

Viscosity and dispersion in an evolutionary model of learning*

Ivar Kolstad**

March 2002

Discussion Paper 07/02

Abstract

A two-population evolutionary model of learning is proposed where there is a conflict of interests between populations, and where interaction ranges from complete viscosity to complete dispersion. The long run stochastically stable states under complete viscosity match those of Kandori, Mailath and Rob (1993). With complete dispersion, the long run stochastically stable states match those of Hehenkamp (2001). With fluid interaction, there is a possibility of convention coexistence. Welfare properties of the long run stochastically stable states are examined using utilitarian and Rawlsian measures of welfare.

Keywords: Evolutionary game theory; Viscosity; Learning; Stochastic stability; Equilibrium selection; Convention coexistence

* Incisive and constructive comments from Bertil Tungodden, Oddvar Kaarbøe and Gaute Torsvik are gratefully acknowledged. For all remaining errors, the author is responsible.

** Chr. Michelsen Institute, Fantoftvegen 38, N-5892 Bergen, Norway. Tel: +47 55 57 42 39. E-mail: ivar.kolstad@cmi.no

Introduction

In a biological context, Hamilton (1964) defines viscosity as the tendency of individuals to have a higher rate of interaction with their close relatives than with more distantly related individuals. Myerson, Pollock and Swinkels (1991) formulate this idea in terms of a biological game, where an agent has a higher probability of meeting any agent sharing his strategy than any agent using a different strategy. Taking the limit as the degree of viscosity tends to zero, Myerson et al define a set of fluid population equilibria. Since the set of fluid population equilibria consists only of Nash equilibria, but not of all Nash equilibria, their model can be viewed as a contribution to the refinements literature. Moreover, since all evolutionarily stable strategies are contained as a subset in the set of fluid population equilibria, Myerson et al have also coined a concept of evolutionary stability which serves as an alternative to that of Maynard Smith and Price (1973).

Others have explored the notion of viscosity in ways more or less similar to that of Myerson et al. The idea of strategy correlation, that agents using the same strategies meet more frequently than agents using different strategies, has been explored by Frank (1988) for the prisoner's dilemma game and by Skyrms (1994, 1996) for a larger set of games. Models of local interaction, most notably those of Ellison (1993), Blume (1993, 1995) and Anderlini and Ianni (1996), capture a form of viscosity where agents have a fixed location and interact only with a limited set of neighbours. Oechssler (1997) suggests a model in which a population is divided into groups that interact only internally, but where agents can occasionally leave one group for another. Finally, viscosity is frequently used as a justification for introducing mutant clusters into evolutionary models, such as in Binmore and Samuelson (2001).

Whereas Myerson et al confine themselves to biological games, the aforementioned contributions forcefully underscore the potential importance of viscosity in human interaction. What is striking, however, is that none of these contributions mention the opposite possibility, that similar agents might in certain cases interact less frequently than dissimilar agents, a phenomenon we might term *dispersion*. To appreciate the importance of dispersive interaction, we need only think of interactions such as those of buyers and sellers, of principals and agents, of professors and students, of males

and females, and so on. While it is true that multi-population evolutionary models of learning appear to adopt dispersion as a matter of course, they always do so in the sense of full dispersion, where similar agents never interact. To name a few, Young (1993, 1998), Hahn (2000) and Hehenkamp (2001) all assume that the members of different groups take on distinct roles in the games played. This might prove a good fit for pure buyer-seller relationships, but once the buyers or sellers start interacting among themselves as well, we have a different kind of situation requiring a different kind of analysis. Current models thus capture the cases ranging from viscosity to fluidity, plus the extreme case of full dispersion. In order to attain "a framework general enough to accomodate all kinds of non-random pairing" (Skyrms, 1996), we ought therefore attempt to fill the gap between fluidity and full dispersion.

In this paper, I present an evolutionary model of learning which accomodates the full range of interaction of two distinct populations, from viscosity through fluidity through dispersion. The basic learning process is similar to that of Kandori, Mailath and Rob (1993), as elaborated on by Kandori and Rob (1995), Hahn (2000) and Hehenkamp (2001). These models basically either assume that members of a population only interact with each other (as in the former two contributions) or only interact with members of the other population (as in the latter two contributions). The gap between these two extremes is partly filled in one specific sense by the local interaction model of Ellison (1993), which employs a learning process similar to that of Kandori et al. The below model adopts a more flexible view of non-random interaction, and attempts to fill the entire gap between these specific models of learning.

The paper proceeds as follows. In the next section, the basic model is presented. Two populations of agents play a game of coordination, where agents from different populations prefer different equilibria. Every so often, agents are called upon to revise their strategies, choosing a best reply to the strategy profile of the preceding period. On rare occasions, agents choose a strategy at random. As the probability of such random choices approaches zero, we study the long run probabilities of different population states. States that have a positive probability of being observed in the very long run when noise is virtually absent, we call long run stochastically stable, adopting the term used by Ellison (2000). The three subsequent sections establish

long run stochastically stable states when interaction is dispersive, fluid and viscous, respectively. Interestingly, the results obtained with complete dispersion mirror those of Hehenkamp (2001). Similarly, the results obtained with complete viscosity are akin to those of Kandori, Mailath and Rob (1993). In a sense, then, the results of Hehenkamp and Kandori et al emerge as special cases in the below model. In the case of fluid interaction, I prove that a state where different populations use different strategies can be long run stochastically stable. This possibility of convention coexistence marks a departure from the results of previous models of learning with a similar mutation structure, and adds to the literature on coexistence initiated by Sugden (1995).¹ In a final section, I note that in the evolutionary literature, the debate on welfare properties of long run stochastically stable states has largely been limited to games of common interests, such as in Bergin and Lipman (1996). This section suggests that utilitarian and Rawlsian measures of welfare can be employed in models of conflicting interests, and reports some results on how the long run stochastically stable states fare when gauged by these measures.

The model

In its literal sense, the term viscous is used to describe a liquid that is thick or sticky, and thus hard to pour. Viscosity is thus an apt term for interaction where agents largely stick to a limited set of partners or opponents. By contrast, the term fluid describes a liquid that flows freely or easily. The analogy of fluid interaction thus implies that an agent interacts just as easily or frequently with one opponent as with another. To expand the dichotomous imagery used by biologists to describe interaction, add the term dispersion, which suggest that agents of the same type scatter to interact more frequently with agents of a different type.

An interpretation of the above three terms can be made within the confines of a two-population model. Consider two distinct populations 1 and 2 of finite sizes N_1 and

¹ Anderlini and Ianni (1996) assume that errors only occur when agents attempt to use a different strategy from that of the preceding period, which implies a non-ergodic dynamic process where in some cases coexistence is an absorbing state.

N_2 , respectively. The following figure provides a description of the three different modes of interaction.

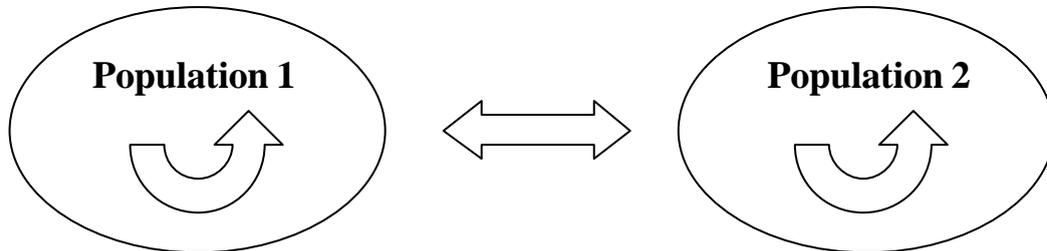


Figure 1. Structure of interaction

As the arrows indicate, the members of a population can interact with agents from their own population, and/or agents from the other population. If populations only interact internally, i.e. members from different populations never meet, interaction is completely viscous. In terms of evolutionary models of learning, Kandori, Mailath and Rob (1993) in essence adopt this assumption by studying single-population interaction. Conversely, if populations only interact externally, i.e. members of the same population never meet, interaction is completely dispersive. Young (1993) and Hehenkamp (2001) propose multi-population models of learning that exhibit this feature. If agents interact as often with any member from one population as from another, interaction is fluid, which is analogous to the definition suggested by Myerson, Pollock and Swinkels (1991).

Specifically, the notion of a round-robin tournament is used to describe interaction. In a round-robin tournament, agents are paired a number of times so that each agent meets each other agent exactly once. Tournaments of this kind are an easy way of having agents interact with the population average, which simplifies the modelling of strategy revision, as discussed below. In the current model, however, we want the frequency with which agents interact with members of their own population and members of the other population to vary. To this end, we imagine that an agent participates in a series of round-robin tournaments with his own and the other population. In each period, agents play r rounds of round-robin tournaments with their own population, and s rounds of tournaments with the other population. Each agent

thus interacts with the average of each population, but not necessarily with the average across populations.

The quotient $p = r/s$ captures the frequency with which agents interact with any member of their own population relative to any member of the opposite population, and p is thus a measure of the degree of viscosity (or dispersion) in interaction. Interaction is fluid if an agent plays an equal number of rounds with each population, i.e. if $p = 1$. If he plays more rounds with his own population, i.e. $p > 1$, interaction is viscous, where $p \rightarrow \infty$ implies complete viscosity. Fewer rounds played with your own population, $p < 1$, implies dispersive interaction, and complete dispersion as $p \rightarrow 0$. For $p \in \langle 0, \infty \rangle$, this formulation in principle allows the study of any form of interaction from completely dispersed through completely viscous.

Another way of modelling interaction that would also be amenable to the notion that agents interact with the population average, is to assume that all agents are paired once, and that the probability of meeting any agent from the same population is the same, though the probability of meeting agents from different populations may differ. The problem with this approach is that populations of different sizes would then exhibit different levels of viscosity, and the interaction of the larger population could never reach a level of full dispersion. To understand why, assume that $N_1 = 100$ and $N_2 = 50$. With complete dispersion, the probability of meeting a member of the opposite population is one, which means that if population 2 exhibits full dispersion, all fifty members of that population are paired with members of population 1. To add up, this must mean that 50 members of population 1 interact with members of population 2, which means that the probability of meeting a member of the opposite population is only $2/3$ for agents from population 1. Nor can that probability be raised above $2/3$, since there are no more potential agents from population 2 with whom agents from population 1 can be paired.

