuelson (1994, Theorem 4) for the more general case. \square

6. A COUNTER-EXAMPLE

Here we briefly turn to the case of control-cost functions that are non-nice. It is not difficult to see that such cost functions do exist.¹¹ For this purpose, define the sequence $\langle (x_n, y_n) \rangle_{n=1}^{\infty}$ in \mathbb{R}_+^2 by $x_1 = 1$, $y_1 = \frac{1}{3}$, and for all integers $n > 1$,

$$\begin{cases} x_n = 2x_{n-1} & \text{if } n \text{ is even} \\ y_n = y_{n-1}^a & \text{if } n \text{ is even} \\ x_n = 3x_{n-1} + 2/y_{n-1} & \text{if } n \text{ is odd} \\ y_n = y_{n-1}/2 & \text{if } n \text{ is odd} \end{cases} \quad (10)$$

where $a > 1$. Clearly $\langle x_n \rangle$ is an increasing sequence in \mathbb{R}_+ going to plus infinity, and $\langle y_n \rangle$ is a decreasing sequence in \mathbb{R}_+ going to zero. Hence, there exist differentiable (to any order) and strictly decreasing functions $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\varphi(0) = \frac{1}{2}$,

¹¹The following construction of functions φ is an adaptation of an example suggested to us by Henk Norde, Tilburg University.

and $\varphi(x_n) = y_n$ for all positive integers n . Any such function φ decreases sufficiently on intervals (x_n, x_{n+1}) with n odd in order for $\varphi(2x_n)/\varphi(x_n) \rightarrow 0$ to hold for n odd, as $n \rightarrow \infty$. On other intervals, however, φ decreases sufficiently little in order for its total integral to be infinite (the integral of φ over any interval (x_n, x_{n+1}) with n even is at least 1). The equation $v' = -\varphi^{-1}$ defines a control-cost function v that is non-nice.

What might happen in case of such a control-cost function is that a stochastically stable equilibrium need not exist in Young's model. Consider the following special case of the payoff bi-matrix (1):

	A	B	
A	5, 1	0, 0	(11)
B	0, 0	2, 2	

With p_i denoting the probability that player i assigns to his pure strategy A in the unique mixed equilibrium of this game, we have $p_1 = \frac{2}{3}$ and $p_2 = \frac{2}{7}$. Hence $p_1 + p_2 < 1$, so (A, A) is the risk-dominant equilibrium. However, a mistake by player 2 at this equilibrium incurs the smallest payoff loss in the two equilibria - this player loses only 1 payoff unit, while all other equilibrium payoff losses are at least 2 units. Consequently, the largest mistake probability occurs in state (A, A) .

Let φ be as constructed above, and let $\delta_n = 1/x_n$ for every odd positive integer n . By definition of φ ,

$$\varepsilon_2(\delta, A) = y_n, \quad \varepsilon_1(\delta, B) = \varepsilon_2(\delta, B) = y_n^a, \quad \text{and} \quad \varepsilon_1(\delta, A) = \mathcal{O}(y_n^a), \quad (12)$$

where $y_n \rightarrow 0$ as $n \rightarrow \infty$.¹² Hence, when individuals' memory size m is 1, then the equilibrium (A, A) is more easily upset than equilibrium (B, B) - because of 2's trembles in state (A, A) : State (A, A) is infinitely more easily upset than state (B, B) in the limit as $n \rightarrow \infty$, granted $a > 1$. More generally, for an arbitrary memory size m , state (A, A) is upset if population 1 makes $[m/3]_+$ mistakes, the probability for which is $\varepsilon_1(\delta, A)^{[m/3]_+}$, or if population 2 makes $[5m/7]_+$ mistakes, the probability for which is $\varepsilon_2(\delta, A)^{[5m/7]_+}$.¹³ Likewise, state (B, B) is upset if population 1 makes $[2m/3]_+$ mistakes, the probability for which is $\varepsilon_1(\delta, B)^{[2m/3]_+}$, or if population 2 makes $[2m/7]_+$ mistakes, the probability for which is $\varepsilon_2(\delta, B)^{[2m/7]_+}$. Using the observations in (12) one finds that, as $n \rightarrow \infty$, state (A, A) is infinitely more easily upset than state (B, B) , for any m , granted $a > 3$.¹⁴

We conclude that there exists a sequence of δ 's converging to zero such that the associated subsequence of stationary distributions in the limit places all probability

