Essays on

evolutionary game theory

by

Ivar Kolstad

A dissertation submitted for the degree of dr. oecon

To Sabrina

Acknowledgements

The process of writing a doctoral dissertation has its highs and lows, and violent swings between them. The excitement of a new idea, the frustrating process of expressing it in analytical terms, the immense satisfaction at proving it mathematically replaced by utter dejection as the proof unravels when subjected to closer scrutiny, the drudgery of coming up with a new approach and redrafting yet another version of a paper, and the relief as a scientific paper is finally completed containing an idea perhaps only vaguely related to your original one, make writing a dissertation akin to some kind of mental action sport. On the whole, though, I have enjoyed doing it, the final high of completing a dissertation dims the memory of an arduous few years. I would not recommend it to a manic depressive, however.

Though writing a dissertation can be a solitary endeavour, it cannot be done in total isolation. A number of individuals have made essential contributions to the process of completing this dissertation, and deserve my most heartfelt appreciation. First and foremost, my supervisor Bertil Tungodden has been crucial to the completion of this thesis. Succeeding my previous supervisor Terje Lensberg in January 2001, when time was short and completed papers few, Bertil helped shape and focus my ideas and provided the encouragement needed for me to stay the course. His combination of enthusiasm and analytical skills I find unmatched, to wax effusive: Everyone should have a supervisor like Bertil.

I am also highly indebted to the two other members of the supervising committee. Gaute Torsvik has contributed valuable and to-the-point comments on my work as it has developed. Oddvar Kaarbøe has provided essential feedback in particular on the technical aspects of the thesis. Both of them have contributed time and energy to this project, beyond what I could have reasonably asked or hoped for.

I take this opportunity to thank the Department of Economics at the Norwegian School of Economics and Business Administration, for giving me the opportunity to write this thesis. I also thank my friends and colleagues at the Department for the fun and stimulating time I spent there. In particular, I thank Turid Elvebakk and Dagny Kristiansen for their help on a million practical matters. Thanks are also due the Ethics Programme of the Research Council of Norway, for the opportunity to attend their terrific course series on ethics. Though this

thesis is less on ethics than I had originally intended, my association with the Ethics Programme has significantly deepened my understanding of science. I would also like to express my gratitude to Alexander Cappelen for introducing me to the Ethics Programme and for his comments on my work. While doing my conscientious objector's service at the Library at the Norwegian School of Economics, I was awarded time off to start the work that led up to this thesis, for which I am most grateful. Though this service is a flagrant confiscation of time, courtesy of the Norwegian state, the Library was a good place to serve. As my thesis neared completion, I am grateful to my current employer, Chr. Michelsen Institute, for giving me the time needed for the final touches.

On a personal note, I thank my Mom and Dad, Marit and Kjell, for their continued support through the process of writing this dissertation. Thanks also to Bjørn Kvist for giving me a boost at critical junctures, and for being a great friend. Last but by no means least, I am extremely grateful to Beate and Sabrina for their love and patience during these last few years.

Bergen, July 2002
Ivar Kolstad

Contents

Introduction

Evolutionary game theory has provided a fresh perspective on the matter of equilibrium selection in games. Since the Nash equilibrium concept does not necessarily provide a unique prediction of the outcome of a game, finding a solution concept with stronger predictive powers has been a major task of game theory. In the words of Sugden (2001, p. 115), "the Holy Grail is a solution concept which, for every game, picks out one and only one combination of strategies as the solution". Though falling somewhat short of the Holy Grail, key solution concepts proposed in evolutionary game theory do sharpen predictions, in what appears to be a less *ad hoc* manner than some of the attempts of classical game theory. This sharpening of predictions comes at the expense of some degree of rationality, in a way that might or might not bring theory closer to how human beings make their choices.

The two essential elements of evolutionary models are adaptation and mutation.[1] Basically, then, evolutionary game theory views decision making as analogous to a biological process of evolution. Adaptation captures an idea that agents imitate the actions of others, or choose their best action given how others have been acting. Mutations represent idiosyncrasies in behaviour, such as errors or experimentation, by which agents diverge from the process of adaptation. An initial solution concept proposed by Maynard Smith and Price (1973), that of evolutionary stability, focuses only on whether an equilibrium is robust to the invasion of a small number of mutants. Though excluding some Nash equilibria, the concept of evolutionary stability does not select among strict Nash equilibria, for instance those of a standard 2x2 coordination game.

By contrast, the solution concept suggested by Foster and Young (1990), stochastic stability, captures the outcome in the very long run when each agent has a small chance of mutating at any time.[2] There is thus a small probability that many agents mutate simultaneously. This solution has a much sharper predictive capacity than evolutionary stability, for instance facilitating selection between the two strict Nash equilibria of a coordination game (Young, 1993). However, for more complex games, there is a possibility that the long run outcome is a

---

[1] The matching of agents is important for the outcome, but is here subsumed under adaptation.

[2] The term long run stochastic stability, coined by Ellison (2000), is used henceforth instead of stochastic stability, to keep terms consistent.

limit cycle, where the process cycles between different strategy combinations forever. A unique solution is thus not offered for every game possible.

This thesis examines the nature of the strategy combinations, represented in evolutionary models as states of play, selected through evolutionary dynamics. In the standard models of Young (1993) and Kandori, Mailath and Rob (1993), the very long run outcome in a 2x2 coordination game entails play according to the risk dominant equilibrium. In other words, in the very long run the models predict that all agents play strategies consistent with an equilibrium that is not necessarily Pareto optimal. Ellison (1993) extents this result to the case where agents interact locally.

From various angles, the first three essays of this thesis considers ways in which evolutionary selection leads to states in which all agents do not pursue the same strategies, where there is no single conventional strategy everywhere and for everyone. In other words, these essays explore the possibility of a coexistence of conventions in evolutionary models. The models of essay 1 and 3 consider only adaptive dynamics, while essay 2 includes persistent mutations in the manner of Young and Kandori, Mailath and Rob. Essay 4 adds balance to the thesis by discussing the interpretation of the evolutionary concepts of adaptation and mutation in terms of human decision making, and points out how differences in the interpretation of mutations can lead to the selection of different equilibria.

In essay 1, "What constitutes a convention? Implications for the coexistence of conventions", a model due to Sugden (1995) is reviewed. In this model a coordination game is repeatedly played at different locations in a continuous social space. Players receive noisy signals of the location of their game, and thus adapt to past play in a region around the actual location of their game. Play in one location is thus influenced by the history of play in other locations, which makes it possible for conventions to spread across locations. Sugden suggests that in a model of this kind, there can be a stationary state of convention coexistence only if interaction is non-uniform across social space, i.e. only if the game is played more frequently in some locations than in others. Essay 1 argues, however, that this result is based on a definition of conventions focussing on the expectations rather than the actions of players. An alternative definition of conventions is suggested, which permits convention coexistence when interaction is uniform.

Essay 2, "Viscosity and dispersion in an evolutionary model of learning", presents a model of adaptation and mutation in which members of two distinct populations preferring different equilibria interact. Interaction ranges from complete viscosity, where agents interact only with members of their own population, to complete dispersion, where agents interact only with members of the other population. The idea of dispersive interaction is a conceptual extension of the viscous-fluid continuum used by Myerson, Pollock and Swinkels (1991).

With complete viscosity, the long run stochastically stable state has each population playing its preferred equilibrium. In a sense, this results matches that of Kandori, Mailath and Rob (1993). With complete dispersion, the long run stochastically stable states match those of Hehenkamp (2001). The population most difficult to dissuade from playing its preferred equilibrium, imposes this equilibrium on the other population, unless the other population is sufficiently much larger, in which case its preferred equilibrium is played by both populations. When interaction is fluid, which means that agents interact as frequently with any agent from their own population as with any agent from the other population, there is a possibility of convention coexistence, where each population plays according to its preferred equilibrium.

When different populations with conflicting interests interact, the Pareto principle provides insufficient guidance on which states are preferable from a welfare point of view. To evaluate the welfare properties of the long run stochastically stable states under dispersive, fluid and viscous interaction, utilitarian and Rawlsian measures are therefore used. While the long run stochastically stable state is preferable from a utilitarian and Rawlsian perspective when interaction is completely viscous, the same does not necessarily hold when interaction is completely dispersive or fluid. In other words, whether evolutionary selection and normative criteria diverge, depends *inter alia* on the degree of viscosity in interaction.

In the literature on urban labour markets, neighbourhood effects and worker mobility have been used as explanations for the pattern of employment in cities. In essay 3, "Social origins of a work ethic: Norms mobility and urban unemployment", an analytical framework for studying the joint impact of these two factors is constructed. Neighbourhood effects are modelled by letting workers' employment decisions be influenced by the decisions of the agents closest to them in a simply construed social space. Mobility takes the form of workers trading locations in social space, similar to the neighbourhood segregation model of Schelling

3

(1971). However, a range of different rules for when two workers trade locations are considered, some of which assume that workers move where the unemployment situation is better, while others assume that the composition of neighbourhoods in terms of the education of workers determines mobility. Each mobility rule produces a different set of long run outcomes in terms of unemployment levels and segregation of workers with different levels of education. However, full employment and total segregation is found to be a long run outcome for most of the mobility rules. The results of the model are sensitive to certain of its specifications, yet it provides an illustration of how evolutionary models can be applied to the study of local labour markets.

The fourth and final essay, "Evolution with endogeneous mutations", asks some fundamental questions about whether an evolutionary model of adaptation and mutation provides an adequate representation of decision making. One aspect of this is how to interpret mutations. Bergin and Lipman (1996) prove that selection in the models of Young (1993) and Kandori, Mailath and Rob (1993) depends critically on the assumption that mutations are equiprobable in all states. If any variation in mutation probabilities is allowed, any equilibrium can be selected. What we take mutations to represent, and consequently how their probabilities can be assumed to vary across states, thus becomes vital for the predictions of an evolutionary model.

Van Damme and Weibull (1998) provide one approach to endogenizing mutations, by stating that if we interpret mutations as mistakes, then agents would want to avoid mistakes that are associated with larger payoff losses. Conversely, if mutations capture experimentation, agents experiment less is states where payoffs are high. Van Damme and Weibull prove that if agents focus on payoff losses when determining their mistake probabilities, and reducing these probabilities is costly, then the selection results of Young and Kandori, Mailath and Rob are upheld. However, it is no less intuitively appealing to make the reverse claim about mistakes and experiments; that agents avoid experimenting when the payoff losses thereof are large, and that agents try to make fewer mistakes in states where payoffs are high. Essay 4 suggests one way in which to model mistakes in this manner, and proves that the result of van Damme and Weibull can thereby be reversed.

References

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

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

Foster, D. and Young, H. Peyton (1990), "Stochastic evolutionary game dynamics", *Theoretical Population Biology*, 38, 219-232

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

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

Schelling, Thomas C. (1971), "Dynamic models of segregation", *Journal of mathematical sociology*, 1, 143-186

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

Sugden, Robert (2001), "The evolutionary turn in game theory", *Journal of economic methodology*, 8:1, 113-130

van Damme, Eric and Jörgen W. Weibull (1998), "Evolution with mutations driven by control costs", mimeo, Center for economic research, Tilburg University

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

What constitutes a convention?

Implications for the coexistence of conventions[*]

Ivar Kolstad[**]

March 2002

Abstract

A model of repeated play of a coordination game, where stage games
have a location in social space, and players receive noisy signals of the
true location of their games, is reviewed. Sugden (1995) suggests that
in such a model, there can be a stationary state of convention
coexistence only if interaction is non-uniform across social space. This
paper shows that an alternative definition of conventions, which links
conventions to actions rather than expectations, permits convention
coexistence when interaction is uniform.

Keywords: Convention; Coordination game; Equilibrium selection

---

Introduction

The old adage "When in Rome, do as the Romans" advises us to adopt the behavioural patterns of the people in whose presence we find ourselves. Whether of necessity or for pleasure, the proverb suggests that it is somehow advantageous to mimic the actions and mannerisms of our social surroundings. The proverb thus prescribes a change in behaviour as we move from one social sphere to another, say, from Rome to Paris, from the cinema to the theatre, from the mail room to the board room and so on. In chameleon-like fashion, we should change our language from Italian to French in the first case, our code of dress from casual wear to formal wear in the second, the formality of our speech from less to more in the third case.

How do the behavioural patterns we observe come into existence in the first place? Why do individuals who find themselves in similar social surroundings often share a common way of doing things; speaking the same language, wearing similar clothes, employing common means of exchange such as money, and using common standards of measure? The evolutionary approach suggests that these *conventions* form through the repeated interaction of individuals. The basic idea is as follows. If there is some advantage to acting in a manner similar to others, and if a shared history of actions is used as a gauge of how others will act, then present actions will reflect past actions, and over time a pattern could form in which one way of acting becomes dominant. According to this line of thinking, then, interdependent individual actions form a collective pattern, a convention, through the indirect observance of precedent. In Rome you speak Italian because it eases communication with others who, based on what you know about Rome, are more likely to speak Italian than any other language.[1]

Evolutionary models of learning adopt the above perspective in one form or another, to explain how agents who adapt to or learn from the actions of their environment, can end up using the same type of action. In the model of Kandori, Mailath and Rob (1993), agents adapt by choosing a best reply to the distribution of actions in the

---

[1] According to this perspective, agents coordinate by watching the past actions of each other, rather than by communicating with each other. This approach is therefore better suited to situations where communication is costly relative to the cost of switching between different actions. In the current example, we are therefore implicitly assuming that travellers are sufficiently well versed in different languages as to make communication on which language to use a waste.

population in the preceding period. Young (1993, 1998), on the other hand, assumes that agents observe a limited sample of the actions taken in a given number of preceding periods, and choose a best reply to this sample. Both Kandori et al and Young do in fact have even greater ambitions than showing that one convention or another will arise through the adaptation of agents, they also want to establish which convention will be chosen. To this end, they introduce a small probability of error into the strategy implementation of agents, and show that as this error probability grows arbitrarily small, one particular convention might be observed with near certainty in the very long run. For populations playing 2x2 coordination games, both Kandori et al and Young obtain the result that the convention thus selected entails play according to the risk dominant equilibrium of the game, as defined by Harsanyi and Selten (1988).[2]

The models of Kandori et al and Young are global interaction models, where an agent has a positive probability of interacting with any other agent in the population.[3] However, Ellison (1993) shows that the risk dominant equilibrium is also selected in a local interaction model where agents have fixed locations on a circle and adapt to the actions of a limited set of neighbours only, adaptation taking a form similar to Kandori et al. Judging from the results of this model, local interaction seems to leave little room for differences in conventions across locations. These results are, however, due to the persistent errors in the strategy implementation of agents. There exist local interaction models without this particular feature that do permit convention coexistence in simple coordination games. Anderlini and Ianni (1996) assume that errors only occur when agents attempt to use a different strategy than they did in the preceding period, which produces a non-ergodic dynamic process whose absorbing states do in some cases contain different strategies at different locations. In a model without implementation errors, Goyal and Janssen (1997) assume that agents can at some cost choose both strategies, thus always achieving coordination, and show that for intermediate cost levels, convention coexistence can be a stationary state.

---

[2] Note that for more complicated games, the models of Kandori et al and Young might differ in their predictions of the very long run outcome, as demonstrated by Jacobsen, Jensen and Sloth (2000).
[3] Sometimes the term uniform, rather than global, is used to describe interaction of this kind, see e.g. Ellison (1993). As the term uniform interaction is used in another context here, I choose the term global interaction to avoid confusion.