The chosen way of modelling matching also differs from that of Myerson et al. Their basic take on viscosity is to say that with probability \mathbf{b} , an agent gets an opponent from his own population, whereas with probability $(1 - \mathbf{b})$ his opponent is drawn at

random from the overall population, i.e. both his own and the other population. The main drawback to this approach is that it only allows the study of the cases ranging from fluid interaction ($b \rightarrow 0$) to completely viscous interaction ($b \rightarrow 1$). There is no natural way in which to expand this framework to the case of dispersive interaction. In an appendix, however, I show that for the range covered, matching according to Myerson et al yields results similar to those of the round-robin matching regime proposed above.

Given the round-robin matching regime, matched agents play a game with two strategies A and B. The game is essentially one of coordination, where a player prefers to use the same strategy as his opponent. However, we assume that the populations differ with respect to which pair of similar strategies is preferable, there is thus a conflict of interests between populations. Hence, regardless of the identity of his opponent, let an agent from population 1 receive payoffs according to the following matrix, where $a > 1$

		Opponent		
		A	B	
Player from population 1	A	a	0	(P1)
	B	0	1	

Similarly, the payoffs to an agent from population 2 can be represented as, for $b > 1$

		Opponent		
		A	B	
Player from population 2	A	1	0	(P2)
	B	0	b	

Thus, whenever two members of population 1 are matched, they play a coordination game, where they both prefer strategy profile (A,A).

		Player from population 1		
		A	B	
Player from population 1	A	a,a	$0,0$	(G1)
	B	$0,0$	$1,1$	

The coordination game is pure in the sense of Kandori and Rob (1995). Equilibrium (A,A) is thus both Pareto dominant and risk dominant.²

Similarly, when two members of population 2 meet, they play a pure coordination game where both prefer (B,B)

		Player from population 2		
		A	B	
Player from population 2	A	$1,1$	$0,0$	(G2)
	B	$0,0$	b,b	

In this game, equilibrium (B,B) is Pareto and risk dominant.

² Harsanyi and Selten (1988) define risk dominance in the following way. Consider any 2x2 game with two strict Nash equilibria U and V , where the losses to players 1 and 2 from unilaterally deviating from the equilibria are (u_1, u_2) and (v_1, v_2) , respectively. U risk dominates V if $u_1 \cdot u_2 > v_1 \cdot v_2$, and V risk dominates U if the opposite inequality holds.

Finally, when members of opposite population interact, they play a battle of the sexes game, where the agent from population 1 prefers profile (A,A) and the agent from population 2 prefers (B,B)

		Player from population 2		
		A	B	
Player from population 1	A	$a,1$	$0,0$	(G3)
	B	$0,0$	$1,b$	

In game (G3), no equilibrium Pareto-dominates the other. Without loss of generality, we assume that $a > b$, which makes equilibrium (A,A) risk dominant in this game.³ Population 1 thus has a stronger preference for its preferred strategy profile than the corresponding preference of population 2.

Denote by z_1^t the number of agents playing A in population 1 in period t , and let z_2^t represent the number of agents playing A in population 2 in period t . The vector $\mathbf{z}^t = (z_1^t, z_2^t)$ thus captures the state of the system at time t . The state space \mathbf{O} is discrete and finite

$$\mathbf{O} = \{\mathbf{z} = (z_1, z_2) : 0 \leq z_i \leq N_i, i = 1, 2\} \quad (1)$$

To ease subsequent discussion, let $\mathbf{z}^{AA} \equiv (N_1, N_2)$ represent the state in which all agents play strategy A, and let $\mathbf{z}^{BB} \equiv (0, 0)$ capture the state in which all play B. Similarly, in state $\mathbf{z}^{AB} \equiv (N_1, 0)$ all members of population 1 play A and all members of population 2 play B. Conversely, in state $\mathbf{z}^{BA} \equiv (0, N_2)$ B is played by everyone in population 1 and A is played by everyone in population 2.

³ Cardinal interpersonal comparability of payoffs must be assumed for this statement to be meaningful.

The state vector \mathbf{z}^t evolves as follows. In-between periods, each agent has a probability $\mathbf{d} \in \langle 0,1 \rangle$ of being called upon to revise his strategy. If called upon, an agent chooses the strategy which maximizes his expected payoffs in the next period, given a belief that all other agents will play as they did in the preceding period. We thus have a stochastic best-reply learning dynamic, where agents are myopic in only heeding the most recent actions of others. Moreover, the dynamic is what Hehenkamp (2001) defines as individualistic, since it leaves open the possibility that none or some or all agents in a population revise at any given time.

From payoff maximization, it follows that a revising agent from population 1 chooses strategy A if the relative frequency with which he expects to encounter A-players exceeds $\mathbf{a}_1 \equiv \frac{1}{1+a}$, where $\mathbf{a}_1 < 0.5$. An agent from population 1 therefore chooses A if

$$\frac{p \cdot z_1 + z_2}{p \cdot N_1 + N_2} > \mathbf{a}_1 \quad (2)$$

If the opposite relation holds, an agent from population 1 chooses B.

Similarly, a revising agent from population 2 chooses strategy A if the frequency with which he meets A-players is above $\mathbf{a}_2 \equiv \frac{b}{1+b}$, where $0.5 < \mathbf{a}_2 < 1 - \mathbf{a}_1$. A revising agent from population 2 will thus choose strategy A if

$$\frac{p \cdot z_2 + z_1}{p \cdot N_2 + N_1} > \mathbf{a}_2 \quad (3)$$

Conversely, B is chosen if the opposite relation holds. In the case where agents are indifferent between A and B, i.e. (2) or (3) hold with equality, we may assume a coin toss determines the strategy chosen.

Inequalities (2) and (3) capture the directions of change in the model. For ease of subsequent exposition, rewrite these two inequalities as:

$$z_2 > \mathbf{a}_1(pN_1 + N_2) - pz_1 \quad (4)$$

$$z_2 > \frac{\mathbf{a}_2(pN_2 + N_1)}{p} - \frac{z_1}{p} \quad (5)$$

By means of these inequalities, we can draw a stability diagram for the system. Figure 2 provides an illustration of such a diagram:

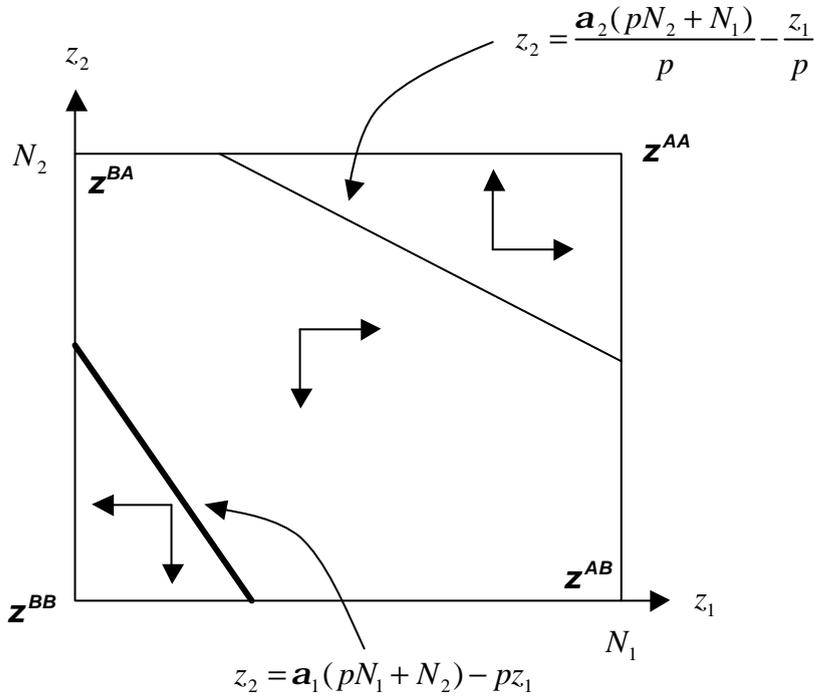


Figure 2. Illustration of stability diagram

The number of A-players in each population, z_1 and z_2 , is measured along the respective axes. This implies that \mathbf{z}^{BB} lies at the origin, and \mathbf{z}^{AA} at the upper right corner of the rectangle formed by the axes and population sizes. Similarly, \mathbf{z}^{AB} is at the lower right corner of the rectangle, and \mathbf{z}^{BA} at its upper left corner. The thicker of the two sloping lines represents the demarcation line between areas where z_1 increases (above the line) and decreases (below the line). The thinner of the two lines in the same manner demarcates the areas where z_2 increases (above) and decreases

(below). An absorbing state is a state which once reached, the process never leaves. In the case depicted in figure 2, there would be three absorbing states, \mathbf{z}^{AA} , \mathbf{z}^{BB} and \mathbf{z}^{AB} . Note that the assumption that agents toss a coin when indifferent implies that states along the demarcation lines are not absorbing.

From inequalities (4) and (5) we see that the two lines of demarcation have inverse slopes, p and $1/p$, respectively. And as p changes, the two lines pivot around the points $(\mathbf{a}_1 N_1, \mathbf{a}_1 N_2)$ and $(\mathbf{a}_2 N_1, \mathbf{a}_2 N_2)$, respectively. As p increases, the thicker line gets steeper, and as $p \rightarrow \infty$ it becomes vertical. Remember that an increased p means that agents interact more frequently with their own population. From a state in which members of population 1 are indifferent between strategy A and B, if the number of A-players in population 1 is reduced by one, the number of extra A-players needed in population 2 for population 1 to remain indifferent, increases with p . In other words, the more frequently you interact with your own population, the larger a change in the behaviour of the opposite population is needed to offset a given change of behaviour in your own population. Conversely, as p decreases, a given change in the behaviour of your own population is offset by smaller changes in the behaviour of the opposite population. Hence, the thicker line in figure 2 gets flatter as p decreases, and for $p \rightarrow 0$ it becomes horizontal. A similar line of arguments applied to population 2 tells us that the thinner line grows less steep as p increases, grows horizontal for $p \rightarrow \infty$, and vertical for $p \rightarrow 0$.

Finally, to gauge the relative attraction of multiple absorbing states, we introduce mutations into the decision making of agents. This takes the form that in each period each agent has a small probability ϵ of choosing strategies at random from a uniform distribution over the two strategies. This random choice then trumps any previous choice of strategy. In sum, then, we have a perturbed stochastic process. For a given level of viscosity p , let $\mathbf{P}(p, \epsilon)$ be the transition matrix implied by the above learning process including mutations. In other words, element ij of $\mathbf{P}(p, \epsilon)$ is the probability of going from state i to state j from one period to the next. For any given p , we can then represent the process by a transition matrix $\mathbf{P}(p, \epsilon)$ on a state space \mathbf{O} , for which we use the shorthand formulation $(\mathbf{O}, \mathbf{P}(p, \epsilon))$.