¹²To see this, note that if $\delta = 1/x_n$ then $\varepsilon_i(\delta, \omega) = \varphi[l_i(\omega)x_n]$. Hence, $\varepsilon_2(\delta, A) = \varphi(x_n) = y_n$, and $\varepsilon_1(\delta, B) = \varepsilon_2(\delta, B) = \varphi(2x_n) = \varphi(x_{n+1}) = y_n^a$. Moreover, $\varepsilon_1(\delta, A) = \varphi(5x_n) \geq \varphi(x_{n+2}) = y_{n+2} = y_n^a/2$, since φ is decreasing and $5x_n < x_{n+2} = 3x_{n+1} + 2/y_{n+1} > 6x_n$.

¹³Here $[z]_+$ denotes the smallest integer exceeding z .

¹⁴The condition is that a exceeds $f(m) = [5m/7]_+ / [2m/7]_+$. The function f achieves its maximum value, 3, at $m = 3$, and approaches $\frac{5}{2}$ as $m \rightarrow \infty$.

mass on state "(B, B)". At the same time, proposition 2 implies that there exists another sequence of δ 's such that the associated subsequence of stationary distributions places all probability mass on state "(A, A)" in the limit. Hence, the overall limit of stationary distributions does not exist for this control-cost function.

Another way of formulating our result thus is that, for every control-cost function, the risk-dominant equilibrium belongs to the limit *set* of supports of stationary distributions, and that this limit set is a singleton if the control-cost function is nice.¹⁵

7. CONCLUSION

Bergin and Lipman (1996) showed that if mutation rates are state dependent, then the long-run equilibrium depends on exactly how these rates do depend on the state. They conclude that the causes of mutations need to be modeled in order to derive justifiable restrictions on how mutation rates depend on the state. In particular, they suggest that one might investigate the consequences of letting the probability of mistakes be related to the payoff losses resulting from these mistakes. This is exactly what we have done in this paper. We have developed a model in which mistakes are endogenously determined, and shown that this model vindicates the original results obtained by Kandori, Mailath and Rob (1993) and Young (1993): the risk-dominant equilibrium is selected in the long run in all generic 2×2 -coordination games.

The model analyzed in this paper, although allowing for mistakes, is based on strong rationality assumptions. Mistakes arise because players choose to make them (since it is too costly to avoid them). Our individuals are hence unboundedly rational when it comes to decision making. Their lack of rationality is only procedural: At no cost or disutility can they choose their own mistake probabilities in every population state. This is a very strong rationality assumption. However, we believe that our conclusion is robust in this respect. For we obtain the same limit result as in the KMR and Young models, in which no individual takes any account of control costs. The effect of introducing control-costs was seen to only "deepen" the "basin of attraction" of the risk-dominant equilibrium, and hence speeding up the convergence to it. The limit result should therefore also be valid in intermediate cases of rationality.

REFERENCES

- [1] Bergin J. and B. Lipman (1996): "Evolution with state-dependent mutations", *Econometrica* 64, 943-956.
- [2] Blume L. (1993): "The statistical mechanics of strategic interaction", *Games and Economic Behavior* 5, 387-424.
- [3] Blume L. (1994): "How noise matters", mimeo., Cornell University.

¹⁵By the *limit set* of supports of stationary distributions, for a given control-cost function, we here mean the union of the supports of probability distributions μ , each of which is the limit to some subsequence of stationary distributions (i.e., associated with some sequence of positive δ 's converging to zero).

- [4] van Damme, E. (1987): *Stability and Perfection of Nash Equilibria*. Springer Verlag (Berlin).
- [5] Harsanyi J. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*. MIT Press (Cambridge, USA).
- [6] Kandori M., Mailath G. and R. Rob (1993): "Learning, mutation, and long run equilibria in games", *Econometrica* 61, 29-56.
- [7] Maruta T. (1997): "Binary games with state dependent stochastic choice", mimeo., Osaka Prefecture University.
- [8] Myerson R. (1978): "Refinements of the Nash equilibrium concept", *International Journal of Game Theory* 7, 73-80.
- [9] Robles J. (1998): "Evolution with changing mutation rates", *Journal of Economic Theory* 79, 207-223.
- [10] Samuelson L. (1994): "Stochastic stability in games with alternative best replies", *Journal of Economic Theory* 64, 35-65
- [11] Young P. (1993): "The evolution of conventions", *Econometrica* 61, 57-84.