On the other hand, we have local interaction models which deem contagion of a particular strategy throughout a population likely. Blume (1995) shows that if there is spatial variation in the initial condition and randomness in the order in which agents revise their strategies, then we get coordination on the risk dominant equilibrium. Lee and Valentinyi (2000) similarly prove that if initially each agent has a positive probability of playing the risk dominant strategy and the population is sufficiently large, the risk dominant equilibrium is realized almost with certainty. In a more general setting, Morris (2000) shows that for any local interaction structure, there exists some contagion threshold, and coexistence is possible if agents do not choose to play according to the risk dominant equilibrium, whenever the probability with which their opponent does so is below this threshold.

A common feature of all the local interaction models discussed above, is that agents have fixed locations in some social space. The proverb "When in Rome, do as the Romans" suggests, however, that there is some manner of local interaction that these models do not properly address. The proverb advises a change in behaviour as we move from place to place, and we therefore need mobile agents to analyze social adaptation of this kind. Sugden (1995) presents a model in which interaction is global in the sense that agents have a chance of meeting all other agents in a population, yet local in the sense that each meeting has a random location in a social space. Agents are matched repeatedly at varying locations to play a coordination game, and adapt to the past history of play at the location at which they find themselves. To make the evolution of play at different locations interdependent, agents are assumed to have an imperfect understanding of their current location. Sugden concludes that in this model, a coexistence of conventions is possible if and only if the frequency of interaction across social space varies in a certain way. If there is no variation, i.e. if interaction is uniform across locations, there can be no coexistence of conventions.

In this paper, I argue that we can expand Sugden's coexistence result to include the case of uniform interaction without unduly altering the fundamental structure of his model. Specifically, the result that coexistence is impossible under uniform interaction hinges on a definition of conventions that focuses on the expectations rather than the actions of agents. I show that if we adopt a more reasonable definition based on what agents do rather than what they expect others to do, a coexistence of

conventions is possible even if there is no spatial variation in the frequency of interaction. With a different and weaker definition of conventions, we thus strengthen the case for coexistence initially made by Sugden.

The paper is organized as follows. In the next section, Sugden's model of convention formation is outlined. In section three, his definition of a convention is reviewed, and his results on coexistence are derived, with a detailed look at why coexistence is impossible when interaction is uniform. In section four, the main reasons for challenging Sugden's definition of a convention are given, and an alternative definition is presented. Section five shows that under this alternative definition, convention coexistence is possible in the case of uniform interaction. A final section raises the important point of robustness of the coexistence outcomes.

Sugden's spatial model of convention formation

To intuitively understand the model presented by Sugden (1995), let us use a simple example to sketch the situation facing the agents of the model and the manner in which they behave. Suppose you have been invited to a party, and have to decide what to wear. You want to blend in with the other guests, so the first thing you do is form an opinion of who else is likely to come, and what they are likely to wear. You know the identity of your hosts, and who they are likely to invite, but this still leaves you with only an imprecise idea of the mix of people you will face at the party. Suppose that in the past you have observed that the way people dress depends on certain of their personal characteristics, let us say their age. You then combine your imprecise understanding of the average age of the people invited with your expectation of how people of that age will dress, and choose the garment that best matches the resulting estimate.

Now, the way you and others dress at this party, influences the ideas you and others have about what people wear what kind of clothes. So the way you and others dress for the next party with a similar mix of people, will be influenced by what people wear at this party. Moreover, since everyone has an imprecise idea of the mix at this party, they might adapt to different ideas of the average party-goer. The garments

normally worn at parties with one mix of people might therefore influence the garments worn at parties with a different mix of people. The kind of question Sugden's model is designed to answer, is whether this will lead to a sitation in which the code of dress is the same for all parties regardless of the age of those invited, or whether we can have a stable situation in which dress codes vary with age.

Sugden frames this basic idea in terms of a model in which agents are repeatedly matched to play a coordination game, where each stage game has a random location in a social space. The players do not know the exact location of their game, instead, they receive a signal of their location which is close to but not necessarily spot on their actual location. The players have a common understanding of the past pattern of play at the various possible signals, and they are able to compute a probability distribution for their opponent's signal given the signal they themselves have received. Based on this information, each player calculates the probability with which his opponent will choose either strategy, and chooses the strategy which maximizes his expected payoff.

In more formal terms, consider a large population of identical agents. In each period, a pair of agents is drawn at random from the population to play the following game

<div align="center">

Player 2

		A	B
		A	B
Player 1	A	a,a	c,d
	B	d,c	b,b

</div>

(G1)

We assume that $a > d$ and $b > c$, which makes (G1) a coordination game with two Nash equilibria in pure strategies, (A,A) and (B,B). Moreover, we assume that

<div align="center">6</div>

$a - d > b - c$, which implies that (A,A) is the risk dominant equilibrium as defined by Harsanyi and Selten (1988).[4]

Players choose the strategy that maximizes their expected payoff. From the above payoff matrix, we see that a player is indifferent between strategies A and B if the probability with which his opponent chooses A is α, where

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

For probabilites greater than α, players prefer strategy A. And for probabilities lower than α, they prefer strategy B. Note that since strategy A is risk dominant, $\alpha < 0.5$. This implies that players may choose strategy A even if the probability with which their opponent does so is below 50%. In a sense, the players are more easily persuaded to choose strategy A than strategy B, as the former strategy requires a lower probability that their opponent acts similarly.

Each game is assigned a random location in a social space. Social space is continuous, consisting of all points on the real line from 0 to 1. The location of a game is a random variable y in the interval $[0,1]$. The probability that a game is assigned to a location less than or equal to y, is represented by $F(y)$. The corresponding density function $f(y)$, which denotes the frequency of interaction at each location, is assumed to be continuous, with $f(y) > 0$ for all $y \in [0,1]$. In other words, all points on the real line from 0 to 1, have a positive probability of being host to the game in any given period.

---

[4] 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.
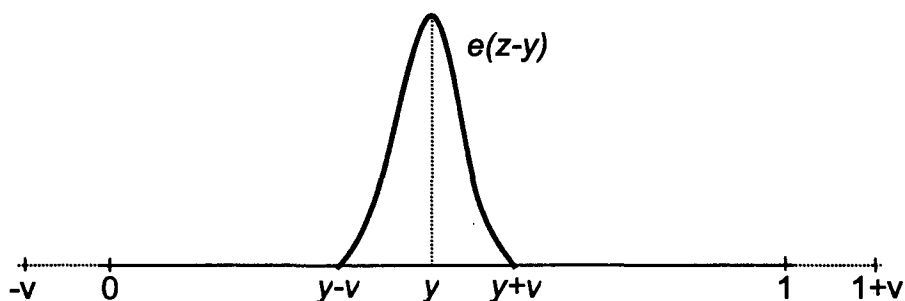
Figure 1. Social space, location of games and distribution of signals

Each player receives a signal z of the location of the game. Figure 1 depicts the probability distribution of signals, given the location y of a game. The signal of a player never falls more than a small distance v from the true location y of a game. Signals closer to y do not have lower probabilities than signals further away from y, and signals equally far from y are equally probable. Formally, the distance between a signal and the true location of a game, $z - y$, is a random variable with density function $e(z - y)$. For some small positive v, $e(z - y) > 0$ if and only if $z - y \in (-v, v)$. The density function $e(z - y)$ is continuous, symmetric around a mean of 0, non-decreasing in the interval $[-v, 0]$ and non-increasing in the interval $[0, v]$. The signals of the two players are assumed to be stochastically independent. Note that if the game is played at a location less than v from 0 or 1, players may receive signals lower than 0 or higher than 1. The signal space is thus wider than the social space, and contains all points in the interval $[-v, 1 + v]$.

Knowing the distribution of games in social space, and the distribution of his signal around the true location of a game, a player can compute a probability function for the true location of a game given his own signal. A player is also aware of the distribution of his opponent's signal around the location of the game, and can calculate a probability function for the signal of his opponent given his own signal. Let $H(x|z)$ be a cumulative probability function which states the probability that the signal of his opponent is less than or equal to x, given his own signal z. The corresponding density function $h(x|z)$ thus represents the probability that the other player receives signal x when a player receives signal z. Note that since a player's signal is at most a distance v above or below the location of a game, the signals of two players are at most $2v$

apart. The function $h(x|z)$ is thus positive if and only if $x \in (z - 2v, z + 2v)$, i.e. in an interval of width $4v$.

A player expects his opponent to conform to past behaviour at the signal his opponent has received. There is thus a kind of bounded rationality at play, where players expect others to make simple decisions based on their signals, while they all actually let more complicated evaluations of their opponent's actions determine their own. A state of play function $g_t(z)$ denotes, for all feasible signals $z \in [-v, 1 + v]$, the probability that a player receiving signal z at time t will choose strategy A. This function captures past play in the sense that it increases for signals at which A is chosen, reaching a maximum of 1 after a finite number of periods in which A is played repeatedly. For signals at which B is chosen, the state of play function decreases and reaches a minimum of 0 after a finite number of periods where B is repeatedly chosen. We can thus define the state space Ω as

$$\Omega = \{g(z) : 0 \le g(z) \le 1, \forall z \in [-v, 1 + v]\} \tag{2}$$

Weighing the probabilities $g(.)$ that strategy A is chosen at different signals with the probabilities $h(.|z)$ that an opponent receives these various signals, a player arrives at a probability that his opponent chooses A given his own signal z. Formally, the probability $\pi(z)$ that your opponent will choose strategy A when you receive signal z is

$$\pi(z) = \int h(x|z)g(x)dx \tag{3}$$

Maximizing expected payoffs, a player thus chooses strategy A if $\pi(z) > \alpha$, and B if $\pi(z) < \alpha$.

The choices of the players in turn feed into the state of play function, and potentially influence play in future periods. We are interested in the stationary states of the system, which can be defined as follows.

9

DEFINITION 1

A state $\bar{g}(z) \in \Omega$ is a stationary state if and only if the following holds:

If $g_t(z) = \bar{g}(z)$ then $g_s(z) = \bar{g}(z)$ for all $s > t$ and all $z \in [-v, 1+v]$

In other words, we are at a stationary state when the state of play function stays the same forever after we have reached this state.

Uniform interaction and coexistence

In the context of the above model, Sugden suggests that a convention is realized at some signal z (or as he puts it, universally followed at z), when two conditions are met. Firstly, a player receiving signal z must observe the convention with certainty. Secondly, the opponent of a player receiving signal z must observe the convention with certainty. In other words, we have an A-convention when for some signal z, both $g(z)$ and $\pi(z)$ equal one. Similarly, we have a B-convention when for some signal z both $g(z)$ and $\pi(z)$ equal zero. Finally, to have a coexistence of conventions we must have an A-convention at some signal, a B-convention at some other signal, and this state of play must be a stationary state.

Interaction is uniform when the frequency of interaction at each location is the same, $f(y) = 1$ for all y. In any given period, then, a game has an equal chance of being assigned a location anywhere on the real line from 0 to 1. With uniform interaction, and given the above definition of a convention, no state in which there exist two different conventions can be a stationary state, as implied by the following proposition.

PROPOSITION 1

Suppose $f(y) = 1$ for all y.

If $\hat{g}(.)$ is a state of play function with the following properties for some signals $z', z'' \in [-v, 1+v]$:

i) $\hat{g}(z) = 1$ for all $z \in (z', z'')$

ii) $\hat{g}(z) \neq 1$ for some $z \notin (z', z'')$

iii) $z'' - z' \geq 4v$ or $[z'' - z' = 2v$ and $z' = -v]$ or $[z'' - z' = 2v$ and $z'' = 1 + v]$.

Then $\hat{g}(.)$ is not a stationary state.

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

Proposition 1 rules out coexistence in the following way. For an A-convention to exist, there must be some signal z at which a player is certain that his opponent chooses A, i.e. $\pi(z) = 1$ for some $z \in [-v, 1+v]$. From equation (3), we see that this implies that A must be played with certainty, $g(.) = 1$, at all signals his opponent has a positive probability of receiving. Since the signal of his opponent can fall anywhere within a distance of $2v$ from his own, this means that A must be played with certainty in a region of width $4v$.[5] However, from proposition 1 we see that if A is played with certainty in a region of this width, we are not at a stationary state if somewhere else A is not played with certainty. With uniform interaction, then, a state in which there is an A-convention somewhere but not everywhere, is not stationary.

There is a simple intuitive reason for this result. Consider a state in which A is played with certainty at all signals between z' and z'', where z' and z'' are at least $4v$ apart. A player receiving a signal at the edge of the region, say at z', calculates a probability distribution $h(.|z')$ for his opponent's signal which can be illustrated as follows

---

[5] If we are considering a signal at the edge of signal space, $z = -v$ or $z = 1 + v$, then the region need only be $2v$ wide. This is reflected in part iii) of the proposition.
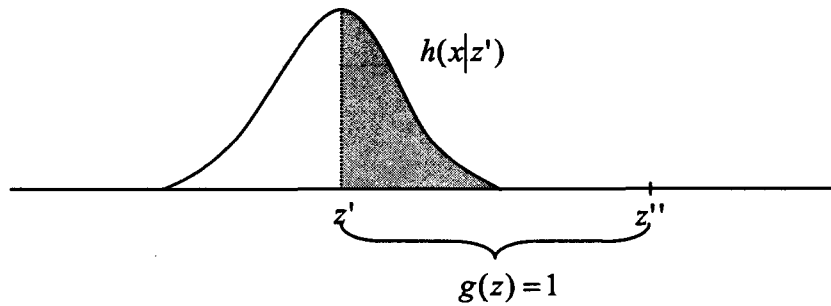
11

Figure 2. Probability distribution of opponent's signal at edge of A-region

When interaction is uniform, $h(\cdot|z')$ has a nice symmetric form around z'. For a player receiving signal z', half the bulk of $h(\cdot|z')$ falls within the region in which A is played with certainty. In other words, for a player receiving a signal at the edge of a region where A is played with certainty, the probability that his opponent receives a signal inside the region is 0.5. From equation (3), this means that the probability with which his opponent plays A, $\pi(z')$, is at least 0.5. Since $\alpha < 0.5$ and thus $\pi(z') > \alpha$, the player at the edge therefore strictly prefers strategy A. Moreover, by continuity, the same is true for a player receiving a signal ever so slightly to the left of z'. The state of play function $g(z)$ therefore increases for signals at the lower edge of the region. A similar argument tells us that players receiving a signal at the upper edge of the region z'', also strictly prefer strategy A. The region in which A is played with certainty thus expands in both directions, and keeps doing so until A is played with certainty throughout signal space.

If the distribution of games is uniform, then, the only stationary state which contains an A-convention is a state which contains only an A-convention. Any state in which there is both an A-convention and a B-convention eventually collapses as the space commanded by the A-convention gradually expands. In a sense, the definition of an A-convention used by Sugden, requires a region where A is played which is above the critical size at which conventions are able to coexist when interaction is uniform.

When interaction is not uniform, however, two conventions can stably coexist. Note that if there are variations in the frequency of interaction across locations, then the probability distribution depicted in figure 1 need not be symmetric. If the variations are of a certain order, a player getting a signal at the edge of a region where A is played with certainty, might then calculate the probability of his opponent's getting a

signal within the region as being equal to α. In this case, if B is played with certainty to the other side of his signal, the player is indifferent between strategies A and B. Neither region thus expands, and we can have a stationary state with coexistent conventions.