We are interested in where the process $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$ goes in the very long run when noise is very small. For given p , we therefore study the probability distribution over population states in \mathbf{O} , as time goes to infinity and noise \mathbf{e} to zero. States that have a positive probability in this distribution, are what Kandori, Mailath and Rob (1993) call long run equilibria, and what Young (1993) calls stochastically stable states. Below, the compromise term of long run stochastic stability proposed by Ellison (2000) is used to denote such states. Standard methods of computing long run stochastically stable states are used in the following analysis of the model, and described in more detail in an appendix.

Dispersive interaction

Let us start by analyzing the case where interaction is dispersive, in other words when interaction with any member of the opposite population is more frequent than with any member of your own population. In the limit, when interaction with your own population is so rare as to be relatively non-existent ($p \rightarrow 0$), dispersion is complete. The following proposition describes some key properties of the long run behaviour of the process in this case.

PROPOSITION 1:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p \rightarrow 0$. Then:

- i) For $N_1 = N_2$ sufficiently large, \mathbf{z}^{AA} is the unique long run stochastically stable state.*
- ii) For N_1 sufficiently large, there exists some $\hat{N}_2 > N_1$ such that for all $N_2 > \hat{N}_2$, \mathbf{z}^{BB} is the unique long run stochastically stable state.*

A formal proof of the proposition is given in an appendix, as are the proofs of later propositions.

For $p \rightarrow 0$, interaction takes the form of agents from different populations playing the battle of the sexes game (G3). Proposition 1i) basically states that if populations

are equally large, players conform to the risk dominant equilibrium (A,A) of this game in the long run. However, as part ii) of the proposition points out, if the population preferring the risk dominated equilibrium (B,B) is sufficiently much larger than the other population, the risk dominated equilibrium is played in the long run. These results are essentially the same as those captured by proposition 2 in Hehenkamp (2001).

The intuitive reason for the above results can be explained as follows. When $p \rightarrow 0$, inequalities (4) and (5) reduce to:

$$z_2 > \mathbf{a}_1 N_2 \tag{6}$$

$$z_1 > \mathbf{a}_2 N_1 \tag{7}$$

In a stability diagram, this implies that the line demarcating increases and decreases in z_1 is horizontal, and the corresponding line for z_2 is vertical. As in figure 2, the thick line in the below figure represents the former demarcation line, and the thin line the latter.

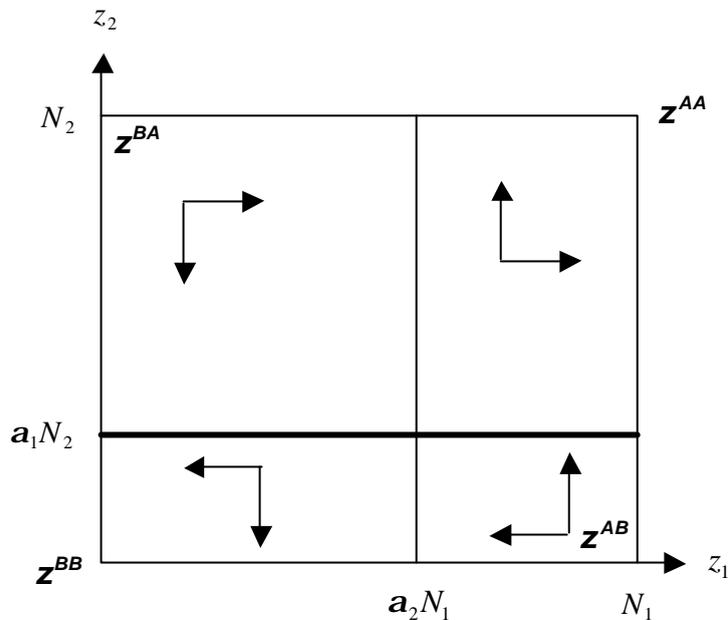


Figure 3. Stability diagram for the case of complete dispersion.

As the diagram reveals, no matter where we start out, we eventually reach either state \mathbf{z}^{AA} or state \mathbf{z}^{BB} . This includes starting points on the demarcation lines, since indifferent agents tossing a coin can then shift the process off the lines. For $p \rightarrow 0$, the unperturbed process thus has two absorbing states, \mathbf{z}^{AA} and \mathbf{z}^{BB} . As established by Young (1993), for perturbed processes of the kind studied here, the long run stochastically stable state must be one (or both) of these. In computing the long run stochastically stable state, we note that a number of unlikely mutations are needed to leave one absorbing state for the other. When the probability of mutations is infinitely small, transitions between states that require more mutations are infinitely less likely than transitions that require fewer mutations. With two absorbing states, the frequency with which we can expect to observe either one of them in the very long run, reflects the difficulty with which that state can be left for the other absorbing state. The long run stochastically stable state is thus the absorbing state that requires more mutations to leave for the other absorbing state, than vice versa.

A closer look at figure 3 tells us that to leave \mathbf{z}^{AA} , enough agents must mutate to strategy B to bring us to a state on either of the two demarcation lines. One of the populations then has B as a best reply, and we might therefore eventually reach \mathbf{z}^{BB} . If the populations are of equal size, $N_1 = N_2$, we can represent this in figure 3 by letting the units on both axes be of similar size. Clearly, then, we must move fewer units to the left from \mathbf{z}^{AA} to reach the thin demarcation line, compared to the number of units we would have to move downwards from \mathbf{z}^{AA} to reach the thick demarcation line. The easiest way to leave \mathbf{z}^{AA} , i.e. the way requiring the least mutations, is horizontally to the left, which takes $(1 - \mathbf{a}_2)N_1$ mutations to strategy B in population 1. Similarly, to leave \mathbf{z}^{BB} , we could either move upwards to the thick demarcation line or to the right to the thin line. Clearly, with similarly sized units on both axis, fewer mutations are needed to move up to the thick line, so to leave \mathbf{z}^{BB} we need a minimum of $\mathbf{a}_1 N_2$ mutations. Since we have assumed that $\mathbf{a}_1 < (1 - \mathbf{a}_2)$, it is thus easier to leave \mathbf{z}^{BB} for \mathbf{z}^{AA} than vice versa, which means that \mathbf{z}^{AA} is long run stochastically stable for equal population sizes.

If population 2 is larger than population 1, $N_2 > N_1$, we must represent the units in figure 3 differently. We can then think of the units on the second axis as being smaller than those on the first, agents are in a sense packed more densely on the second axis than the first. And if the density with which they are packed on the second axis grows sufficiently large, the number of units from \mathbf{z}^{BB} upwards to the thick demarcation line, exceeds the number of units from \mathbf{z}^{AA} leftwards to the thin demarcation line. In other words, if population 2 is sufficiently much larger than population 1, fewer mutations are needed to leave \mathbf{z}^{AA} for \mathbf{z}^{BB} than vice versa. This holds even if population 2 becomes so much larger that the easiest way to leave \mathbf{z}^{BB} is rightwards to the thin demarcation line, since this always implies a greater number of mutations than the transition from \mathbf{z}^{AA} to the thin line.

Incidentally, the reason why the sizes N_1 and N_2 of the two populations must be sufficiently large for the above results to hold, is as follows. Imagine that each population consisted of only one agent, $N_1 = N_2 = 1$. In that case, we could leave any absorbing state by means of a single mutation, and the criterion of long run stochastic stability would therefore not discriminate between absorbing states. By requiring that N_1 and N_2 be sufficiently large, we are in effect making sure that the units on the axes of figure 3 are sufficiently fine-grained for such a distinction between absorbing states to be made.

Fluid interaction

When an agent interacts as frequently with any agent from his own population, as with any agent from the other population, interaction is fluid. In the current model, this means that an agent engages in the same number of round-robin tournaments with both populations, $p = 1$. For a revising agent, this means that the actions of all agents in the preceding period receive the same weight in determining the optimal strategy. Since population sizes may differ, this means that the larger population has a greater impact on the decision of a revising agent than the smaller population. For fluid interaction, the following proposition holds.

PROPOSITION 2:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p = 1$. Then:

For N_1 and N_2 sufficiently large, there exists some $\hat{\mathbf{a}}_2$ such that for all $\mathbf{a}_2 > \hat{\mathbf{a}}_2$, \mathbf{z}^{AB} is long run stochastically stable.

On the face of it, proposition 2 states that if population 2 prefers strategy profile (B,B) sufficiently strongly over (A,A), then the state in which the two populations adhere to different conventions, \mathbf{z}^{AB} , is long run stochastically stable. Note, however, that on the assumption that $a > b$, $(1 - \mathbf{a}_1)$ is bounded below by \mathbf{a}_2 , which implies that the upper bound of \mathbf{a}_1 must decrease as \mathbf{a}_2 increases. An implication of proposition 2 is therefore that \mathbf{z}^{AB} is long run stochastically stable if \mathbf{a}_1 and \mathbf{a}_2 are sufficiently far apart. In other words, \mathbf{z}^{AB} is observed with certainty in the long run if both populations prefer their desired strategy profiles sufficiently strongly.

The intuitive argument underlying the proposition is as follows. With $p = 1$, inequalities (4) and (5) reduce to:

$$z_2 > \mathbf{a}_1(N_1 + N_2) - z_1 \quad (8)$$

$$z_2 > \mathbf{a}_2(N_2 + N_1) - z_1 \quad (9)$$

The lines demarcating the areas where z_1 increases and decreases, and z_2 increases and decreases, are now parallel and have a slope of -1 , with the latter line above and to the right of the former. In the below diagram, the thick line partitions the areas where z_1 increases and decreases, and the thin line the areas where z_2 increases and decreases.

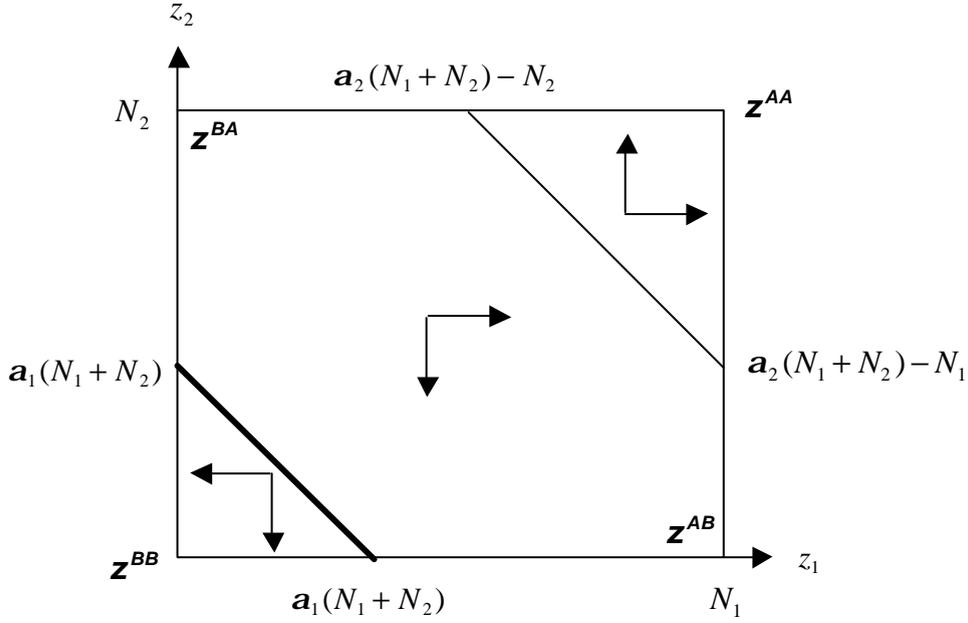


Figure 4. Sketch of stability diagram with fluid interaction.

In the above figure, there are three absorbing states, \mathbf{z}^{AA} , \mathbf{z}^{BB} and \mathbf{z}^{AB} . Note that for $p=1$, there need not be more than two absorbing states, \mathbf{z}^{AA} and \mathbf{z}^{BB} , since with unequal population sizes, $\mathbf{a}_1(N_1 + N_2)$ could exceed N_1 , or $\mathbf{a}_2(N_1 + N_2) - N_1$ could be negative. However, for given population sizes, we can construct a case in which there are three absorbing states by increasing \mathbf{a}_2 and hence decreasing the upper bound on \mathbf{a}_1 , which slides the two demarcation lines apart. At some level of $\mathbf{a}_2 < 1$, we thus get lines that cross in the manner of figure 4.

The same argument can be used to explain why \mathbf{z}^{AB} is the long run stochastically stable state if \mathbf{a}_2 is sufficiently large. As before, the long run stochastically stable state must be in the set of absorbing states. For \mathbf{a}_2 sufficiently large, there are three absorbing states, and a long run stochastically stable one can be characterized as being harder to leave for either of the other absorbing states, than it is to reach from the absorbing state from which it is hardest to reach. That \mathbf{z}^{AB} can have this property for large \mathbf{a}_2 , we can illustrate by examining the case where $\mathbf{a}_2 \rightarrow 1$. This implies that $\mathbf{a}_1 \rightarrow 0$. The demarcation lines then slide as far as they go into opposite corners of the above figure. To leave either of the two states \mathbf{z}^{AA} or \mathbf{z}^{BB} for \mathbf{z}^{AB} now requires only one mutation. However, to leave \mathbf{z}^{AB} for either of the two other absorbing states requires a minimum of $\min\{N_1, N_2\}$ mutations. For large population sizes, this means

that it is harder to leave \mathbf{z}^{AB} for either of the other two absorbing states, than it is to leave either of the other two for \mathbf{z}^{AB} . With fluid interaction, then, the state of coexistent conventions \mathbf{z}^{AB} can be long run stochastically stable if populations adhere sufficiently strongly to their preferred strategy profiles. This possibility of coexistence marks a departure from previous evolutionary models of learning, such as those of Kandori, Mailath and Rob (1993), Young (1993) and Hehenkamp (2001), whose results do not permit long run convention coexistence.

Viscous interaction

When interaction is viscous, an agent meets any member of his own population more frequently than any member of the opposite population. In our formulation, more rounds of round-robin are played with members of your own population than with the other population. In the limit, when interaction with the other population is comparatively non-existent, i.e. $p \rightarrow \infty$, we have complete viscosity. The following proposition captures the evolution of play in this instance.

PROPOSITION 3:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p \rightarrow \infty$. Then: If and only if N_1 and N_2 are sufficiently large, \mathbf{z}^{AB} is the unique long run stochastically stable state.

In other words, when the two populations virtually never interact, each population adopts its preferred strategy, regardless of the strength of that preference. This result is akin to the main result of Kandori, Mailath and Rob (1993), who find that a single population interacting only with itself will end up playing according to the risk dominant equilibrium. By implication, according to their model, two separate populations having different risk dominant equilibria, will thus play differently in the long run. This mirrors the case where $p \rightarrow \infty$, since we have two virtually separate populations playing games (G1) and (G2), respectively.

The procedure of comparing how easily absorbing states are left for and reached from other absorbing states is inconclusive in this instance, and proving that \mathbf{z}^{AB} is the only long run stochastically stable state is therefore a more complex operation. As shown in an appendix, the formal proof relies on a comparison of the ease with which an absorbing state can be reached by way of all the other absorbing states. In a sense, the absorbing state that can be reached with the least number of mutations in this manner, is long run stochastically stable.

It is difficult to illustrate this result by the simple means used in previous sections. However, the result mirrors the fact that \mathbf{z}^{AB} has the largest basin of attraction of the absorbing states in this case, i.e. there are more states from which we transit to \mathbf{z}^{AB} with certainty in a finite number of periods, than to any other absorbing state. This we can demonstrate graphically. For $p \rightarrow \infty$, inequalities (4) and (5) can be rewritten as:

$$z_1 > \mathbf{a}_1 N_1 \tag{10}$$

$$z_2 > \mathbf{a}_2 N_2 \tag{11}$$

The demarcation line which distinguishes between increases and decreases in z_1 is now vertical, and the line which separates increases and decreases in z_2 is horizontal. Let a thick and a thin line represent these two demarcation lines. The stability diagram then looks as follows.

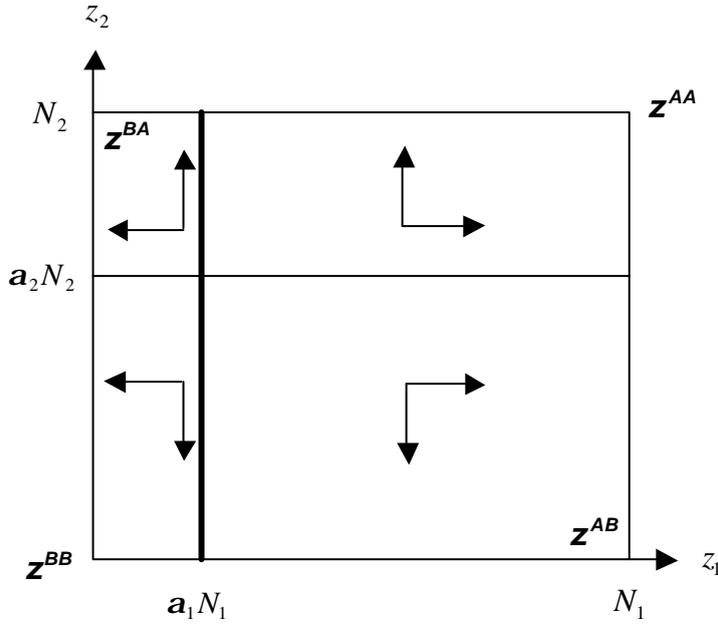


Figure 5. Stability diagram for the case of complete viscosity.

As the diagram shows, there are four absorbing states, \mathbf{z}^{AA} , \mathbf{z}^{BB} , \mathbf{z}^{AB} and \mathbf{z}^{BA} . The size of the basin of attraction of each is the area within which the learning dynamic brings us to that state. The four states thus have basins of attraction of sizes $(1 - \mathbf{a}_1)N_1 \cdot (1 - \mathbf{a}_2)N_2$, $\mathbf{a}_1N_1 \cdot \mathbf{a}_2N_2$, $(1 - \mathbf{a}_1)N_1 \cdot \mathbf{a}_2N_2$ and $\mathbf{a}_1N_1 \cdot (1 - \mathbf{a}_2)N_2$, respectively. Since $\mathbf{a}_1 < (1 - \mathbf{a}_2) < \mathbf{a}_2 < (1 - \mathbf{a}_1)$, state \mathbf{z}^{AB} thus has the larger basin of attraction of the four. In a sense, then, \mathbf{z}^{AB} is the absorbing state which is easiest to reach by way of the other absorbing states.

Welfare properties of the long run stochastically stable states

The welfare properties of long run stochastically stable states in learning models, has been the topic of much debate. For the processes studied by Kandori, Mailath and Rob (1993) and Young (1993), the long run stochastically stable states in coordination games entail play according to the risk dominant equilibria. This is potentially troublesome, since a risk dominant equilibrium can be Pareto dominated by another equilibrium. However, Bergin and Lipman (1996) show that by allowing mutation probabilities to vary between states, the Pareto dominant equilibrium can be selected in the very long run. In other evolutionary models such as that of Binmore, Samuelson

and Vaughan (1995), the structure of the payoffs determine whether we end up in the risk dominant or Pareto dominant equilibrium.

Less has been said about the welfare properties of games of conflicting interest, where gains for one player entail losses for another. In such games, the Pareto principle has no cutting power, and we need some other criterion by which to evaluate welfare, a criterion which weighs the relative payoffs of different populations. One such criterion would be a classical utilitarian one, where the better outcome is that which produces the highest payoffs summed over all individuals (see e.g. Harsanyi (1977)). In the current context, we can gauge the total payoffs in any state \mathbf{z} by the sum of the average payoffs $\mathbf{p}_i(\mathbf{z})$ to each population $i=1,2$, weighted by the size of each population N_i .⁴ If we define the relation $\mathbf{z}' \succ^U \mathbf{z}''$ as meaning that state \mathbf{z}' is strictly better than state \mathbf{z}'' from a utilitarian point of view, this relation is characterized as follows

$$\mathbf{z}' \succ^U \mathbf{z}'' \quad \text{iff} \quad N_1 \cdot \mathbf{p}_1(\mathbf{z}') + N_2 \cdot \mathbf{p}_2(\mathbf{z}') > N_1 \cdot \mathbf{p}_1(\mathbf{z}'') + N_2 \cdot \mathbf{p}_2(\mathbf{z}'') \quad (12)$$

Another criterion is the Rawlsian leximin principle, which claims that the better state is the one where the worst off group has the highest payoff, and if the worst off group is equally well off in two states, the better state is the one where the second worst off group has the highest payoff, and so on (Rawls, 1971). Let us define a relation $\mathbf{z}' \succ^R \mathbf{z}''$ as denoting that state \mathbf{z}' is strictly better from a Rawlsian perspective than \mathbf{z}'' . If we simplify slightly by equating groups with populations in our model, this relation has the following property⁵

$$\begin{aligned} \mathbf{z}' \succ^R \mathbf{z}'' \\ \text{if } [\min \{\mathbf{p}_1(\mathbf{z}'), \mathbf{p}_2(\mathbf{z}')\} > \min \{\mathbf{p}_1(\mathbf{z}''), \mathbf{p}_2(\mathbf{z}'')\}] \\ \text{or } [\min \{\mathbf{p}_1(\mathbf{z}'), \mathbf{p}_2(\mathbf{z}')\} = \min \{\mathbf{p}_1(\mathbf{z}''), \mathbf{p}_2(\mathbf{z}'')\} \\ \text{and } \max \{\mathbf{p}_1(\mathbf{z}'), \mathbf{p}_2(\mathbf{z}')\} > \max \{\mathbf{p}_1(\mathbf{z}''), \mathbf{p}_2(\mathbf{z}'')\}] \end{aligned} \quad (13)$$

⁴ Average payoffs means average both across encounters within a period and across the members of population i . To meaningfully sum these we must assume cardinal interpersonal comparability of payoffs.