What constitutes a convention – actions speak louder than expectations

The definition of convention existence used by Sugden prevents conventions coexistence when interaction is uniform. A convention only exists if there is some signal where a player can be sure that his opponent observes the conventional strategy. And if there is a signal where a player can be sure his opponent chooses the risk dominant strategy A, then no other convention is stable. The element of certainty in expectations used in the definition of conventions is thus what kills coexistence. It is therefore fitting to ask whether it is reasonable to put so much emphasis on expectations when defining conventions.

Intuitively, the definition used by Sugden seems to include more than a definition of conventions need include. A commonly cited definition of conventions due to Lewis (1969) suggests that "a convention is a pattern of behavior that is customary, expected and self-enforcing". A convention denotes a behavioural pattern, a regularity in the actions taken by a set of agents. The basic units that form a convention are thus the actions of individual agents, not their expectations. Expectations do form a basis on which to choose actions, but it is regularities in the actions chosen that are of interest, not regularities in expectations. Expectations are only of derivative importance, in perpetuating the regularities in actions needed for a convention to persist.

This is certainly the view taken in other parts of the evolutionary literature. Conventions are defined on the basis of state of play vectors, matrices or functions, and expectations are an element of what keeps conventions in place (see e.g. Young, 1993, 1996). A conventional definition of conventions would thus focus on strategies, and impose no stricter requirements on expectations than that they perpetuate strategy choices. In the context of Sugden's model, this means that requiring players to be absolutely certain their opponents choose a particular strategy, is too strict a demand

to impose in a definition of conventions. For a player to do what has generally been done at his signal, he need only deem it sufficiently probable that his opponent chooses similarly. A more reasonable definition would thus substitute an idea of sufficient probability in expectations for that of absolute certainty.

At a signal where A is generally played, $g(z) = 1$, players need only expect their opponents to play A with probability $\pi(z)$ above α, to keep playing A at this signal. Where B is generally played, $g(z) = 0$, we need only $\pi(z) < \alpha$ for B to continue being played. We thus arrive at an alternative definition of conventions which does not include more than such a definition need include: There is an A-convention if for some signal z, A is played with certainty by players at this signal, $g(z) = 1$, and A is the optimal choice for a player receiving this signal, $\pi(z) > \alpha$. Similarly, there is a B-convention if for some signal z, $g(z) = 0$ and $\pi(z) < \alpha$. In accordance with Sugden's idea of coexistence, we have a coexistence of conventions when an A-convention exists at some signal, a B-convention exists at another signal, and this state is a stationary state.

Uniform interaction and coexistence revisited

If we adopt the alternative definition of a convention, states of convention coexistence can be stationary states when interaction is uniform, as the following proposition implies.

PROPOSITION 2

Suppose $f(y) = 1$ for all y.

If $\hat{g}(.)$ is a state of play function with the following properties for some signals $z', x \in (v, 1 - v)$:

i) $\hat{g}(z) = 1$ for all $z \in (z', x)$

ii) $\hat{g}(z) = 0$ for all $z \notin (z', x)$

Then there exists some signal $x = z''$ for which $\hat{g}(.)$ is a stationary state.

The proposition says that with a uniform frequency of games in social space, a state in which strategy A is played with certainty within some region, and B is played with certainty everywhere else, is a stationary state provided the region where A is played is of a certain width. Clearly, such a state meets the requirements of coexistence under the above alternative definition. By using a more reasonable definition of convention coexistence, we thus get a result which is stronger in the sense that it deems coexistence possible even if interaction is uniform.

The intuition behind the proposition is as follows. Imagine that we are in a state $\hat{g}(.)$ where A is played with certainty in some region of signal space z' to x, and B is played with certainty at all signals outside this region. Consider a player who receives a signal at the lower end z' of the region where A is played with certainty. The probability distribution of his opponent's signal can be illustrated as follows
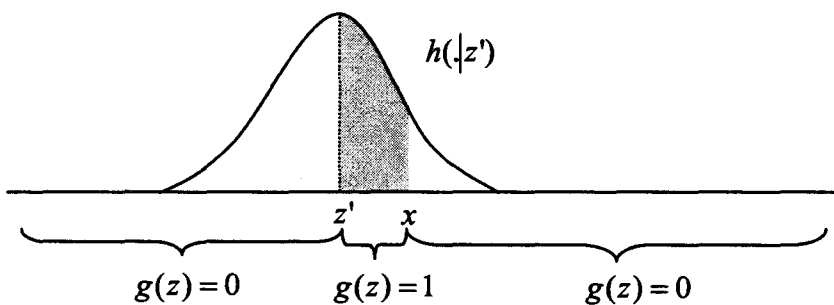
Figure 3. Probability distribution of opponent's signal at border between A- and B-playing regions.

For a player receiving signal z', the shaded area represents the probability that his opponent gets a signal in the region where A is played with certainty. The location of x determines how large this probability is. The further away x is from z', the larger is this probability, with a maximum of 0.5 if x is a distance $2v$ or more from z'. Due to the fact that A is played with probability one between z' and x, and probability zero elsewhere, the shaded area also equals $\pi(z')$, the probability that A is played by the opponent of a player receiving signal z'. Now, imagine that we first let x be a distance $2v$ above z', which implies $\pi(z') = 0.5$. If we start sliding x towards z', $\pi(z')$ decreases, and due to the continuity of $h(\cdot|z')$, at some point $x = z''$, we get

15

$\pi(z') = \alpha$. The player at the border z' between two regions where A and B is played, is now indifferent between the two strategies. Due to the fact that $h(\cdot|z')$ is symmetric and the same for all signals when interaction in uniform, the player at the other border z'' is also indifferent between A and B.

For a player receiving a signal inside the region z' to z'', A is the optimal strategy. The reason is that if we place the centre of the $h(\cdot|z)$-curve anywhere between z' and z'', the weight this curve puts on the region in which A is played, is greater than if the curve centred on one of the edges of that region. In other words, the probability $\pi(z)$ that your opponent plays A is greater for signals inside the region than at its edges, and we thus have $\pi(z) > \alpha$ for all signals between z' and z''. Moreover, a similar argument tells us that $\pi(z) < \alpha$ for signals ouside this region, and the optimal choice for a player receiving such a signal is strategy B. Consequently, for $x = z''$, state $\hat{g}(.)$ is a stationary state.

Concluding remarks

Sugden (1995) argues that in a model where agents are matched repeatedly to play a coordination game, where games have a location in a social space, and players do not know the exact location of their game, conventions can coexist only if the frequency of interaction varies across locations. We might interpret this as saying that if everyone acts according to the rule "when in Rome, do as the Romans" or "when at a party, dress the age of the other party-goers", the possibility that over time codes of conduct or dress would remain different in different surroundings, is limited. However, this paper argues that Sugdens's definition of a convention focuses too much on the expectations of the players rather than their actions. If instead we adopt a definition where their actions are the key element, we get the result that coexistence is possible even if interaction is equally frequent at all locations in social space.

The stationary state of coexistence established above is, however, only one type of stationary state. The state in which A is played with certainty across the space of

signals, or the one in which B is played with certainty at all signals, are also stationary states. Moreover, as the state of coexistence can crumble if the state of play function is perturbed only slightly, this state might be less robust to different kinds of perturbations than other stationary states. Ideally, we ought to test the different states for robustness. One way to do so is to follow the approach of Young (1993) and Kandori, Mailath and Rob (1993), and introduce a small probability that agents choose their inoptimal strategy. However, the processes studied by Young and Kandori et al are defined on a finite and discrete state space, and as the model studied here has a continuous state space, their algorithm for identifying robust states is not applicable in this case. A different way of assessing robustness which might be more attuned to the present context, is to use the approach of Blume (1995) and see whether variations in initial conditions make some states more likely than others.

Finally, a note on the dispersion of signals in the above model. When the dispersion of signals v is small and interaction uniform, one can have a string of correctly sized segments playing A in a social space where B is otherwise played. As long as these segments are at least $2v$ apart, they do not exert a joint influence strong enough to alter the state of play. A general lesson from the above model is therefore that the more certain players are of their true location, the greater can the variation in conventions across social space be. Conversely, the greater is the confusion about one's correct location, the less variation in conventions is possible. In the extreme, if v is large in comparison to social space, there can be no coexistence of conventions, even by the alternative definition. The impact of the dispersion of signals on the maximum number of regions with different conventions has to do with the influence play in one location has on play in another. The more confusion about true locations, the greater is the range of locations that influences play in any one location.

Appendix: Proof of propositions

PROOF OF PROPOSITION 1:

This proposition is proved by reference to theorem 2 in Sugden (1995), which basically establishes the following:

Let $\tilde{g}(.)$ be a state of play function which has the following properties for some signals $z', z'' \in [-v, 1+v]$:

i) $\tilde{g}(z) = 1$ for all $z \in (z', z'')$

ii) $z''-z' \geq 4v$ or $[z''-z' = 2v \text{ and } z' = -v]$ or $[z''-z' = 2v \text{ and } z'' = 1+v]$

iii) $1 - H(z'|z') \geq \alpha$

iv) $H(z''|z'') \geq \alpha$

If at any time t, the state of play function is $\tilde{g}(.)$, then

a) For all $s > t$, $g_s(z) = \tilde{g}(z) = 1$ for all $z \in (z', z'')$

b) If $1 - H(z'|z') > \alpha$, then for some finite $s > t$, $g_s(z) = 1$ for some $z \in [-v, z')$

c) If $H(z''|z'') > \alpha$, then for some finite $s > t$, $g_s(z) = 1$ for some $z \in (z'', 1+v]$

In state $\tilde{g}(.)$, strategy A is played with certainty in a region of width at least $4v$ (or at least $2v$ at the edges of signal space), and players receiving signals at the edge of this region perceive the chance of their opponent's receiving a signal inside the region as at least α. Part a) then says that in all later periods, strategy A will keep being played with certainty at all signals inside the region. Part b) and c) say that if a player receiving a signal at either edge of the region deems the probability of his opponent's getting a signal inside the region as strictly higher than α, then the region in which A is played with certainty will expand at this edge.

With uniform interaction, $f(y) = 1$ for all y, for a player receiving signal z, the probability that the signal of his opponent is above z is 0.5, as is the probability that the signal is below z. Due to the fact that $\alpha < 0.5$, we thus have

$1 - H(z|z) = H(z|z) > \alpha$ for any signal $z \in [-v, 1+v]$. By iterated application of the above theorem, this means that a region of width 4v in which strategy A is played with certainty, will expand until A is played with certainty at all signals. By definition 1, this implies that a state $\hat{g}(.)$ fitting the description of proposition 1 is not a stationary state if interaction is uniform.□

PROOF OF PROPOSITION 2:

When $f(y) = 1$ for all y, $h(x|z)$ has the following properties:

i) $h(x|z)$ is symmetric around mean z

ii) $h(x|z)$ is non-decreasing in the interval $[-v, z)$

iii) $h(x|z)$ is non-increasing in the interval $(z, 1+v]$

iv) $h(z'+a|z') = h(z''+a|z'')$ for all $a \in (-2v, 2v)$ and $z', z'' \in (v, 1-v)$

Consider the interval (z', x). Clearly:

$$\pi(z') = H(x|z') - H(z'|z') \text{ for } x \geq z' \tag{A1}$$

Since $H(x|z)$ is continuous in x, $\pi(z')$ is continuous in x, and has a maximum value of 0.5 and a minimum value of 0. There thus exists some point $x = z''$ at which $\pi(z') = \alpha$.

Moreover by properties i) and iv):

$$\pi(z'') = H(z'|z'') - H(z|z'') = H(z'|z') - H(z|z') = \pi(z') \tag{A2}$$

19

Finally, when $f(y) = 1$:

$$\frac{\partial \pi(a)}{\partial a} = \frac{\partial \left[H(z'|a - H(z|a)\right]}{\partial a} = h(z|a) - h(z'|a) \tag{A3}$$

Thus, from properties i), ii) and iii), $\pi(a)$ is non-decreasing for $a \in \left(-v, \frac{z'+z''}{2}\right)$ and non-increasing for $a \in \left(\frac{z'+z''}{2}, 1 + v\right)$. Which implies:

$$\pi(a) \geq \pi(z') \text{ for all } a \in (z', z'') \tag{A4}$$

$$\pi(a) \leq \pi(z') \text{ for all } a \notin [z', z''] \tag{A5}$$

For a state of play function $\hat{g}(.)$ such as that of proposition 2, agents between z' and z'' continue playing A, agents below z' or above z'' keep playing B, and agents at z' or z'' are indifferent. All of which makes $\hat{g}(.)$ a stationary state by definition 1.□

References:

Anderlini, Luca and Antonella Ianni (1996), "Path dependence and learning from neighbors", *Games and economic behavior*, 13, 141-177

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

Goyal, Sanjeev and Maarten C. W. Janssen (1997), "Non-exclusive conventions and social coordination", *Journal of economic theory*, 77, 34-57

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

Jacobsen, Hans Jørgen, Mogens Jensen and Birgitte Sloth (2000), "KMR and Young precesses select different equilibria, two examples", mimeo, University of Copenhagen

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

Lee, In Ho and Ákos Valentinyi (2000), "Noisy contagion without mutation", *Review of economic studies*, 67, 47-56

Lewis, David K. (1969), *Convention – a philosophical study*, Harvard University Press, Cambridge, Mass.