⁵ Ordinal interpersonal comparability of payoffs must be assumed in this case.

Let us evaluate the long run stochastically stable states established above according to these criteria. In the case of full dispersion, $p \rightarrow 0$, the following result holds.

PROPOSITION 4:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p \rightarrow 0$. Then: For N_1 sufficiently large, there exists some $\tilde{N}_2 > \hat{N}_2 > N_1$ such that for $N_2 \in \langle \hat{N}_2, \tilde{N}_2 \rangle$, \mathbf{z}^{BB} is the unique long run stochastically stable state, while $\mathbf{z}^{AA} \succ^U \mathbf{z}^{BB}$ and $\mathbf{z}^{AA} \succ^R \mathbf{z}^{BB}$.

The proposition says that there is a range of relative population sizes within which the long run stochastically stable state is not the absorbing state producing the maximum total payoff, nor is it the state leaving the worst off population better off. In other words, for some population sizes, the evolutionary process selects a state which is inoptimal from a utilitarian and from a Rawlsian point of view.

A simple way to understand the fact that a state is selected which does not maximize total payoff, is to note that utilitarianism and the evolutionary process implicitly maximize different things. According to utilitarianism, \mathbf{z}^{BB} is better than \mathbf{z}^{AA} if the sum of the payoffs of the two populations is larger in the former state, i.e. if

$$N_1 \cdot 1 + N_2 \cdot b > N_1 \cdot a + N_2 \cdot 1 \quad (14)$$

From the discussion and proof of proposition 1, we know that the evolutionary process selects \mathbf{z}^{BB} if $\frac{N_2}{N_1} > \frac{(1 - a_2)}{a_1}$. Using the definitions of a_1 and a_2 , we can rewrite this inequality as

$$N_2 \cdot b - N_1 \cdot 1 > N_1 \cdot a - N_2 \cdot 1 \quad (15)$$

In other words, the evolutionary process selects \mathbf{z}^{BB} if the difference between the total payoffs of the best off population and the worst off population in that state is greater than the difference between the best off and worst off population in \mathbf{z}^{AA} .

Where utilitarianism maximizes the sum of payoffs, the evolutionary process thus implicitly maximizes the difference in payoffs between the better and worse off population. As a consequence, while utilitarianism is egalitarian in letting everyone count for one, the evolutionary process is fiercely inegalitarian in letting the worse off agents count negatively.

The reason for the divergence between utilitarianism and the evolutionary process, is that the evolutionary process selects the state more robust to mutations, which is not necessarily the state that yields the highest total payoff.⁶ Rewriting (15), we get that the evolutionary process selects \mathbf{z}^{BB} if

$$\frac{N_2}{N_1} > \frac{(1-a_2)}{a_1} = \frac{a+1}{b+1} \quad (16)$$

The degree to which population 2 must be larger than population 1, depends on the ease with which population 2 switches to B, compared to the ease with which population 1 switches to A. This in turn proves a matter of how large the sum of payoffs over the two states \mathbf{z}^{AA} and \mathbf{z}^{BB} is for the two populations. Which state is more robust to mutations thus depends on population sizes and *total* payoffs over the states.

By contrast, a version of (14) tells us that utilitarianism prefers \mathbf{z}^{BB} if

$$\frac{N_2}{N_1} > \frac{a-1}{b-1} \quad (17)$$

Utilitarianism thus focuses on payoff *differences* between the states \mathbf{z}^{AA} and \mathbf{z}^{BB} . The degree to which population 2 must be larger, depends on the loss incurred by each member of population 1 in moving from \mathbf{z}^{AA} to \mathbf{z}^{BB} , compared to the gain to each member of population 2 in moving between the two states. The two populations differ less in their total payoffs across the two states, than in their payoff differences

⁶ See Hehenkamp (2001) for a discussion of the relation between selection in this case and risk dominance.

between the states. By implication, for the lowest population ratios at which \mathbf{z}^{BB} is more robust to mutations than \mathbf{z}^{AA} , the utilitarian principle prefers \mathbf{z}^{AA} over \mathbf{z}^{BB} .

On the Rawlsian leximin principle, \mathbf{z}^{BB} is not preferable to \mathbf{z}^{AA} for any relative population sizes. This stems from the fact that the average payoff of the worst off population is 1 in both \mathbf{z}^{AA} and \mathbf{z}^{BB} , which means that we must compare the payoffs of the best off population in each state, which is a in \mathbf{z}^{AA} and b in \mathbf{z}^{BB} . Given the view the evolutionary process takes of the payoffs of the worst off population, it is not very surprising that the process in some cases selects a state which is worse according to the leximin principle.

Turning to the case of fluid interaction, $p = 1$, we can prove the following result.

PROPOSITION 5:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p = 1$. Then:

For $N_1 = N_2$ sufficiently large, there exists some \hat{a}_2 such that for all $a_2 > \hat{a}_2$, \mathbf{z}^{AB} is long run stochastically stable, while $\mathbf{z}^{AA} \succ^U \mathbf{z}^{AB}$ and $\mathbf{z}^{AB} \succ^R \mathbf{z}^{AA}$.

Remember from proposition 2 that the state of convention coexistence, \mathbf{z}^{AB} , is long run stochastically stable when the populations are sufficiently biased in favour of their preferred strategy. Proposition 5 states that for equal population sizes, if said bias is sufficient for \mathbf{z}^{AB} to be long run stochastically stable, \mathbf{z}^{AB} is worse in terms of total payoff than \mathbf{z}^{AA} , but better in terms of payoff to the worst off population.⁷ Note that in state \mathbf{z}^{AA} , all encounters entail coordination, which means that population 1 earns an average payoff of a , whereas population 2 earns 1. In \mathbf{z}^{AB} , on the other hand, there is only coordination when members of the same population meet, i.e. in half the encounters of each player. Population 1 thus earns on average $a/2$ and population 2 earns $b/2$. The proof of proposition 5 shows that for $b > 3$, \mathbf{z}^{AB} is long run stochastically stable. Thus, from the average payoffs we see that the worse off population 2 is better off in \mathbf{z}^{AB} than \mathbf{z}^{AA} , when \mathbf{z}^{AB} is long run stochastically

⁷ Indeed, one can show that \mathbf{z}^{AB} is worse in terms of total payoff than \mathbf{z}^{BB} if $a - b < 2$.

stable. However, in going from \mathbf{z}^{AA} to \mathbf{z}^{AB} , a member of population 2 improves his average payoff by $\frac{b-2}{2}$. On the other hand, the loss incurred by a member of population 1 from such a transition is $a/2$. Since by assumption $a > b$, the loss to population 1 is thus greater than the gain to population 2, which implies that \mathbf{z}^{AA} is better from a utilitarian perspective than \mathbf{z}^{AB} . In the current context, then, utilitarianism prefers coordination since it has a favourable impact on total payoffs. The result highlights the fact that utilitarianism cares only about the total level of utility, and cares not about how that total is distributed among agents.

Though the results so far are mixed, that is not the case when interaction is completely viscous.

PROPOSITION 6:

Consider the process of learning with noise $(\mathbf{O}, \mathbf{P}(p, \mathbf{e}))$, and suppose $p \rightarrow \infty$. Then:

$$\mathbf{z}^{AB} \succ^U \mathbf{z}^{AA}, \mathbf{z}^{AB} \succ^U \mathbf{z}^{BB} \text{ and } \mathbf{z}^{AB} \succ^R \mathbf{z}^{AA} \succ^R \mathbf{z}^{BB}.$$

In other words, \mathbf{z}^{AB} is better on both criteria when interaction is completely viscous. The simple reason is that with complete viscosity, agents only play their own population, and in the absorbing states in question, they always coordinate with their opponents. In state \mathbf{z}^{AB} , both groups play according to their preferred equilibria, and thus get an average payoff of a and b , respectively. In state \mathbf{z}^{AA} , only population 1 gets to play its preferred equilibrium, which makes average payoffs a and 1, respectively. In state \mathbf{z}^{BB} , only population 2 plays its preferred equilibrium, which makes average payoffs 1 and b , respectively. Even a casual glance at these numbers reveals that \mathbf{z}^{AB} is better than the other two both in terms of total payoff, and in terms of the payoff accruing to the worst off group. With complete viscosity, convention coexistence thus has some merit.

Concluding remarks

The basic argument of this paper is that in modelling the interaction of several distinct populations, we should allow for the possibility that agents might interact more frequently, or less frequently, with members of their own population than with members of another population. The framework proposed above captures this idea in a simple manner, and permits the study of the whole range of two-population interaction, from complete dispersion through complete viscosity. Though the round-robin matching regime could be expanded into a model featuring any number of populations, one need not have more than two populations with conflicting interests to obtain noteworthy results. Notably, in an evolutionary model of learning based on this matching regime, results similar to those of previous evolutionary models emerge in the special cases of complete viscosity and complete dispersion. Moreover, a novel result from this model is that interacting populations can exhibit different conventions in the very long run, there is in other words a possibility of convention coexistence.

The notion that populations have conflicting interests also facilitates a richer understanding of the normative properties of states selected by evolutionary dynamics. What makes one state better than another is less of a trivial matter when interests diverge, and since different normative principles take different views of this matter, we get a more detailed basis on which to evaluate evolutionary selection. As noted previously, selection in evolutionary models of learning focus on robustness against errors or mutations, and the characteristics that make one state normatively preferable to another are only important to the extent that they influence robustness. It is therefore not hard to appreciate why evolutionary selection is sometimes at odds with what is normatively preferable. For utilitarian and Rawlsian views of what is normatively preferable, the results obtained above suggest that whether evolutionary selection and normative principles diverge, depends *inter alia* on the degree of viscosity in interaction.