Morris, Stephen (2000), "Contagion", *Review of economic studies*, 67, 57-78

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

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

Young, H. Peyton (1996), "The economics of convention", *Journal of economic perspectives*, vol. 10, no. 2, 105-122

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

Viscosity and dispersion

in an evolutionary model of learning[*]

Ivar Kolstad[**]

March 2002

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

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

2

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 probabilites 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

3

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.
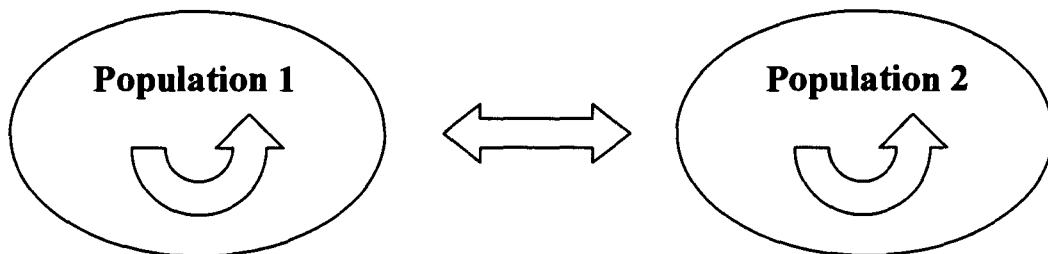
Figure 1. Structure of interaction

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 $p \to \infty$ implies complete viscosity. Fewer rounds played with your own population, $p < 1$, implies dispersive interaction, and complete dispersion as $p \to 0$. 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 β, an agent gets an opponent from his own population, whereas with probability $(1 - \beta)$ 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 ($\beta \to 0$) to completely viscous interaction ($\beta \to 1$). 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 $a > 1$

<center>Opponent</center>

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

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

<center>Opponent</center>

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

<center>7</center>

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
	B	0,0	1,1

(G1)

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
	B	0,0	b,b

(G2)

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
	B	0,0	1,b

(G3)

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 $z^t = (z_1^t, z_2^t)$ thus captures the state of the system at time t. The state space Ω is discrete and finite

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

To ease subsequent discussion, let $z^{AA} \equiv (N_1, N_2)$ represent the state in which all agents play strategy A, and let $z^{BB} \equiv (0,0)$ capture the state in which all play B. Similarly, in state $z^{AB} \equiv (N_1, 0)$ all members of population 1 play A and all members of population 2 play B. Conversely, in state $z^{BA} \equiv (0, N_2)$ 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 $\boldsymbol{z}^t$ evolves as follows. In-between periods, each agent has a probability $\delta \in \langle 0,1 \rangle$ 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 $\alpha_1 \equiv \dfrac{1}{1+a}$, where $\alpha_1 < 0.5$. An agent from population 1 therefore chooses A if

$$\frac{p \cdot z_1 + z_2}{p \cdot N_1 + N_2} > \alpha_1 \tag{2}$$

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 $\alpha_2 \equiv \dfrac{b}{1+b}$, where $0.5 < \alpha_2 < 1 - \alpha_1$. A revising agent from population 2 will thus choose strategy A if

$$\frac{p \cdot z_2 + z_1}{p \cdot N_2 + N_1} > \alpha_2 \tag{3}$$

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:

$$z_2 > \alpha_1(pN_1 + N_2) - pz_1 \tag{4}$$

$$z_2 > \frac{\alpha_2(pN_2 + N_1)}{p} - \frac{z_1}{p} \tag{5}$$

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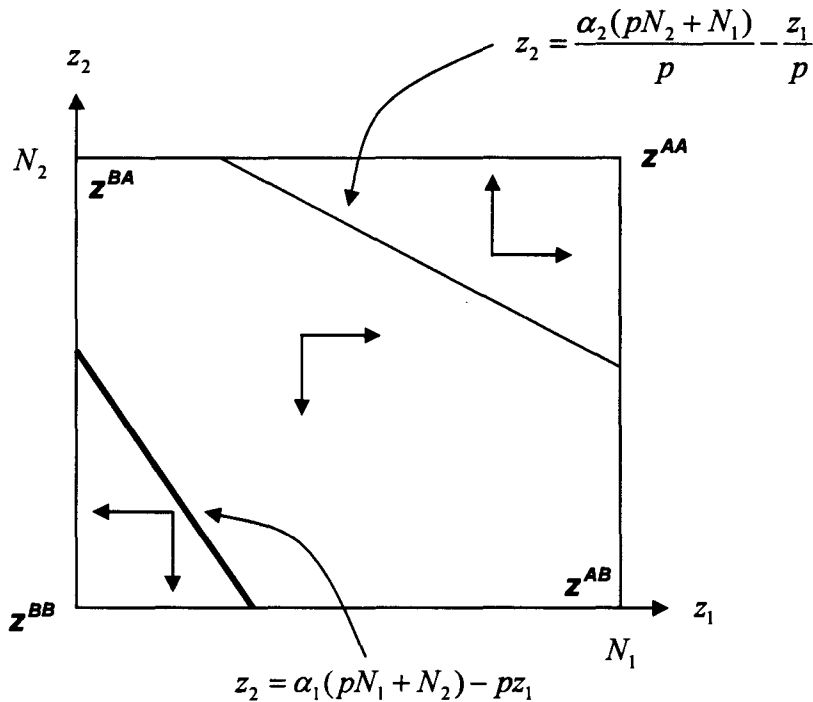

Figure 2. Illustration of stability diagram

The number of A-players in each population, z_1 and z_2, is measured along the respective axes. This implies that z^{BB} lies at the origin, and z^{AA} at the upper right corner of the rectangle formed by the axes and population sizes. Similarly, z^{AB} is at the lower right corner of the rectangle, and 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

11

(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, z^{AA}, z^{BB} and 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 $\frac{1}{p}$, respectively. And as p changes, the two lines pivot around the points $(\alpha_1 N_1, \alpha_1 N_2)$ and $(\alpha_2 N_1, \alpha_2 N_2)$, respectively. As p increases, the thicker line gets steeper, and as $p \to \infty$ 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 $p \to 0$ 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 $p \to \infty$, and vertical for $p \to 0$.

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 $\mathbf{P}(p, \varepsilon)$ be the transition matrix implied by the above learning process including mutations. In other words, element ij of $\mathbf{P}(p, \varepsilon)$ 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 $\mathbf{P}(p, \varepsilon)$ on a state space Ω, for which we use the shorthand formulation $(\Omega, \mathbf{P}(p, \varepsilon))$.

We are interested in where the process $(\Omega, \mathbf{P}(p, \varepsilon))$ goes in the very long run when noise is very small. For given p, we therefore study the probability distribution over population states in Ω, as time goes to infinity and noise ε 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 ($p \to 0$), 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 $(\Omega, \mathbf{P}(p, \varepsilon))$*, and suppose* $p \to 0$*. 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 $p \to 0$, interaction takes the form of agents from different populations playing the battle of the sexes game (G3). Proposition 1i) basically states that if populations

13

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:

$$z_2 > \alpha_1 N_2 \tag{6}$$

$$z_1 > \alpha_2 N_1 \tag{7}$$

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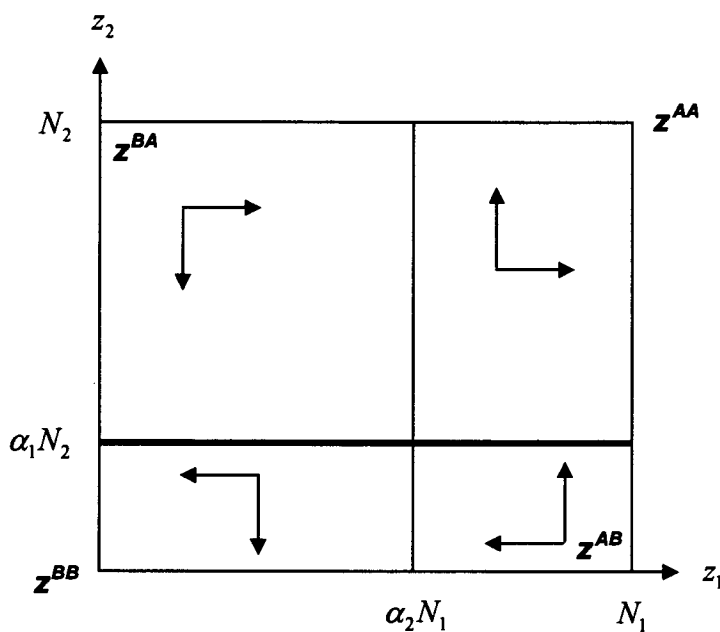
Figure 3. Stability diagram for the case of complete dispersion.

14

As the diagram reveals, no matter where we start out, we eventually reach either state z^{AA} or state 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 $p \to 0$, the unperturbed process thus has two absorbing states, z^{AA} and 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 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 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 z^{AA} to reach the thin demarcation line, compared to the number of units we would have to move downwards from z^{AA} to reach the thick demarcation line. The easiest way to leave z^{AA}, i.e. the way requiring the least mutations, is horizontally to the left, which takes $(1-\alpha_2)N_1$ mutations to strategy B in population 1. Similarly, to leave 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 z^{BB} we need a minimum of $\alpha_1 N_2$ mutations. Since we have assumed that $\alpha_1 < (1-\alpha_2)$, it is thus easier to leave z^{BB} for z^{AA} than vice versa, which means that 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 z^{BB} upwards to the thick demarcation line, exceeds the number of units from 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 z^{AA} for z^{BB} than vice versa. This holds even if population 2 becomes so much larger that the easiest way to leave z^{BB} is rightwards to the thin demarcation line, since this always implies a greater number of mutations than the transition from 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 $(\Omega, \mathbf{P}(p, \varepsilon))$, *and suppose* $p = 1$. *Then:*

For N_1 *and* N_2 *sufficiently large, there exists some* $\hat{\alpha}_2$ *such that for all* $\alpha_2 > \hat{\alpha}_2$, $\mathbf{z}^{AB}$ *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 $a > b$, $(1 - \alpha_1)$ is bounded below by α_2, which implies that the upper bound of α_1 must decrease as α_2 increases. An implication of proposition 2 is therefore that $\mathbf{z}^{AB}$ is long run stochastically stable if α_1 and α_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:

$$z_2 > \alpha_1 (N_1 + N_2) - z_1 \tag{8}$$

$$z_2 > \alpha_2 (N_2 + N_1) - z_1 \tag{9}$$

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.
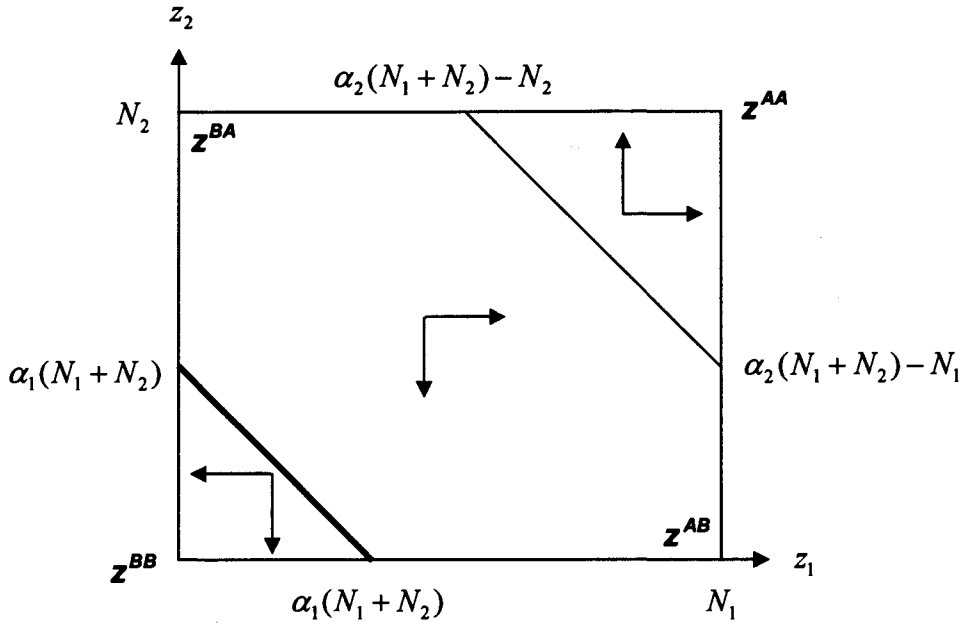
17

Figure 4. Sketch of stability diagram with fluid interaction.

In the above figure, there are three absorbing states, z^{AA}, z^{BB} and z^{AB}. Note that for $p = 1$, there need not be more than two absorbing states, z^{AA} and z^{BB}, since with unequal population sizes, $\alpha_1(N_1 + N_2)$ could exceed N_1, or $\alpha_2(N_1 + N_2) - N_1$ could be negative. However, for given population sizes, we can construct a case in which there are three absorbing states by increasing α_2 and hence decreasing the upper bound on α_1, which slides the two demarcation lines apart. At some level of $\alpha_2 < 1$, we thus get lines that cross in the manner of figure 4.

The same argument can be used to explain why z^{AB} is the long run stochastically stable state if α_2 is sufficiently large. As before, the long run stochastically stable state must be in the set of absorbing states. For α_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 z^{AB} can have this property for large α_2, we can illustrate by examining the case where $\alpha_2 \to 1$. This implies that $\alpha_1 \to 0$. The demarcation lines then slide as far as they go into opposite corners of the above figure. To leave either of the two states z^{AA} or z^{BB} for z^{AB} now requires only one mutation. However, to leave 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 z^{AB} for either of the other two absorbing states, than it is to leave either of the other two for z^{AB}. With fluid interaction, then, the state of coexistent conventions 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 $(\Omega, \mathbf{P}(p, \varepsilon))$, *and suppose* $p \rightarrow \infty$. *Then:*

If and only if N_1 *and* N_2 *are sufficiently large,* 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 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 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 z^{AB} with certainty in a finite number of periods, than to any other absorbing state. This we can demonstrate graphically. For $p \to \infty$, inequalities (4) and (5) can be rewritten as:

$$z_1 > \alpha_1 N_1 \tag{10}$$

$$z_2 > \alpha_2 N_2 \tag{11}$$

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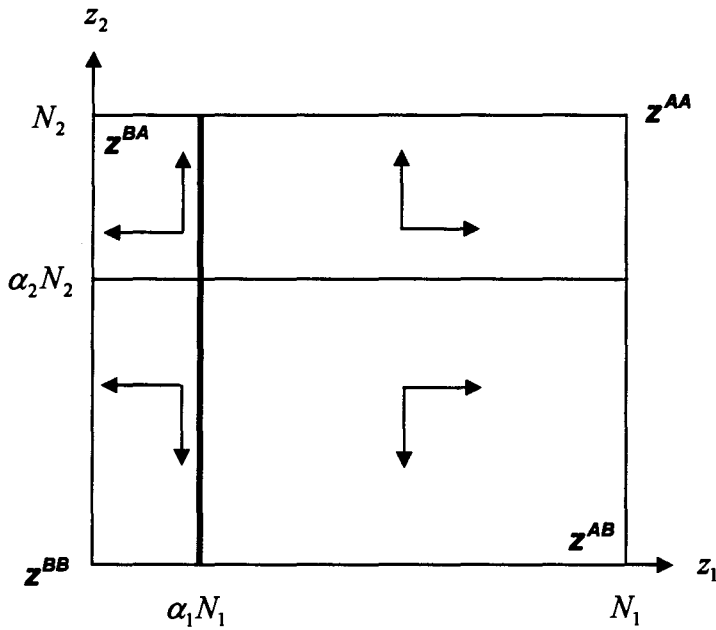
Figure 5. Stability diagram for the case of complete viscosity.

As the diagram shows, there are four absorbing states, z^{AA}, z^{BB}, z^{AB} and 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 $(1-\alpha_1)N_1 \cdot (1-\alpha_2)N_2$, $\alpha_1 N_1 \cdot \alpha_2 N_2$, $(1-\alpha_1)N_1 \cdot \alpha_2 N_2$ and $\alpha_1 N_1 \cdot (1-\alpha_2)N_2$, respectively. Since $\alpha_1 < (1-\alpha_2) < \alpha_2 < (1-\alpha_1)$, state z^{AB} thus has the larger basin of attraction of the four. In a sense, then, 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

21

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 z by the sum of the average payoffs $\pi_i(z)$ to each population $i=1,2$, weighted by the size of each population N_i.[4] If we define the relation $z' \succ^U z''$ as meaning that state z' is strictly better than state z'' from a utilitarian point of view, this relation is characterized as follows

$$z' \succ^U z'' \quad \text{iff} \quad N_1 \cdot \pi_1(z') + N_2 \cdot \pi_2(z') > N_1 \cdot \pi_1(z'') + N_2 \cdot \pi_2(z'') \quad (12)$$

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 $z' \succ^R z''$ as denoting that state z' is strictly better from a Rawlsian perspective than z''. If we simplify slightly by equating groups with populations in our model, this relation has the following property[5]

$$
\begin{aligned}
&z' \succ^R z'' \\
&\text{if } [\min\{\pi_1(z'), \pi_2(z')\} > \min\{\pi_1(z''), \pi_2(z'')\}] \\
&\text{or } [\min\{\pi_1(z'), \pi_2(z')\} = \min\{\pi_1(z''), \pi_2(z'')\} \\
&\quad \text{and } \max\{\pi_1(z'), \pi_2(z')\} > \max\{\pi_1(z''), \pi_2(z'')\}]
\end{aligned} \quad (13)
$$

---

[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, $p \to 0$, the following result holds.

PROPOSITION 4:

Consider the process of learning with noise $(\Omega, \mathbf{P}(p, \varepsilon))$, and suppose $p \to 0$. Then:

For N_1 sufficiently large, there exists some $\widetilde{N}_2 > \hat{N}_2 > N_1$ such that for $N_2 \in \langle \hat{N}_2, \widetilde{N}_2 \rangle$, z^{BB} is the unique long run stochastically stable state, while $z^{AA} \succ^U z^{BB}$ and $z^{AA} \succ^R z^{BB}$.

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, z^{BB} is better than z^{AA} if the sum of the payoffs of the two populations is larger in the former state, i.e. if

$$N_1 \cdot 1 + N_2 \cdot b > N_1 \cdot a + N_2 \cdot 1 \tag{14}$$

From the discussion and proof of proposition 1, we know that the evolutionary process selects z^{BB} if $\dfrac{N_2}{N_1} > \dfrac{(1 - \alpha_2)}{\alpha_1}$. Using the definitions of α_1 and α_2, we can rewrite this inequality as

$$N_2 \cdot b - N_1 \cdot 1 > N_1 \cdot a - N_2 \cdot 1 \tag{15}$$

In other words, the evolutionary process selects 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 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 z^{BB} if

$$\frac{N_2}{N_1} > \frac{(1-\alpha_2)}{\alpha_1} = \frac{a+1}{b+1} \tag{16}$$

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 z^{AA} and 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 z^{BB} if

$$\frac{N_2}{N_1} > \frac{a-1}{b-1} \tag{17}$$

Utilitarianism thus focuses on payoff *differences* between the states z^{AA} and 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 z^{AA} to 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 z^{BB} is more robust to mutations than z^{AA}, the utilitarian principle prefers z^{AA} over z^{BB}.

On the Rawlsian leximin principle, z^{BB} is not preferable to 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 z^{AA} and z^{BB}, which means that we must compare the payoffs of the best off population in each state, which is a in z^{AA} and b in 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 $(\Omega, P(p, \varepsilon))$, *and suppose* $p = 1$. *Then:*
For $N_1 = N_2$ *sufficiently large, there exists some* $\hat{\alpha}_2$ *such that for all* $\alpha_2 > \hat{\alpha}_2$, z^{AB}
is long run stochastically stable, while $z^{AA} \succ^U z^{AB}$ *and* $z^{AB} \succ^R z^{AA}$.

Remember from proposition 2 that the state of convention coexistence, 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 z^{AB} to be long run stochastically stable, z^{AB} is worse in terms of total payoff than z^{AA}, but better in terms of payoff to the worst off population.[7] Note that in state z^{AA}, all encounters entail coordination, which means that population 1 earns an average payoff of a, whereas population 2 earns 1. In 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 $b > 3$, z^{AB} is long run stochastically stable. Thus, from the average payoffs we see that the worse off population 2 is better off in z^{AB} than z^{AA}, when z^{AB} is long run stochastically

---

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

stable. However, in going from z^{AA} to z^{AB}, a member of population 2 improves his average payoff by $\dfrac{b-2}{2}$. 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 z^{AA} is better from a utilitarian perspective than 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 $(\Omega, \mathbf{P}(p, \varepsilon))$, *and suppose* $p \to \infty$. *Then:*
$z^{AB} \succ^U z^{AA}$, $z^{AB} \succ^U z^{BB}$ *and* $z^{AB} \succ^R z^{AA} \succ^R z^{BB}$.

In other words, 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 z^{AB}, both groups play according to their preferred equilibria, and thus get an average payoff of a and b, respectively. In state z^{AA}, only population 1 gets to play its preferred equilibrium, which makes average payoffs a and 1, respectively. In state 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 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 Ω, 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 $\varepsilon > 0$, there is a positive probability of moving from any state in Ω to any other state in Ω in a finite number of periods, which by definition means the process is irreducible. Let μ be a probability distribution over the states in Ω, and $\mathbf{P}(p, \varepsilon)$ 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:

$$\mu \mathbf{P}(p, \varepsilon) = \mu \tag{A1}$$

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 z and back again in any positive number of periods, for any state z in Ω.[8] Let $v^t(z|z^0)$ be the probability that at time t we are in state z, when at time 0 we were in state z^0. For an aperiodic and irreducible process the following result holds:

$$\lim_{t \to \infty} v^t(z|z^0) = \mu^{\varepsilon}(z) \tag{A2}$$

The probability of being in a certain state z as time goes to infinity, thus converges to the probability $\mu^{\varepsilon}(z)$ 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 $z \in \Omega$, "let N_z be the set of all integers $n \geq 1$ such that there is a positive probability of moving from z to z in exactly n periods. The process is aperiodic if for every z, the greatest common divisor of N_z is unity."

The stationary distribution μ^ε is difficult to compute, so we focus on the case where the probability of mutations is arbitrarily small, $\lim_{\varepsilon \to 0} \mu^\varepsilon$. States z which have the property that $\lim_{\varepsilon \to 0} \mu^\varepsilon(z) > 0$, we define as long run stochastically stable. A fundamental result by Young (1993) establishes that $\lim_{\varepsilon \to 0} \mu^\varepsilon$ exists, and equals a stationary distribution of the corresponding process without mutations, $\varepsilon = 0$. Note that the process without mutations is not irreducible, which means that it can have several stationary distributions. For any such stationary distribution, μ^0, the states z that have positive probabilities in this distribution, $\mu^0(z) > 0$, 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, $\lim_{\varepsilon \to 0} \mu^\varepsilon$ equals the stationary distribution μ^{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 limits 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 z' is accessible from z, if there is a positive probability of reaching z' from 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 potental, 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 Z, a radius $R(Z)$ and a coradius $CR(Z)$. In the current context of equiprobable mutations, the radius $R(Z)$ of a limit set Z is the minimum number of mutations needed to leave Z and enter a state from which another limit set is accessible. The radius thus provides a measure of how easily Z can be left for another limit set. To compute the

coradius $CR(Z)$, you take the minimum number of mutations needed to leave each of the limit sets different from Z for a state from which Z is accessible, and let the coradius equal the maximum of these. The coradius thus measures how easily Z can be reached from the other limit sets, specifically from the limit set from which Z is most difficult to reach. Ellison proves that if $R(Z) > CR(Z)$, then the states in 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 Z has the property that from each limit set different from Z there is a unique sequence of directed edges between limit sets leading to Z. As an example, imagine that there are three limit sets; A, B and C. There are thus three trees rooted at A:
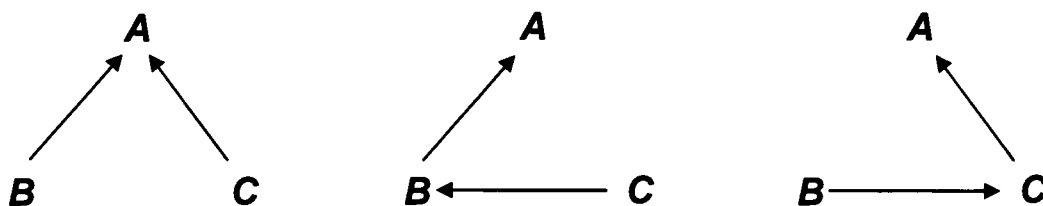
*Figure A1. Trees rooted at **A** with three limit sets **A**, **B** and **C**.*

Similarly, there are three trees rooted at B, and three trees rooted at 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 $p \to 0$, let us first compute the limit sets. For $\varepsilon = 0$, z^{AA} and z^{BB} are absorbing states of the process. In 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 z^{BB}, everyone plays B, and any revising player chooses strategy B. No other states are thus accessible from z^{AA} and 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 z^{AA}, z^{BB}, z^{AB} or z^{BA} are thus accessible from any other state. However, since in 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 z^{AB} to z^{BB} has positive probability. Similarly, in state z^{BA}, all population 1 agents might be drawn for revision, having A as their optimal choice. From z^{BA} a transition to z^{AA} thus has positive probability. In sum, this means that the absorbing states z^{AA} or z^{BB} are accessible from any other state in Ω, 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

31

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 z^{AA} to the a state from which z^{BB} is accessible, is along one of the axes. The minimum number of mutations needed to leave z^{AA} for z^{BB} is therefore $\min\{(1-\alpha_1)N_2,(1-\alpha_2)N_1\}$. In other words:

$$R(z^{AA}) = CR(z^{BB}) = \min\{(1-\alpha_1)N_2,(1-\alpha_2)N_1\} \qquad \text{(A3)}$$

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

$$R(z^{BB}) = CR(z^{AA}) = \min\{\alpha_1 N_2,\alpha_2 N_1\} \qquad \text{(A4)}$$

Since $\alpha_1 < (1-\alpha_2) < \alpha_2 < (1-\alpha_1)$, for $N_1 = N_2$ we have

$$R(z^{AA}) = (1-\alpha_2)N_1 > CR(z^{AA}) = \alpha_1 N_2 \qquad \text{(A5)}$$

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

For $N_2 > N_1$, we still have $CR(z^{BB}) = (1-\alpha_2)N_1$. However, for $\dfrac{N_2}{N_1} > \dfrac{(1-\alpha_2)}{\alpha_1}$, we get the following relation:

$$R(z^{BB}) = \min\{\alpha_1 N_2,\alpha_2 N_1\} > CR(z^{BB}) = (1-\alpha_2)N_1 \qquad \text{(A6)}$$

For sufficiently large $\dfrac{N_2}{N_1}$, z^{BB} is thus the unique long run stochastically stable state.□

PROOF OF PROPOSITION 2:

For $p = 1$, the following arguments apply. For $\varepsilon = 0$, z^{AA} and z^{BB} are still absorbing states, for the same reasons as in the proof of proposition 1. Note that in state 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 $\dfrac{N_1}{N_1 + N_2}$. If this frequency exceeds α_1, revising agents from population 1 keep playing A, i.e. if

$$\alpha_1 < \frac{N_1}{N_1 + N_2} \tag{A7}$$

Similarly, the frequency with which population 2 players meet A-players is $\dfrac{N_1}{N_1 + N_2}$. If this frequency does not exceed α_2, revising agents from population 2 keep playing B, that is if:

$$\alpha_2 > \frac{N_1}{N_1 + N_2} \tag{A8}$$

Since $\alpha_1 < (1 - \alpha_2)$, for $N_1, N_2 \in \langle 0, \infty \rangle$ there exists some $\alpha_2 < 1$ for which (A7) and (A8) hold. For sufficiently large α_2, 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 z^{AA}, z^{BB}, z^{AB} or z^{BA} from any state other than these four. In z^{BA}, an agent from either population encounters $\dfrac{N_2}{N_1 + N_2}$ A-players. A revising member of population 1 would continue playing B if

$$\frac{N_2}{N_1 + N_2} < \alpha_1 \tag{A9}$$

33

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

$$\frac{N_2}{N_1 + N_2} > \alpha_2 \qquad (A10)$$

However, since $\alpha_2 > \alpha_1$, (A9) and (A10) are incompatible. Thus, in state z^{BA}, if all agents revise, all members of either or both populations will switch strategies, in which case we transit to z^{AA}, z^{BB} or z^{AB}.

Next, we find the stochastically stable states when there are three absorbing states, z^{AA}, z^{BB} and 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 z^{AB} can be stochastically stable, so let us compute the radius and coradius of this absorbing state. The radius is:

$$R(z^{AB}) = \min\{N_1 - \alpha_1(N_1 + N_2), \alpha_2(N_1 + N_2) - N_1\} \qquad (A11)$$

And the coradius (when we make use of the fact that $\alpha_1 < (1 - \alpha_2)$):

$$CR(z^{AB}) = \max\{(1 - \alpha_2)(N_1 + N_2), \alpha_1(N_1 + N_2)\} = (1 - \alpha_2)(N_1 + N_2) \qquad (A12)$$

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

$$N_1 - \alpha_1(N_1 + N_2) > (1 - \alpha_2)(N_1 + N_2) \Leftrightarrow (1 - \alpha_2) + \alpha_1 < \frac{N_1}{N_1 + N_2} \qquad (A13)$$

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

$$\alpha_2(N_1 + N_2) - N_1 > (1 - \alpha_2)(N_1 + N_2) \Leftrightarrow 2\alpha_2 - 1 > \frac{N_1}{N_1 + N_2} \qquad (A14)$$

34

Since $\alpha_1 < (1 - \alpha_2)$, for $N_1, N_2 \in \langle 0, \infty \rangle$ there exists some $\alpha_2 < 1$ for which (A13) and (A14) hold. For sufficiently large α_2, z^{AB} is thus stochastically stable. $\square$

PROOF OF PROPOSITION 3:

For $p \to \infty$, the limit sets are as follows. As above, z^{AA}, z^{BB}, z^{AB} or z^{BA} are accessible from any other state. For the same reasons as above, z^{AA} and z^{BB} are absorbing states. With completely viscous interaction, in state 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. z^{AB} is therefore an absorbing state. In 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 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.

	z^{AA}	z^{BB}	z^{AB}	z^{BA}
z^{AA}	–	$(1-\alpha_1)N_1 + (1-\alpha_2)N_2$	$(1-\alpha_2)N_2$	$(1-\alpha_1)N_1$
z^{BB}	$\alpha_1 N_1 + \alpha_2 N_2$	–	$\alpha_1 N_1$	$\alpha_2 N_2$
z^{AB}	$\alpha_2 N_2$	$(1-\alpha_1)N_1$	–	$(1-\alpha_1)N_1 + \alpha_2 N_2$
z^{BA}	$\alpha_1 N_1$	$(1-\alpha_2)N_2$	$\alpha_1 N_1 + (1-\alpha_2)N_2$	–

$$(A15)$$

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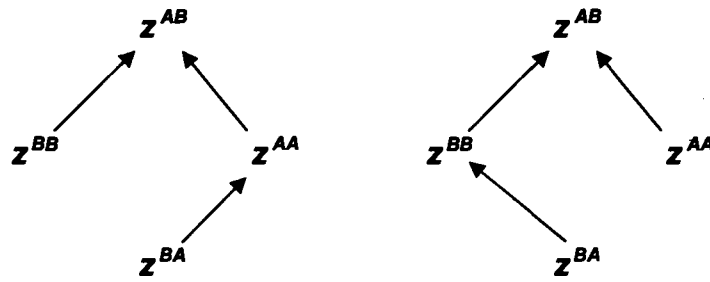
Figure A2. Two trees rooted at z^{AB}.