Appendix A. Proof of propositions 1 through 3

The process defined by strategy revision and mutations is a discrete time Markov process on a finite state space \mathbf{O} , since the probability of transiting between two states from the current period to the next, depends on the properties of no state other than the current. For any $\epsilon > 0$, there is a positive probability of moving from any state in \mathbf{O} to any other state in \mathbf{O} in a finite number of periods, which by definition means the process is irreducible. Let μ be a probability distribution over the states in \mathbf{O} , and $\mathbf{P}(p, \epsilon)$ be the matrix of transition probabilities. For an irreducible process, a standard result for finite Markov chains states that there exists a unique solution to the following equation:

$$\mu \mathbf{P}(p, \epsilon) = \mu \quad (\text{A1})$$

In other words, such a process has a unique stationary distribution, which we term μ^ϵ . Moreover, the process in question is aperiodic, since we can move from state \mathbf{z} and back again in any positive number of periods, for any state \mathbf{z} in \mathbf{O} .⁸ Let $v^t(\mathbf{z} | \mathbf{z}^0)$ be the probability that at time t we are in state \mathbf{z} , when at time 0 we were in state \mathbf{z}^0 . For an aperiodic and irreducible process the following result holds:

$$\lim_{t \rightarrow \infty} v^t(\mathbf{z} | \mathbf{z}^0) = m^\epsilon(\mathbf{z}) \quad (\text{A2})$$

The probability of being in a certain state \mathbf{z} as time goes to infinity, thus converges to the probability $m^\epsilon(\mathbf{z})$ awarded that state by the stationary distribution. The probability that the process reaches any state after a large number of periods, is thus independent of the initial state.

⁸ Young (1998) gives a precise definition of an aperiodic process: For each state $\mathbf{z} \in \mathbf{O}$, "let $N_{\mathbf{z}}$ be the set of all integers $n \geq 1$ such that there is a positive probability of moving from \mathbf{z} to \mathbf{z} in exactly n periods. The process is aperiodic if for every \mathbf{z} , the greatest common divisor of $N_{\mathbf{z}}$ is unity."

The stationary distribution \mathbf{m}^ϵ is difficult to compute, so we focus on the case where the probability of mutations is arbitrarily small, $\lim_{\epsilon \rightarrow 0} \mathbf{m}^\epsilon$. States \mathbf{z} which have the property that $\lim_{\epsilon \rightarrow 0} \mathbf{m}^\epsilon(\mathbf{z}) > 0$, we define as long run stochastically stable. A fundamental result by Young (1993) establishes that $\lim_{\epsilon \rightarrow 0} \mathbf{m}^\epsilon$ exists, and equals a stationary distribution of the corresponding process without mutations, $\epsilon = 0$. Note that the process without mutations is not irreducible, which means that it can have several stationary distributions. For any such stationary distribution, \mathbf{m}^0 , the states \mathbf{z} that have positive probabilities in this distribution, $\mathbf{m}^0(\mathbf{z}) > 0$, constitute a limit set of the process. Young proceeds to prove that the long run stochastically stable states are those contained in the limit sets that have minimum stochastic potential. Stated differently, $\lim_{\epsilon \rightarrow 0} \mathbf{m}^\epsilon$ equals the stationary distribution \mathbf{m}^{0*} which puts positive probability on the limit set having minimum stochastic potential.

To find the long-run stochastically stable states of a process, we thus first find the limit sets when mutations are absent, and then compute the stochastic potential of these limit sets. The search for limit sets is executed as follows. A state \mathbf{z}' is accessible from \mathbf{z} , if there is a positive probability of reaching \mathbf{z}' from \mathbf{z} in a finite number of periods. Two states communicate if each is accessible from the other. A limit set is a set of states such that all states in the set communicate, and no state outside the set is accessible from any state inside the set. A limit set is thus a set of states which once reached, the process never leaves. An absorbing state is a limit set consisting of a singleton state.

To find the limit sets with minimum stochastic potential, i.e. the long run stochastically stable states, we can proceed in two ways, one simple yet in some cases inconclusive, and the other more complex yet conclusive. The simpler method is due to Ellison (2000), who defines two characteristic numbers for each limit set \mathbf{Z} , a radius $R(\mathbf{Z})$ and a coradius $CR(\mathbf{Z})$. In the current context of equiprobable mutations, the radius $R(\mathbf{Z})$ of a limit set \mathbf{Z} is the minimum number of mutations needed to leave \mathbf{Z} and enter a state from which another limit set is accessible. The radius thus provides a measure of how easily \mathbf{Z} can be left for another limit set. To compute the

coradius $CR(\mathbf{Z})$, you take the minimum number of mutations needed to leave each of the limit sets different from \mathbf{Z} for a state from which \mathbf{Z} is accessible, and let the coradius equal the maximum of these. The coradius thus measures how easily \mathbf{Z} can be reached from the other limit sets, specifically from the limit set from which \mathbf{Z} is most difficult to reach. Ellison proves that if $R(\mathbf{Z}) > CR(\mathbf{Z})$, then the states in \mathbf{Z} are long run stochastically stable. This condition is just a sufficient condition for long run stochastic stability, there can thus be long run stochastically stable states which Ellison's method does not identify.

Which brings us to the more complex method which produces a complete characterization of long run stochastically stable states. This is the original method of tree surgery devised by Young (1993). For all limit sets, find the minimum number of mutations needed to go from one limit set to another. Next, for each limit set, construct all possible trees rooted at that set. A tree rooted at limit set \mathbf{Z} has the property that from each limit set different from \mathbf{Z} there is a unique sequence of directed edges between limit sets leading to \mathbf{Z} . As an example, imagine that there are three limit sets; \mathbf{A} , \mathbf{B} and \mathbf{C} . There are thus three trees rooted at \mathbf{A} :

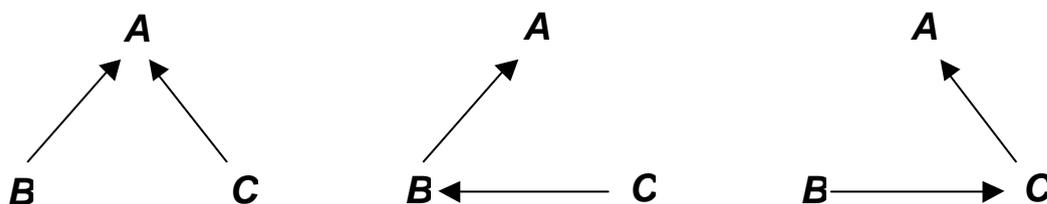


Figure A1. Trees rooted at \mathbf{A} with three limit sets \mathbf{A} , \mathbf{B} and \mathbf{C} .

Similarly, there are three trees rooted at \mathbf{B} , and three trees rooted at \mathbf{C} . For all trees constructed in this manner, we compute the sum of the mutations associated with each directed edge. We then find the tree with the minimum sum of mutations, and the limit set rooted at this tree has minimum stochastic potential. In other words, it contains the stochastically stable states of the process in question.

PROOF OF PROPOSITION 1:

For $p \rightarrow 0$, let us first compute the limit sets. For $e = 0$, \mathbf{z}^{AA} and \mathbf{z}^{BB} are absorbing states of the process. In \mathbf{z}^{AA} , everyone plays A, and any revising player thus expects to meet only A-players in the next period, which makes A his optimal choice. Similarly, in \mathbf{z}^{BB} , everyone plays B, and any revising player chooses strategy B. No other states are thus accessible from \mathbf{z}^{AA} and \mathbf{z}^{BB} , which makes them absorbing states.

Note that all members of a specific population have the same best reply to the population state of the preceding period. In a given period, there is a positive probability that all agents are drawn for strategy revision, in which case all agents from each population choose the same strategy if they have a unique best reply to the preceding state. If the members of a population are indifferent between strategies A and B, i.e. if they have several best replies, they all toss a coin. Since there is a positive probability that all coin tosses show the same result, there is a positive probability that all members of a population choose the same strategy. The states \mathbf{z}^{AA} , \mathbf{z}^{BB} , \mathbf{z}^{AB} or \mathbf{z}^{BA} are thus accessible from any other state. However, since in \mathbf{z}^{AB} a revising player from population 1 expects to meet only B-playing members, he would switch to B. And due to the fact that there is a positive probability that all population 1 agents are drawn for revision, they might all change to B, which means that a transition from \mathbf{z}^{AB} to \mathbf{z}^{BB} has positive probability. Similarly, in state \mathbf{z}^{BA} , all population 1 agents might be drawn for revision, having A as their optimal choice. From \mathbf{z}^{BA} a transition to \mathbf{z}^{AA} thus has positive probability. In sum, this means that the absorbing states \mathbf{z}^{AA} or \mathbf{z}^{BB} are accessible from any other state in \mathbf{O} , which implies that no other state can be contained in a limit set.

With two limit sets, the above two methods of computing the long run stochastically stable states are equivalent. The reason is that each limit set is at the root of only one tree, and the mutations associated with the single directed edge of this tree equals the coradius of this limit set and the radius of the other limit set. For the limit set at the root of the tree with the minimum number of mutations, the radius thus exceeds the coradius. The radius-coradius method in this way provides a full characterization of

long run stochastically stable states in this instance, and we can therefore use it to identify the unique long run stochastically stable states of the process.

As figure 3 reveals, the shortest route from \mathbf{z}^{AA} to the a state from which \mathbf{z}^{BB} is accessible, is along one of the axes. The minimum number of mutations needed to leave \mathbf{z}^{AA} for \mathbf{z}^{BB} is therefore $\min \{(1 - \mathbf{a}_1)N_2, (1 - \mathbf{a}_2)N_1\}$. In other words:

$$R(\mathbf{z}^{AA}) = CR(\mathbf{z}^{BB}) = \min \{(1 - \mathbf{a}_1)N_2, (1 - \mathbf{a}_2)N_1\} \quad (\text{A3})$$

Similarly, the shortest way from \mathbf{z}^{BB} to \mathbf{z}^{AA} is along either axis, which makes

$$R(\mathbf{z}^{BB}) = CR(\mathbf{z}^{AA}) = \min \{\mathbf{a}_1N_2, \mathbf{a}_2N_1\} \quad (\text{A4})$$

Since $\mathbf{a}_1 < (1 - \mathbf{a}_2) < \mathbf{a}_2 < (1 - \mathbf{a}_1)$, for $N_1 = N_2$ we have

$$R(\mathbf{z}^{AA}) = (1 - \mathbf{a}_2)N_1 > CR(\mathbf{z}^{AA}) = \mathbf{a}_1N_2 \quad (\text{A5})$$

In other words, for equal population sizes, \mathbf{z}^{AA} is the unique long run stochastically stable state.