The sum of mutations for the transitions of each tree is $2\alpha_1 N_1 + (1 - \alpha_2)N_2$ and $\alpha_1 N_1 + 2(1 - \alpha_2)N_2$, 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 z^{AB}-tree thus has the minimum total mutations associated with it, which means that 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 $p \to 0$ and define $\hat{N}_2$ as follows

$$\frac{\hat{N}_2}{N_1} \equiv \frac{(1-\alpha_2)}{\alpha_1} = \frac{1 - \frac{b}{1+b}}{\frac{1}{1+a}} = \frac{1+a}{1+b} \qquad (B1)$$

From proposition 1, we know that if for $N_2 > \hat{N}_2$, 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 z^{AA}, and a payoff of 1 in state z^{BB}. Similarly, a player from population 2 earns a payoff of 1 per encounter in z^{AA} and b in z^{BB}. Thus, we know that $z^{AA} \succ^U z^{BB}$ if

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

However, for $a > b$

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

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

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

37

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

That $z^{AA} \succ^R z^{BB}$, 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 z^{AA} than in z^{BB}.□

PROOF OF PROPOSITION 5:

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

$$R(z^{AB}) = \min\{(1 - 2\alpha_1)N, (2\alpha_2 - 1)N\}$$ (B5)

$$CR(z^{AB}) = 2(1 - \alpha_2)N$$ (B6)

Furthermore, $\alpha_1 < (1 - \alpha_2)$ implies:

$$R(z^{AB}) = (2\alpha_2 - 1)N$$ (B7)

Consequently, we have:

$$R(z^{AB}) > CR(z^{AB}) \Leftrightarrow 2\alpha_2 - 1 > 2 - 2\alpha_2 \Leftrightarrow \alpha_2 > 0.75$$ (B8)

If $\hat{\alpha}_2 = 0.75$, then for all $\alpha_2 > \hat{\alpha}_2$, z^{AB} is long run stochastically stable. Note that $\alpha_2 > 0.75$ 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 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 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 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 $z^{AA} \succ^U z^{AB}$ when

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

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

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

PROOF OF PROPOSITION 6:

Suppose $p \to \infty$. The members of a population thus only play against each other. In state z^{AB}, both populations coordinate on their preferred equilibria, which makes payoffs a in population 1 and b in population 2. In z^{AA}, only population 1 plays according to its preferred equilibrium, which makes payoffs a and 1, respectively. In 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

$$aN_1 + bN_2 > aN_1 + N_2 \Rightarrow z^{AB} \succ^U z^{AA} \tag{B10}$$

$$aN_1 + bN_2 > N_1 + bN_2 \Rightarrow z^{AB} \succ^U z^{BB} \tag{B11}$$

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

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 β of being matched with someone from their own population, and a probability $(1 - \beta)$ 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:

$$\beta \frac{z_1}{N_1} + (1 - \beta)\frac{z_1 + z_2}{N_1 + N_2} > \alpha_1 \tag{C1}$$

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

$$\beta \frac{z_2}{N_2} + (1 - \beta)\frac{z_1 + z_2}{N_1 + N_2} > \alpha_2 \tag{C2}$$

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

$$z_2 > \frac{\alpha_1(N_1 + N_2)}{(1 - \beta)} - \frac{N_1 + \beta N_2}{(1 - \beta)N_1} z_1 \tag{C3}$$

$$z_2 > \frac{\alpha_2 N_2(N_1 + N_2)}{\beta N_1 + N_2} - \frac{(1 - \beta)N_2}{\beta N_1 + N_2} z_1 \tag{C4}$$

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

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

Fluid interaction in this model occurs when $\beta \to 0$. Inequalities (C3) and (C4) then reduce to

$$z_2 > \alpha_1 (N_1 + N_2) - z_1 \tag{C5}$$

$$z_2 > \alpha_2 (N_2 + N_1) - z_1 \tag{C6}$$

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 $\beta \to 1$. 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

Addendum to

"Viscosity and dispersion in an evolutionary model of learning"

Ivar Kolstad

December 2001

The following completes the proof of proposition 3 in "Viscosity and dispersion in an evolutionary model of learning". A figure of all z^{AA}-trees is drawn, and all trees rooted at other limit sets are derived from this figure by switching pairs of limit sets in the figure. The sum of the minimum mutations needed for the transitions of each tree, which below is termed the total resistance of each tree, is computed. I then show that the minimum resistance tree must be a tree rooted at z^{AB}.
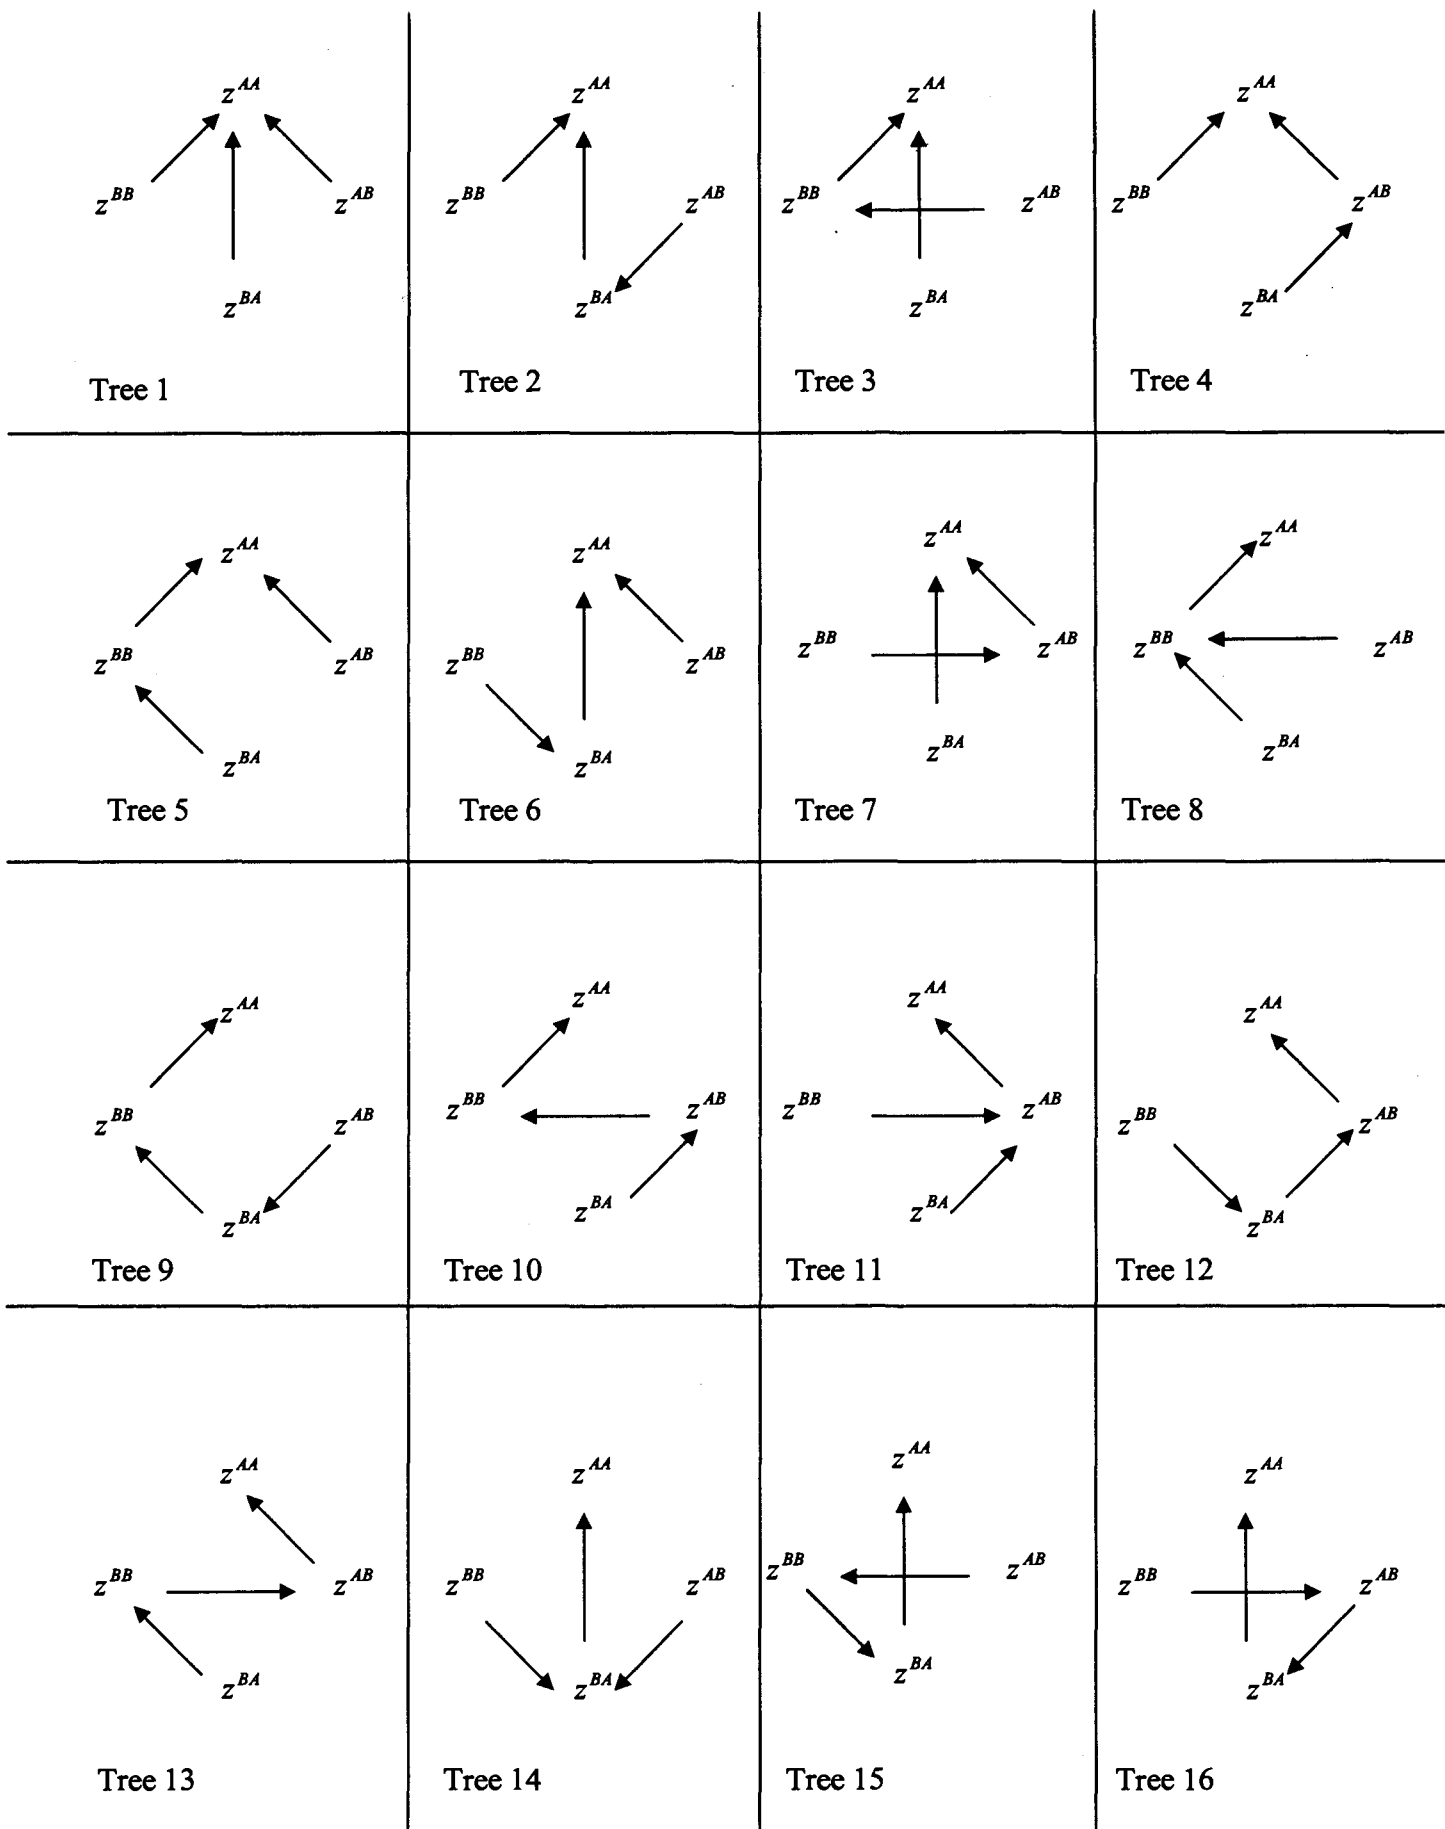
Tree 1

Tree 2

Tree 3

Tree 4

Tree 5

Tree 6

Tree 7

Tree 8

Tree 9

Tree 10

Tree 11

Tree 12

Tree 13

Tree 14

Tree 15

Tree 16

Figure of $\mathbf{z}^{AA}$-trees.

2

Total resistance of z^{AA}-trees:

Let γ_i^{AA} denote the resistance of z^{AA}-tree i in the above figure. We then have:

$$\gamma_1^{AA} = 2(\alpha_1 N_1 + \alpha_2 N_2) \tag{1}$$

$$\gamma_2^{AA} = 2(\alpha_1 N_1 + \alpha_2 N_2) + (1 - \alpha_1)N_1 \tag{2}$$

$$\gamma_3^{AA} = 2\alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_1)N_1 \tag{3}$$

$$\gamma_4^{AA} = 2(\alpha_1 N_1 + \alpha_2 N_2) + (1 - \alpha_2)N_2 \tag{4}$$

$$\gamma_5^{AA} = \alpha_1 N_1 + 2\alpha_2 N_2 + (1 - \alpha_2)N_2 \tag{5}$$

$$\gamma_6^{AA} = \alpha_1 N_1 + 2\alpha_2 N_2 \tag{6}$$

$$\gamma_7^{AA} = 2\alpha_1 N_1 + \alpha_2 N_2 \tag{7}$$

$$\gamma_8^{AA} = \alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_1)N_1 + (1 - \alpha_2)N_2 \tag{8}$$

$$\gamma_9^{AA} = \alpha_1 N_1 + 2\alpha_2 N_2 + (1 - \alpha_1)N_1 + (1 - \alpha_2)N_2 \tag{9}$$

$$\gamma_{10}^{AA} = 2\alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_1)N_1 + (1 - \alpha_2)N_2 \tag{10}$$

$$\gamma_{11}^{AA} = 2\alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_2)N_2 \tag{11}$$

$$\gamma_{12}^{AA} = \alpha_1 N_1 + 2\alpha_2 N_2 + (1 - \alpha_2)N_2 \tag{12}$$

$$\gamma_{13}^{AA} = \alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_2)N_2 \tag{13}$$

$$\gamma_{14}^{AA} = \alpha_1 N_1 + 2\alpha_2 N_2 + (1 - \alpha_1) N_1 \tag{14}$$

$$\gamma_{15}^{AA} = \alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_1) N_1 \tag{15}$$

$$\gamma_{16}^{AA} = 2\alpha_1 N_1 + \alpha_2 N_2 + (1 - \alpha_1) N_1 \tag{16}$$

Since $\alpha_1 < (1 - \alpha_1)$, we easily see that

$$\gamma_i^{AA} > \gamma_7^{AA} \text{ for all } i \in \{1,2,3,4,8,9,10,11,14,15,16\} \tag{17}$$

Similarly, since $(1 - \alpha_2) < \alpha_2$, we have

$$\gamma_i^{AA} > \gamma_{13}^{AA} \text{ for all } i \in \{5,6,12\} \tag{18}$$

The z^{AA}-tree with minimum resistance is thus either tree number 7 or number 13.

Total resistance of z^{BB}-trees:

Let z^{AA} and z^{BB} switch places in the above figure. We then get the 16 z^{BB}-trees. Denote by γ_i^{BB} the total resistance of tree i derived this way.

$$\gamma_1^{BB} = 2(1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 \tag{19}$$

$$\gamma_2^{BB} = 2(1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \tag{20}$$

$$\gamma_3^{BB} = (1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \tag{21}$$

$$\gamma_4^{BB} = 2(1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_1 N_1 \tag{22}$$

$$\gamma_5^{BB} = 2(1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_1 N_1 \tag{23}$$

$$\gamma_6^{BB} = 2(1-\alpha_1)N_1 + (1-\alpha_2)N_2 \tag{24}$$

$$\gamma_7^{BB} = (1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 \tag{25}$$

$$\gamma_8^{BB} = (1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_1 N_1 + \alpha_2 N_2 \tag{26}$$

$$\gamma_9^{BB} = 2(1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_1 N_1 + \alpha_2 N_2 \tag{27}$$

$$\gamma_{10}^{BB} = (1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_1 N_1 + \alpha_2 N_2 \tag{28}$$

$$\gamma_{11}^{BB} = (1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_1 N_1 \tag{29}$$

$$\gamma_{12}^{BB} = 2(1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_1 N_1 \tag{30}$$

$$\gamma_{13}^{BB} = (1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_1 N_1 \tag{31}$$

5

$$\gamma_{14}^{BB} = 2(1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_2 N_2 \qquad (32)$$

$$\gamma_{15}^{BB} = (1-\alpha_1)N_1 + (1-\alpha_2)N_2 + \alpha_2 N_2 \qquad (33)$$

$$\gamma_{16}^{BB} = (1-\alpha_1)N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \qquad (34)$$

Since $(1-\alpha_2) < \alpha_2$, we have

$$\gamma_i^{BB} > \gamma_7^{BB} \text{ for all } i \in \{1,2,3,4,8,9,10,11,14,15,16\} \qquad (35)$$

And due to $\alpha_1 < (1-\alpha_1)$, we have

$$\gamma_i^{BB} > \gamma_{13}^{BB} \text{ for all } i \in \{5,6,12\} \qquad (36)$$

The z^{BB}-tree with minimum resistance is thus either tree number 7 or number 13.

Total resistance of z^{AB}-trees:

As in the preceding section, if we let z^{AA} and z^{AB} switch places in the above figure, we get the 16 z^{AB}-trees. Denote by γ_i^{AB} the total resistance of tree i.

$$\gamma_1^{AB} = 2\alpha_1 N_1 + 2(1-\alpha_2)N_2 \tag{37}$$

$$\gamma_2^{AB} = 2\alpha_1 N_1 + (1-\alpha_2)N_2 + (1-\alpha_1)N_1 \tag{38}$$

$$\gamma_3^{AB} = 2\alpha_1 N_1 + 2(1-\alpha_2)N_2 + (1-\alpha_1)N_1 \tag{39}$$

$$\gamma_4^{AB} = 2\alpha_1 N_1 + (1-\alpha_2)N_2 \tag{40}$$

$$\gamma_5^{AB} = \alpha_1 N_1 + 2(1-\alpha_2)N_2 \tag{41}$$

$$\gamma_6^{AB} = \alpha_1 N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \tag{42}$$

$$\gamma_7^{AB} = 2\alpha_1 N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \tag{43}$$

$$\gamma_8^{AB} = \alpha_1 N_1 + 2(1-\alpha_2)N_2 + (1-\alpha_1)N_1 \tag{44}$$

$$\gamma_9^{AB} = \alpha_1 N_1 + (1-\alpha_2)N_2 + (1-\alpha_1)N_1 \tag{45}$$

$$\gamma_{10}^{AB} = 2\alpha_1 N_1 + (1-\alpha_2)N_2 + (1-\alpha_1)N_1 \tag{46}$$

$$\gamma_{11}^{AB} = 2\alpha_1 N_1 + (1-\alpha_2)N_2 + \alpha_2 N_2 \tag{47}$$

$$\gamma_{12}^{AB} = \alpha_1 N_1 + (1-\alpha_2)N_2 + \alpha_2 N_2 \tag{48}$$

$$\gamma_{13}^{AB} = \alpha_1 N_1 + 2(1-\alpha_2)N_2 + \alpha_2 N_2 \tag{49}$$

$$\gamma_{14}^{AB} = \alpha_1 N_1 + (1-\alpha_2)N_2 + (1-\alpha_1)N_1 + \alpha_2 N_2 \tag{50}$$

$$\gamma_{15}^{AB} = \alpha_1 N_1 + 2(1-\alpha_2)N_2 + (1-\alpha_1)N_1 + \alpha_2 N_2 \tag{51}$$

$$\gamma_{16}^{AB} = 2\alpha_1 N_1 + (1-\alpha_2)N_2 + (1-\alpha_1)N_1 + \alpha_2 N_2 \tag{52}$$

Here, since $\alpha_1 < (1-\alpha_1)$, we have the following

$$\gamma_i^{AB} > \gamma_4^{AB} \text{ for all } i \in \{1,2,3,7,8,9,10,11,14,15,16\} \tag{53}$$

Since $(1-\alpha_2) < \alpha_2$, we get

$$\gamma_i^{AB} > \gamma_5^{AB} \text{ for all } i \in \{6,12,13\} \tag{54}$$

The z^{AB}-tree with minimum resistance is therefore tree 4 or 5.

Total resistance of z^{BA}-trees:

Finally, let z^{AA} and z^{BA} switch places in the preceding figure, which transforms the trees into z^{BA}-trees. The resistance γ_i^{BA} of z^{BA}-tree i is:

$$\gamma_1^{BA} = 2(1-\alpha_1)N_1 + 2\alpha_2 N_2 \tag{55}$$

$$\gamma_2^{BA} = (1-\alpha_1)N_1 + 2\alpha_2 N_2 \tag{56}$$

$$\gamma_3^{BA} = 2(1-\alpha_1)N_1 + \alpha_2 N_2 \tag{57}$$

$$\gamma_4^{BA} = (1-\alpha_1)N_1 + 2\alpha_2 N_2 + (1-\alpha_2)N_2 \tag{58}$$

$$\gamma_5^{BA} = 2(1-\alpha_1)N_1 + 2\alpha_2 N_2 + (1-\alpha_2)N_2 \tag{59}$$

$$\gamma_6^{BA} = 2(1-\alpha_1)N_1 + 2\alpha_2 N_2 + \alpha_1 N_1 \tag{60}$$

$$\gamma_7^{BA} = 2(1-\alpha_1)N_1 + \alpha_2 N_2 + \alpha_1 N_1 \tag{61}$$

$$\gamma_8^{BA} = 2(1-\alpha_1)N_1 + \alpha_2 N_2 + (1-\alpha_2)N_2 \tag{62}$$

$$\gamma_9^{BA} = (1-\alpha_1)N_1 + 2\alpha_2 N_2 + (1-\alpha_2)N_2 \tag{63}$$

$$\gamma_{10}^{BA} = (1-\alpha_1)N_1 + \alpha_2 N_2 + (1-\alpha_2)N_2 \tag{64}$$

$$\gamma_{11}^{BA} = (1-\alpha_1)N_1 + \alpha_2 N_2 + \alpha_1 N_1 + (1-\alpha_2)N_2 \tag{65}$$

$$\gamma_{12}^{BA} = (1-\alpha_1)N_1 + 2\alpha_2 N_2 + \alpha_1 N_1 + (1-\alpha_2)N_2 \tag{66}$$

$$\gamma_{13}^{BA} = 2(1-\alpha_1)N_1 + \alpha_2 N_2 + \alpha_1 N_1 + (1-\alpha_2)N_2 \tag{67}$$

9

$$\gamma_{14}^{BA} = (1-\alpha_1)N_1 + 2\alpha_2 N_2 + \alpha_1 N_1 \qquad (68)$$

$$\gamma_{15}^{BA} = 2(1-\alpha_1)N_1 + \alpha_2 N_2 + \alpha_1 N_1 \qquad (69)$$

$$\gamma_{16}^{BA} = (1-\alpha_1)N_1 + \alpha_2 N_2 + \alpha_1 N_1 \qquad (70)$$

As $(1-\alpha_2) < \alpha_2$, we have:

$$\gamma_i^{BA} > \gamma_{10}^{BA} \text{ for all } i \in \{1,2,4,5,6,8,9,11,12,13,14\} \qquad (71)$$

And since $\alpha_1 < (1-\alpha_1)$, we get:

$$\gamma_i^{BA} > \gamma_{16}^{BA} \text{ for all } i \in \{3,7,15\} \qquad (72)$$

Thus, the z^{BA}-tree with minimum resistance is tree number 10 or 16.

Minimum resistance trees:

Due to the fact that $(1 - \alpha_2) < \alpha_2$, a comparison of (7) and (40), and (13) and (41) reveals that

$$\gamma_7^{AA} > \gamma_4^{AB} \tag{73}$$

and

$$\gamma_{13}^{AA} > \gamma_5^{AB} \tag{74}$$

Moreover, since $\alpha_1 < (1 - \alpha_1)$, when we compare (25) and (41), and (31) and (40), we get

$$\gamma_7^{BB} > \gamma_5^{AB} \tag{75}$$

and

$$\gamma_{13}^{BB} > \gamma_4^{AB} \tag{76}$$

Finally, a comparison of (64) and (41), and (70) and (40) tells us that

$$\gamma_{10}^{BA} > \gamma_5^{AB} \tag{77}$$

and

$$\gamma_{16}^{BA} > \gamma_4^{AB} \tag{78}$$

Consequently, the minimum resistance tree is a z^{AB}-tree. $\square$

11

Social origins of a work ethic:

Norms, mobility and urban unemployment[*]

Ivar Kolstad[**]

March 2002

Abstract

Neighbourhood effects and worker mobility have been proposed as explanations for the pattern of employment in cities. This paper presents a theoretical framework within which the joint impact of these two factors can be analyzed. The evolution of unemployment patterns is modelled as a stochastic process, where workers sometimes make employment decisions influenced by local norms, and sometimes decisions of where to live based on neighbourhood characteristics. A long run outcome of full employment and complete segregation is found to be robust to a wide range of process specifications. More non-segregated long run outcomes are possible if mobility decisions are based on neighbourhood employment rates than if they are based on other neighbourhood characteristics.

Keywords: Norms; Mobility; Unemployment; Urban studies

Introduction

In many cities, there is a concentration of social problems to certain inner-city areas. The disproportionate presence of poverty, crime and unemployment in central urban areas has been extensively documented by a number of empirical studies (e.g. Glaeser, Kahn and Rappaport (2000), Glaeser, Sacerdote and Scheinkman (1996), Fieldhouse (1998), Immergluck (1998), Reingold (1999) and Raphael (1998)). A number of explanantions for this spatial pattern of social problems have been suggested, many of which view unemployment as a key problem. The explanatory factors used are commonly some variation on notions of opportunity, influence and/or mobility. The traditional spatial mismatch hypothesis of Kain (1968) argues that suburban job growth has increased the distance an average inner city worker must commute to work, and thereby increased their costs of employment. In a much cited contribution, Wilson (1987) suggests that unemployment is part of a greater tangle of social problems, but deems salient the exodus of good role models from poorer neighbourhoods as an explanation of these problems. There have also been suggestions of an inflow of poor people to inner city areas, due to lower housing costs (see e.g. Glaeser, 1999) or access to public transportation for those too poor to own a car (Glaeser, Kahn and Rappaport, 2000). O'Regan and Quigley (1998) find human capital and exposure to the employed the most important factors for employment, from which we can surmise that low skill levels in inner cities lead to unemployment, which leads to more unemployment as others follow suit. Conversely, Bertrand, Luttmer and Mullainathan (1999) find that the probability of being on welfare increases with your exposure to social networks in which welfare use is more common.

The notion that there are neighbourhood or peer group effects in the spatial pattern of unemployment, has lately received much attention. Though empirical studies documenting neighbourhood effects face some methodological challenges, there nevertheless seems to be a consensus that such effects are real and important (see e.g. Glaeser, 1999). If indeed there are neigbourhood effects in employment, one implication is that the spatial distribution of employment might exhibit multiple equilibria. Or, in the words of Glaeser and Scheinkman (2000), there might exists a social multiplier, where small changes in the fundamental causes of individual

2

employment might have a large impact on the aggregate level of employment. That small changes in employment policy might have a large impact on aggregate employment, is obviously something of which policy makers should take note. However, to correctly heed these neighbourhood effects, we need some theoretical framework within which to study their implications.

In this paper, I propose a theoretical framework by means of which we can study the joint impact of neighbourhood influence and worker mobility on the level and spatial distribution of urban unemployment. The purpose of the paper is to provide a foundation for a systematic treatment of the issues involved, rather than to draw precise policy implications. The notion of influence used is consistent with that of Lindbeck, Nyberg and Weibull (1999), where workers are assumed to be influenced by a social norm against being unemployed, and where the strength of the norm depends on the number of employed workers. In contrast to Lindbeck et al, however, I assign workers locations on a social grid, and assume that each worker is influenced only by his closest neighbours. The norms are thus local, rather than global, in scope. The existence of a social grid also permits the study of worker mobility, which in the model takes the form of pairs of workers exchanging locations, as in the neighbourhood segregation model of Schelling (1971). A variety of ways in which workers might decide to move is explored, some of which are consistent with the idea of Wilson (1987) that good role models leave depressed neighbourhoods, some of which are not.

The basic approach of the paper is to model the locations and employment status of workers as a stochastic process, where workers are repeatedly drawn at random to make either decisions of whether or not to be employed, or of whether or not to move to another location. The limit sets of the process are taken to represent the patterns of employment and worker locations we can expect to see in the long run, when the process has run for a sufficiently large number of periods. The objective of the paper is to see how different assumptions about the manner in which workers make employment and mobility decisions can lead to different long-run outcomes. Though essentially a model of interdependent preferences, many of the elements of the model developed below were inspired by models in the field of evolutionary game theory, specifically those of Kandori, Mailath and Rob (1993), Ellison (1993) and Young

3

(1993, 1998). In particular, the notion of local interaction is similar to that introduced by Ellison. Since the time horizon within which it makes sense to study employment and mobility decisions is restricted, however, we focus only on long run outcomes and do not introduce error terms into decisions to select between long run outcomes, the way the aforementioned models do. The basic model also has similarities to that of Bala and Goyal (2001), but has a different object of study.

In the following section I present an initial version of the model in which workers are immobile and segregated according to their level of education, and have their employment decisions influenced by their neighbours. In section three, this simple introductory version is used to illustrate that reducing the sample of neighbours observed by a worker when making employment decisions, in effect works as a means of selection between long run outcomes. In particular, if sample sizes are below a certain level, the state of full employment is the only possible long run outcome. Section four specifies the general model in which workers make both employment decisions and decisions of where to live, and section five suggests a range of different ways in which decisions of whether or not to move can be made. In section six, I show that a state of full employment and complete segregation according to education is a long run outcome for almost all of the motives for mobility specified. Moreover, if workers move to locations that are strictly better on some characteristic, very different long run outcomes are possible, including states of full employment, states of full unemployment among those with a low level of education, states of full segregation according to education and states of full integration. In addition, more non-segregated long run outcomes are possible when mobility decisions are based on neighbourhood employment rates rather than neighbourhood composition in terms of education.

A model of neighbourhood effects in a segregated city

Consider a finite population of N workers, who inhabit equally many locations of a circular city. The workers are heterogeneous in some characteristic $e \in \{L, H\}$, which we take to be education, though it might also be productivity or some other

characteristic. Denote by N_1 the number of workers with a high level of education, $e = H$. And let $N_2 = N - N_1$ be the number of workers with a low level of education, $e = L$. Initially, we will assume that workers are completely segregated in terms of education, with the high education workers occupying positions 1 through N_1, and the low education workers inhabiting positions $N_1 + 1$ through N. In this formulation, workers thus do not have a choice of where to live.