For $N_2 > N_1$, we still have $CR(\mathbf{z}^{BB}) = (1 - \mathbf{a}_2)N_1$. However, for $\frac{N_2}{N_1} > \frac{(1 - \mathbf{a}_2)}{\mathbf{a}_1}$, we

get the following relation:

$$R(\mathbf{z}^{BB}) = \min \{\mathbf{a}_1N_2, \mathbf{a}_2N_1\} > CR(\mathbf{z}^{BB}) = (1 - \mathbf{a}_2)N_1 \quad (\text{A6})$$

For sufficiently large $\frac{N_2}{N_1}$, \mathbf{z}^{BB} is thus the unique long run stochastically stable

state.?

PROOF OF PROPOSITION 2:

For $p = 1$, the following arguments apply. For $\mathbf{e} = 0$, \mathbf{z}^{AA} and \mathbf{z}^{BB} are still absorbing states, for the same reasons as in the proof of proposition 1. Note that in state \mathbf{z}^{AB} , $z_1 = N_1$ and $z_2 = 0$. From inequality (2), we thus see that the frequency with which a member of population 1 encounters A-players is $\frac{N_1}{N_1 + N_2}$. If this frequency exceeds \mathbf{a}_1 , revising agents from population 1 keep playing A, i.e. if

$$\mathbf{a}_1 < \frac{N_1}{N_1 + N_2} \quad (\text{A7})$$

Similarly, the frequency with which population 2 players meet A-players is $\frac{N_1}{N_1 + N_2}$.

If this frequency does not exceed \mathbf{a}_2 , revising agents from population 2 keep playing B, that is if:

$$\mathbf{a}_2 > \frac{N_1}{N_1 + N_2} \quad (\text{A8})$$

Since $\mathbf{a}_1 < (1 - \mathbf{a}_2)$, for $N_1, N_2 \in \langle 0, \infty \rangle$ there exists some $\mathbf{a}_2 < 1$ for which (A7) and (A8) hold. For sufficiently large \mathbf{a}_2 , \mathbf{z}^{AB} is thus an absorbing state.

As before, there is a positive probability that all agents revise simultaneously. As agents from the same population have the same optimal strategy, there is thus a positive probability that we end up in \mathbf{z}^{AA} , \mathbf{z}^{BB} , \mathbf{z}^{AB} or \mathbf{z}^{BA} from any state other than these four. In \mathbf{z}^{BA} , an agent from either population encounters $\frac{N_2}{N_1 + N_2}$ A-players. A revising member of population 1 would continue playing B if

$$\frac{N_2}{N_1 + N_2} < \mathbf{a}_1 \quad (\text{A9})$$

Similarly, a revising population 2 agent would keep playing A if

$$\frac{N_2}{N_1 + N_2} > \mathbf{a}_2 \quad (\text{A10})$$

However, since $\mathbf{a}_2 > \mathbf{a}_1$, (A9) and (A10) are incompatible. Thus, in state \mathbf{z}^{BA} , if all agents revise, all members of either or both populations will switch strategies, in which case we transit to \mathbf{z}^{AA} , \mathbf{z}^{BB} or \mathbf{z}^{AB} .

Next, we find the stochastically stable states when there are three absorbing states, \mathbf{z}^{AA} , \mathbf{z}^{BB} and \mathbf{z}^{AB} . As figure 4 reveals, the easiest way to leave one absorbing state for another, is along the axes. We want to prove that \mathbf{z}^{AB} can be stochastically stable, so let us compute the radius and coradius of this absorbing state. The radius is:

$$R(\mathbf{z}^{AB}) = \min \{N_1 - \mathbf{a}_1(N_1 + N_2), \mathbf{a}_2(N_1 + N_2) - N_1\} \quad (\text{A11})$$

And the coradius (when we make use of the fact that $\mathbf{a}_1 < (1 - \mathbf{a}_2)$):

$$CR(\mathbf{z}^{AB}) = \max \{(1 - \mathbf{a}_2)(N_1 + N_2), \mathbf{a}_1(N_1 + N_2)\} = (1 - \mathbf{a}_2)(N_1 + N_2) \quad (\text{A12})$$

If the first element in the radius expression is the minimum of the two, then the radius exceeds the coradius when:

$$N_1 - \mathbf{a}_1(N_1 + N_2) > (1 - \mathbf{a}_2)(N_1 + N_2) \Leftrightarrow (1 - \mathbf{a}_2) + \mathbf{a}_1 < \frac{N_1}{N_1 + N_2} \quad (\text{A13})$$

If the second element is the minimum, the radius exceeds the coradius when:

$$\mathbf{a}_2(N_1 + N_2) - N_1 > (1 - \mathbf{a}_2)(N_1 + N_2) \Leftrightarrow 2\mathbf{a}_2 - 1 > \frac{N_1}{N_1 + N_2} \quad (\text{A14})$$

Since $\mathbf{a}_1 < (1 - \mathbf{a}_2)$, for $N_1, N_2 \in \langle 0, \infty \rangle$ there exists some $\mathbf{a}_2 < 1$ for which (A13) and (A14) hold. For sufficiently large \mathbf{a}_2 , \mathbf{z}^{AB} is thus stochastically stable.?

PROOF OF PROPOSITION 3:

For $p \rightarrow \infty$, the limit sets are as follows. As above, \mathbf{z}^{AA} , \mathbf{z}^{BB} , \mathbf{z}^{AB} or \mathbf{z}^{BA} are accessible from any other state. For the same reasons as above, \mathbf{z}^{AA} and \mathbf{z}^{BB} are absorbing states. With completely viscous interaction, in state \mathbf{z}^{AB} any revising player from population 1 expects to meet only A-players from his own population, and therefore keeps playing A. Any revising player from population 2 expects to meet only B-players from population 2, and hence keeps playing B. \mathbf{z}^{AB} is therefore an absorbing state. In \mathbf{z}^{BA} , agents from population 1 meet only B-players and keep playing B, while agents from population 2 meet only A-players and keep playing A, so \mathbf{z}^{BA} is also an absorbing state.

The radius-coradius method does not identify the long run stochastically stable state in this case. We therefore use the more elaborate tree surgery method. The following matrix reflects the minimum number of mutations needed to transit from the absorbing states in the rows to those in the columns, cf. figure 5.

	\mathbf{z}^{AA}	\mathbf{z}^{BB}	\mathbf{z}^{AB}	\mathbf{z}^{BA}
\mathbf{z}^{AA}	–	$(1 - \mathbf{a}_1)N_1 + (1 - \mathbf{a}_2)N_2$	$(1 - \mathbf{a}_2)N_2$	$(1 - \mathbf{a}_1)N_1$
\mathbf{z}^{BB}	$\mathbf{a}_1N_1 + \mathbf{a}_2N_2$	–	\mathbf{a}_1N_1	\mathbf{a}_2N_2
\mathbf{z}^{AB}	\mathbf{a}_2N_2	$(1 - \mathbf{a}_1)N_1$	–	$(1 - \mathbf{a}_1)N_1 + \mathbf{a}_2N_2$
\mathbf{z}^{BA}	\mathbf{a}_1N_1	$(1 - \mathbf{a}_2)N_2$	$\mathbf{a}_1N_1 + (1 - \mathbf{a}_2)N_2$	–

(A15)

With four absorbing states, there are 16 trees rooted at each absorbing state, 64 trees in all. I do not recount the total mutations associated with each of them here, but these calculations are available upon request. Note, however, the following two trees rooted at \mathbf{z}^{AB}

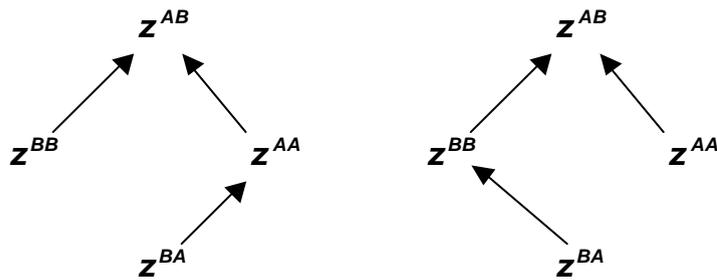


Figure A2. Two trees rooted at \mathbf{z}^{AB} .

The sum of mutations for the transitions of each tree is $2\mathbf{a}_1N_1 + (1 - \mathbf{a}_2)N_2$ and $\mathbf{a}_1N_1 + 2(1 - \mathbf{a}_2)N_2$, respectively. A comparison with the total mutations of all other trees, reveals that they all have a sum of mutations higher than either of these two trees. A \mathbf{z}^{AB} -tree thus has the minimum total mutations associated with it, which means that \mathbf{z}^{AB} is the unique long run stochastically stable state. Note that this result holds if and only if populations sizes are sufficiently large, if for instance each population consisted of only one agent, the minimum mutation tree for any absorbing state would involve three mutations, which implies that there is no unique long run stochastically stable state. ?

Appendix B: Proof of propositions 4 through 6

PROOF OF PROPOSITION 4:

Suppose $p \rightarrow 0$ and define \hat{N}_2 as follows

$$\frac{\hat{N}_2}{N_1} \equiv \frac{(1 - \mathbf{a}_2)}{\mathbf{a}_1} = \frac{1 - \frac{b}{1+b}}{\frac{1}{1+a}} = \frac{1+a}{1+b} \quad (\text{B1})$$

From proposition 1, we know that if for $N_2 > \hat{N}_2$, \mathbf{z}^{BB} is the unique long run stochastically stable state.

Since agents only play with the other population, a player from population 1 earns a payoff of a from each encounter in \mathbf{z}^{AA} , and a payoff of 1 in state \mathbf{z}^{BB} . Similarly, a player from population 2 earns a payoff of 1 per encounter in \mathbf{z}^{AA} and b in \mathbf{z}^{BB} . Thus, we know that $\mathbf{z}^{AA} \succ^U \mathbf{z}^{BB}$ if

$$a \cdot N_1 + 1 \cdot N_2 > 1 \cdot N_1 + b \cdot N_2 \Leftrightarrow \frac{N_2}{N_1} < \frac{a-1}{b-1} \quad (\text{B2})$$

However, for $a > b$

$$\frac{a-1}{b-1} > \frac{1+a}{1+b} \quad (\text{B3})$$

In other words, if we define \tilde{N}_2 as follows

$$\frac{\tilde{N}_2}{N_1} \equiv \frac{a-1}{b-1} \quad (\text{B4})$$

Then $\tilde{N}_2 > \hat{N}_2$, and for $N_2 \in \langle \hat{N}_2, \tilde{N}_2 \rangle$, \mathbf{z}^{BB} is long run stochastically stable and $\mathbf{z}^{AA} \succ^U \mathbf{z}^{BB}$.

That $\mathbf{z}^{AA} \succ^R \mathbf{z}^{BB}$, is a trivial implication of the average payoffs discussed above. The worse off population is equally badly off in both states, whereas the better off population is better off in state \mathbf{z}^{AA} than in \mathbf{z}^{BB} .?

PROOF OF PROPOSITION 5:

Suppose $p = 1$. For equal population sizes $N_1 = N_2 = N$, (A11) and (A12) become:

$$R(\mathbf{z}^{AB}) = \min \{(1 - 2\mathbf{a}_1)N, (2\mathbf{a}_2 - 1)N\} \quad (\text{B5})$$

$$CR(\mathbf{z}^{AB}) = 2(1 - \mathbf{a}_2)N \quad (\text{B6})$$

Furthermore, $\mathbf{a}_1 < (1 - \mathbf{a}_2)$ implies:

$$R(\mathbf{z}^{AB}) = (2\mathbf{a}_2 - 1)N \quad (\text{B7})$$

Consequently, we have:

$$R(\mathbf{z}^{AB}) > CR(\mathbf{z}^{AB}) \Leftrightarrow 2\mathbf{a}_2 - 1 > 2 - 2\mathbf{a}_2 \Leftrightarrow \mathbf{a}_2 > 0.75 \quad (\text{B8})$$

If $\hat{\mathbf{a}}_2 = 0.75$, then for all $\mathbf{a}_2 > \hat{\mathbf{a}}_2$, \mathbf{z}^{AB} is long run stochastically stable. Note that $\mathbf{a}_2 > 0.75$ implies $a > b > 3$.

Since populations are equally large, an agent plays as many encounters with his own population as the opposite one. In state \mathbf{z}^{AB} , an agent from population 1 coordinates with his own kind for a payoff of a half the time, and miscoordinates with the other population for zero payoff half the time, which makes his average payoff in state \mathbf{z}^{AB}

$a/2$. The same happens to agents from population 2, except they coordinate with their own kind for a payoff of b , which makes their average payoff $b/2$. In state \mathbf{z}^{AA} , there is always coordination, and agents from population 1 earn a from each encounter, and agents from population 2 earn 1. In comparing total payoffs in the two states, we have $\mathbf{z}^{AA} \succ^U \mathbf{z}^{AB}$ when

$$(a+1) \cdot N > \frac{a+b}{2} \cdot N \Leftrightarrow a > b-2 \quad (\text{B9})$$

In other words, since $a > b$, \mathbf{z}^{AA} is always better in terms of total payoff than \mathbf{z}^{AB} , which also applies to the cases where \mathbf{z}^{AB} is long run stochastically stable.

Population 2 is worse off than population 1 in both state \mathbf{z}^{AA} and state \mathbf{z}^{AB} , and gets an average payoff of $b/2$ in the former state and 1 in the latter. Thus, for $b > 3$, \mathbf{z}^{AB} is the better state according to the leximin principle, $\mathbf{z}^{AB} \succ^R \mathbf{z}^{AA}$. In other words, when \mathbf{z}^{AB} is stochastically stable, it is preferable to \mathbf{z}^{AA} on the leximin principle.?

PROOF OF PROPOSITION 6:

Suppose $p \rightarrow \infty$. The members of a population thus only play against each other. In state \mathbf{z}^{AB} , both populations coordinate on their preferred equilibria, which makes payoffs a in population 1 and b in population 2. In \mathbf{z}^{AA} , only population 1 plays according to its preferred equilibrium, which makes payoffs a and 1, respectively. In \mathbf{z}^{BB} , the converse is true, and payoffs are 1 and b , respectively. If we compare the total payoffs in these three states we thus get

$$aN_1 + bN_2 > aN_1 + N_2 \Rightarrow \mathbf{z}^{AB} \succ^U \mathbf{z}^{AA} \quad (\text{B10})$$

$$aN_1 + bN_2 > N_1 + bN_2 \Rightarrow \mathbf{z}^{AB} \succ^U \mathbf{z}^{BB} \quad (\text{B11})$$

In terms of minimal average payoffs in each state, the worst off population is better off in state \mathbf{z}^{AB} than in either of the other two. The worst off population is equally well off in \mathbf{z}^{AA} and \mathbf{z}^{BB} , but the better off population is better off in \mathbf{z}^{AA} than in \mathbf{z}^{BB} . The ordering of the leximin principle is thus $\mathbf{z}^{AB} \succ^R \mathbf{z}^{AA} \succ^R \mathbf{z}^{BB} .?$

Appendix C: Matching according to Myerson et al.

As mentioned earlier, Myerson, Pollock and Swinkels (1991) model viscosity in a different way. In their model, players have a probability \mathbf{b} of being matched with someone from their own population, and a probability $(1 - \mathbf{b})$ of being matched at random with the total population. Myerson et al defined a population (or kin group) as a set of agents sharing the same strategy, but let us explore the analogous idea when populations are characterized by different payoffs from interaction as studied above. Introducing matching according to Myerson et al into this framework, we get that a member of population 1 chooses strategy A if:

$$\mathbf{b} \frac{z_1}{N_1} + (1 - \mathbf{b}) \frac{z_1 + z_2}{N_1 + N_2} > \mathbf{a}_1 \quad (\text{C1})$$

Similarly, an agent from population 2 chooses strategy A if:

$$\mathbf{b} \frac{z_2}{N_2} + (1 - \mathbf{b}) \frac{z_1 + z_2}{N_1 + N_2} > \mathbf{a}_2 \quad (\text{C2})$$

These inequalities correspond to (2) and (3) in the main model. Let us rewrite the inequalities in the following way:

$$z_2 > \frac{\mathbf{a}_1(N_1 + N_2)}{(1 - \mathbf{b})} - \frac{N_1 + \mathbf{b}N_2}{(1 - \mathbf{b})N_1} z_1 \quad (\text{C3})$$

$$z_2 > \frac{\mathbf{a}_2 N_2 (N_1 + N_2)}{\mathbf{b}N_1 + N_2} - \frac{(1 - \mathbf{b})N_2}{\mathbf{b}N_1 + N_2} z_1 \quad (\text{C4})$$

The resulting inequalities correspond to (4) and (5).

Matching according to Myerson et al only makes sense for $\mathbf{b} \in [0,1]$, so we cannot study the case of dispersion.

Fluid interaction in this model occurs when $\mathbf{b} \rightarrow 0$. Inequalities (C3) and (C4) then reduce to

$$z_2 > \mathbf{a}_1(N_1 + N_2) - z_1 \quad (\text{C5})$$

$$z_2 > \mathbf{a}_2(N_2 + N_1) - z_1 \quad (\text{C6})$$

These two inequalities match (8) and (9) exactly. Long run stochastically stable states would thus be derived in exactly the same way, and proposition 2 holds for this kind of matching as well.

Completely viscous interaction means that $\mathbf{b} \rightarrow 1$. In this case (C3) and (C4) reduce to (10) and (11), and the long run stochastically stable states are as in proposition 3.

References

Anderlini, L. and A. Ianni (1996), "Path dependence and learning from neighbours", *Games and economic behavior*, vol. 13, no. 2, 141-177

Bergin, James and Barton L. Lipman (1996), "Evolution with state-dependent mutations", *Econometrica*, Vol. 64, No. 4, 943-956

Binmore, Ken and Larry Samuelson (2001), "Coordinated action in the electronic mail game", *Games and economic behavior*, 35, 6-30

Binmore, Ken, Larry Samuelson and Richard Vaughan (1995), "Musical chairs: Modeling noisy evolution", *Games and economic behavior*, 11, 1-35

Blume, Lawrence E. (1993), "The statistical mechanics of strategic interaction", *Games and economic behavior*, 4, 387-424

Blume, Lawrence E. (1995), "The statistical mechanics of best-response strategy revision", *Games and economic behavior*, 11, 111-145

Ellison, Glenn (1993), «Learning, local interaction, and coordination», *Econometrica*, 61, 1047-71

Ellison, Glenn (2000), "Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution", *Review of economic studies*, 67, 17-45

Frank, Robert H. (1988), *Passions within reason. The strategic role of the emotions*, W. W. Norton & Company, New York

Hahn, Sunku (2000), "The long run equilibrium in a game of 'battle of the sexes'", mimeo, Saitama University

Hamilton, W. D. (1964), "The genetical evolution of social behavior", *Journal of theoretical biology*, 7, 1-52

Harsanyi, John C. (1977), "Morality and the theory of rational behavior", *Social research*, v. 44, iss. 4, 623-656

Harsanyi, J. and R. Selten (1988), *A general theory of equilibrium in games*, Cambridge: MIT Press

Hehenkamp, Burkhard (2001), "Equilibrium selection in the two-population KMR model", *Wirtschaftstheoretische Diskussionsbeiträge*, no. 2001-01, Universität Dortmund

Kandori, Michihiro, George J. Mailath and Rafael Rob (1993), «Learning, mutation, and long run equilibria in games», *Econometrica*, vol. 61, no. 1, 29-56

Kandori, Michihiro and Rafael Rob (1995), "Evolution of equilibria in the long run: A general theory and applications", *Journal of economic theory*, 65, 383-414

Maynard Smith, J. and G. R. Price (1973), "The logic of animal conflict", *Nature*, 246, 15-18

Myerson, Roger B., Gregory B. Pollock and Jeroen M. Swinkels (1991), "Viscous population equilibria", *Games and economic behavior*, 3, 101-109

Oechssler, Jörg (1997), "Decentralization and the coordination problem", *Journal of economic behavior and organization*, vol. 32, 119-135

Rawls, John (1971), *A theory of justice*, The Belknap Press of Harvard University Press, Cambridge, Massachusetts

Skyrms, Brian (1994), "Darwin meets the logic of decision – correlation in evolutionary game-theory", *Philosophy of science*, vol. 61, no.4, 503-528

Skyrms, Brian (1996), *Evolution of the social contract*, Cambridge University Press, New York

Sugden, Robert (1995), "The coexistence of conventions", *Journal of economic behavior and organization*, vol. 28, 241-256

Young, H. Peyton (1993), "The evolution of conventions", *Econometrica*, vol. 61, no. 2, 57-84

Young, H. Peyton (1998), *Individual strategy and social structure – An evolutionary theory of institutions*, Princeton University Press, Princeton, New Jersey