Workers do, however, have a choice between working full time (E) and being unemployed (U). We will assume time is discrete, and in each period each worker has a probability $\delta \in \langle 0,1 \rangle$ of being called upon to revise his current employment status. When revising, a worker perceives the rewards from working as the utility $u(.)$ of consuming his net wages $w(.)$. We assume that wages are increasing in levels of education, $w(H) > w(L)$, and for a worker with education e_i we write the payoffs $\pi_i(E)$ from being employed as

$$\pi_i(E) = u[w(e_i)] \tag{1}$$

The rewards from being unemployed are the utility of consuming unemployment benefits T. There is also a social cost to being unemployed, which depends on the composition of the neighbourhood of a worker in terms of employment. We assume each worker has k neighbours to each side of him on the circle, $2k$ neighbours in all. A revising worker at location i observes a sample $s \in [1, 2k]$ of his neighbours, and assumes the proportion $\bar{q}_i$ of employed workers in this sample is representative for his neighbourhood. The social cost to being unemployed is an increasing function $v(\bar{q}_i)$ of this proportion. The payoffs $\pi_i(U)$ from choosing unemployment can then be written as[1]

$$\pi_i(U) = u(T) - v(\bar{q}_i) \tag{2}$$

---

[1] In equation (2) we assume additive separability. This means that we view the utility from benefits and the social costs as distinct elements which do not influence each other.

The social cost $v(\overline{q}_i)$ might have several interpretations. Lindbeck, Nyberg and Weibull (1999), who employ a similar payoff structure, suggest that the cost $v(\overline{q}_i)$ might reflect some social norm in favour of working, a norm whose strength depends on the number of agents adhering to it. Alternatively, if we view $v(\overline{q}_i)$ as a relative social cost, capturing the difference in socially derived payoffs when unemployed as compared to when employed, $v(\overline{q}_i)$ might represent some advantage in acting similarly to one's neighbours. Being the deviant can expose you to the resentment or distrust of others, but there are also more tangible rewards from acting in a manner similar to others. Being employed while having a network of employed neighbours might for instance provide you with more opportunities for finding a better paying job or with better ways of doing your current job. And if you are unemployed in a neighbourhood of unemployment, your chances of discovering better ways of exploiting the system of benefits might increase.

However social costs are construed, payoffs translate into actions in the following way. A revising worker at location i chooses employment if $\pi_i(E) > \pi_i(U)$, and unemployment if the opposite inequality holds. If $\pi_i(E) = \pi_i(U)$, the worker is indifferent and tosses a coin to select his employment status.

For given forms of the functions $u(.)$, $w(.)$ and $v(.)$ and a given value of the parameter T, we can derive the minimum proportion q_e^* of employed workers needed to induce a worker with education e to choose employment. In other words, there is some q_H^* such that a high education worker chooses employment if $\overline{q}_i > q_H^*$, and unemployment if $\overline{q}_i < q_H^*$. Similarly, there is some q_L^* such that a low education worker chooses employment if $\overline{q}_i > q_L^*$ and unemployment if $\overline{q}_i < q_L^*$. To add some further structure to the model, assume that $q_H^* < 0$, which means that a high education worker always chooses employment no matter how much or how little employment there is in his neighbourhood. This restriction eases analysis, by decreasing the number of states we have to consider. For low education workers, on the other hand, $q_L^* \in \langle 0,1 \rangle$, which implies that their choice of employment status does differ according

to the employment situation of their neighbourhood.[2] Utilities and social costs fitting these restrictions are illustrated by the below figure.
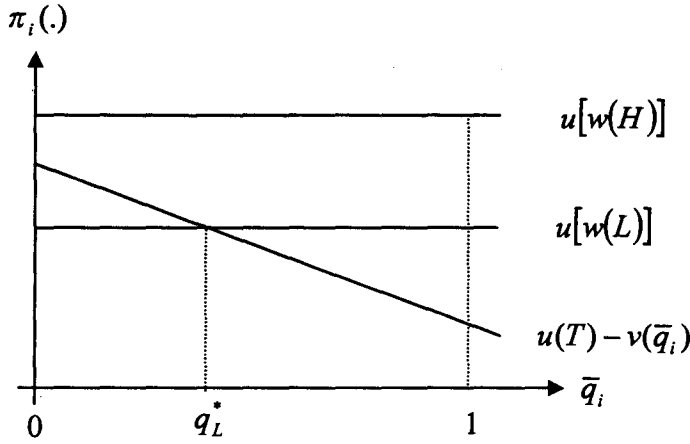
$\pi_i(.)$

Figure 1: Restrictions on payoffs

Given the assumption that agents are immobile and segregated, we can represent the state of play in period t by a vector $\boldsymbol{m}^t$, whose ith element $m_i^t \in \{E,U\}$ is the employment status of the agent at position i on the circle at time t. The state space $\overline{\Omega}$ consists of all state vectors $\boldsymbol{m}$ such that each element in $\boldsymbol{m}$ is either E or U.

$$\overline{\Omega} = \{\boldsymbol{m} : m_i \in \{E,U\}, \forall i \in [1,N]\} \tag{3}$$

For ease of exposition, let us name a few states. Denote by $\boldsymbol{m}^{EE}$ the state in which everyone is employed, i.e.

$$\boldsymbol{m}^{EE} \equiv \{\boldsymbol{m} : m_i = E, \forall i \in [1,N]\} \tag{4}$$

Similarly, let $\boldsymbol{m}^{EU}$ represent the state in which only the high-education agents are employed, while the low-education agents are unemployed

---

[2] The set of which q_L^* is an element does not contain its boundaries, which means that low education workers are not indifferent if everyone in their neighbourhood is employed or unemployed. Including the boundaries would not alter the gist of the results that follow, but would make them significantly less tidy, as the limit sets in the boundary cases could be cycles containing a large number of states.

7

$$m^{EU} \equiv \{m : m_i = E, \forall i \in [1, N_1] \,\&\, m_i = U, \forall i \in [N_1 + 1, N]\} \qquad (5)$$

Given the manner in which agents revise their employment status, the evolution of the state vector m^t constitutes a Markov chain on the state space $\overline{\Omega}$. For any given neighbourhood sample size s, let $\overline{P}(s)$ be the transition matrix implied by the process of revision, where element jk of $\overline{P}(s)$ is the probability of going from state j to state k from one period to the next. For any given s, we can then represent the process by a transition matrix $\overline{P}(s)$ on a state space $\overline{\Omega}$, which we can sum up as $(\overline{\Omega}, \overline{P}(s))$.

Long run behaviour in a segregated city

Our object of study is the evolution of play as agents repeatedly reconsider their employment status. The long run outcomes of this process, i.e. where we end up after the process of revisions has run for a large number of periods, is represented by the limit sets of the process. A limit set is a set of states which once reached, the process never leaves.[3] Even more strictly, an absorbing state is a limit set consisting of only a single state. In other words, once we have reached an absorbing state, we remain in that state in all later periods. A limit set that contains several states, is often referred to as a limit cycle.

For the above process, the following proposition captures the long run behaviour of agents:

PROPOSITION 1

Consider the process $(\overline{\Omega}, \overline{P}(s))$. For N_1 and N_2 sufficiently large:

i) m^{EE} is an absorbing state for all $q_L^* \in \langle 0,1 \rangle$ and all $s \in [1, 2k]$.

ii) m^{EU} is an absorbing state if and only if $q_L^* \in \langle 0.5,1 \rangle$ and $s > \dfrac{k}{q_L^*}$.

iii) There are no other limit sets for any $q_L^* \in \langle 0,1 \rangle$ and $s \in [1, 2k]$.

---

[3] Markov chains and limit sets are defined more rigorously in an appendix.

8

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

The first part of the proposition tells us that the state of full employment, m^{EE}, is an absorbing state for all relevant values of q_L^* and s. The state of full employment is thus robust to variations in these parameters. The intuitive reason m^{EE} is an absorbing state in all these cases, is that a revising worker in this state draws a sample of only employed workers, and thus chooses to remain employed. Once we are in the state of full employment m^{EE}, no worker ever alters his employment status, which means that we stay in m^{EE}.

In contrast, the state of full unemployment in the low education group m^{EU} is only an absorbing state for a restricted range of values of q_L^* and s. Specifically, $q_L^* \in \langle 0.5,1 \rangle$ means that for a low-education worker the required number of unemployed neighbours that would make him choose unemployment is lower than the required number of employed neighbours that would make him choose employment. Moreover, the sample of workers cannot be too small, $q_L^* \in \langle 0.5,1 \rangle$ and $s > \dfrac{k}{q_L^*}$ imply that $s > k$, so workers must sample more than half their neighbourhood for unemployment to be a stable long-term outcome.

The reason m^{EU} is not an absorbing state when low education workers are more easily persuaded to choose employment than unemployment, is as follows. Imagine that a low education worker chooses employment if exactly half or more of his neighbours are employed, $q_L^* = 0.5$. Assume that there are at least k employed high education workers. In any given period there is a chance an unemployed low education worker living next to a high education worker is called upon to revise his employment status. If he samples his entire neighbourhood, $s = 2k$, he perceives a neighbourhood employment rate of 50% and thus chooses employment. If he samples less than his entire neighbourhood, $s < 2k$, there is still a chance that half or more than half his sample are employed, upon which he chooses employment. The same

9

argument applies if in the next period the next unemployed worker on the circle revises his employment status, so there is a chance he chooses employment as well. And thus we can continue around the circle until all low education workers have chosen employment, and we have reached the state of full employment m^{EE}. The state of full employment among those with a low level of education thus unravels as the workers at the edges of the unemployed segment switch to employment.

A similar argument tells us why neighbourhood samples must be of a certain size for m^{EU} to be an absorbing state. Imagine the smallest possible sample size, $s = 1$, and consider once more the unemployed low education worker living next to a high education worker. The sample drawn by this worker might consist of a high education employed worker, which would make him choose employment for any relevant value of q_L^*. The same is true for the next unemployed worker on the circle, and so on until we reach the state of full employment m^{EE}. As the second part of proposition 1 tells us, the minimum sample size needed to prevent such an unraveling of the unemployed segment decreases as it gets harder to make low education workers choose employment.

The unemployed segment does not unravel from its edges in the above manner, if a low education worker chooses employment only if more than half his neighbours do $q_L^* > 0.5$, and if sample sizes are sufficiently large $s > \dfrac{k}{q_L^*}$. In that case, m^{EU} is an absorbing state. Note that this hinges on the size of the low education group being sufficiently large for the unemployed to sustain each other's choices. As the last part of the proposition establishes, m^{EU} and m^{EE} are in fact the only possible limit sets of the process. The reason is that if there are two low education workers living next to each other who differ in their employment status, then they have the same number of employed neighbours. If called upon to revise, at least one of them might therefore want to alter his status. Repeated revisions of this sort can bring us to m^{EE} or m^{EU}.

For certain values of the parameters of the model, we thus have two absorbing states, whereas for other values we have only one. In particular, a notable implication of proposition 1 is that by reducing the sample size of the agents in the model, we can

reduce the number of absorbing states. Reductions of sample size can thus be viewed as a means of selection between absorbing states in the present model. As small sample sizes might be taken to represent imitative behaviour of agents, and larger sample sizes more rational best reply deliberations, a reasonable interpretation of this result is that less rationality entails a more unique prediction of long run outcomes.

A model of neighbourhood effects and mobility

The above assumption of full segregation and immobile agents is rather extreme, yet serves as a useful introductory case. We now abandon this assumption, allowing any initial configuration of residences for high and low education agents, and affording agents the opportunity to switch locations. Workers thus sometimes revise their employment status, and sometimes their place of residence. The choice of employment status takes place much as in the above model, whereas for the choice of residence a range of different rules that might govern mobility are proposed.

As in the preceding model, there are N agents occupying as many locations on a circle, N_1 of whom have a high level of education (H) and the remainder a low level of education (L). The idea that workers sometimes revise their employment status and sometimes their place of residence can be modelled in a variety of ways, yet we choose the following simple variant. In each period there is a random draw, which with probability p puts us in a situation mode (S) and with probability $(1-p)$ puts us in a residence mode (R). The size of p might then reflect the frequency with which choices of employment are made relative to choices of mobility.

In a period in which we are in a situation mode, each worker has a probability $\delta \in \langle 0,1 \rangle$ of being selected to revise his employment status. The choice between employment (E) and unemployment (U) is then made the same way as in the preceding model, with one modification. Having made the above point about sample sizes, we now abandon this element and let $s = 2k$. A worker revising his employment status now observes the proportion of employment in his entire neighbourhood (i.e. across all $2k$ neighbours), and if we denote by q_i the proportion

11

employed in the neighbourhood of the worker currently occupying location i, the payoffs to this worker from unemployment become

$$\pi_i(U) = u(T) - v(q_i) \qquad (6)$$

A revising worker compares these payoffs with the payoffs from employment given by equation (1), and makes the choice which maximizes his payoffs, tossing a coin if indifferent. The restriction on payoffs imposed earlier remain in place, so a high education worker always chooses employment, whereas a low education worker is influenced by the level of employment in his neighbourhood. In the situation mode, no worker changes his place of residence.

In a period where the random draw puts us in a residence mode, two workers are drawn at random to consider switching locations with each other. The basic idea is that a move is made if both find the residence of the other more desirable than their own, or if one of the two finds the residence of the other more desirable and has the means to compensate the other for making the switch. In this respect, the model resembles the residential segregation model of Schelling (1971). In the present model, there is a variety of ways in which workers can assess the desirability of locations. In the next section, we discuss a range of these. The different ways of assessing locations are captured by rules of mobility, stating that two workers exchange locations if they and their neighbourhoods have certain characteristics. If the two do not have the required characterstics, the workers remain in their current locations. In the residence mode, no worker revises his employment status.

With mobile agents, we can represent the state of play at time t by a matrix $\mathbf{M}^t$, whose ith column $m_i^t \in \{E,U\} \times \{H,L\}$ captures the employment status and the level of education of the agent at location i on the circle at time t. The state space Ω consists of all state matrices $\mathbf{M}$ such that each column m_i of $\mathbf{M}$ has E or U in its first row, and H or L in its second row.

$$\Omega = \{\mathbf{M} : m_i \in \{E,U\} \times \{H,L\}, \forall i \in [1,N]\} \qquad (7)$$

12

For expositional convenience, we name a few sets of states. Denote by $\mathbf{M}^{EE}$ the set of states in which all workers are employed.

$$\mathbf{M}^{EE} \equiv \{\mathbf{M} : m_i \in \{E\} \times \{H, L\} \forall i \in [1, N]\} \qquad (8)$$

Moreover, let $\mathbf{M}^{EESEG}$ be the set of states in which every worker is employed, and workers are completely segregated according to their level of education. Note that $\mathbf{M}^{EESEG}$ is a subset of $\mathbf{M}^{EE}$. If A is a set of locations on the circle, and $\mathbf{A}$ the set of all such sets A that contain N_1 adjacent locations on the circle only, then we can define $\mathbf{M}^{EESEG}$ as follows.

$$\mathbf{M}^{EESEG} \equiv \{\mathbf{M} : m_i = (E, H) \forall i \in A \,\&\, M_i = (E, L) \forall i \notin A | A \in \mathbf{A}\} \qquad (9)$$

Similarly, let $\mathbf{M}^{EU}$ denote the set of states in which all workers with a high level of education are employed, and all workers with a low level of education are unemployed.

$$\mathbf{M}^{EU} \equiv \{\mathbf{M} : m_i \in \{(E, H), (U, L)\} \forall i \in [1, N]\} \qquad (10)$$

The set of states in which all high education workers are employed, all low education workers unemployed, and workers are completely segregated according to levels of education, we call $\mathbf{M}^{EUSEG}$. It follows that $\mathbf{M}^{EUSEG}$ is a subset of $\mathbf{M}^{EU}$.

$$\mathbf{M}^{EUSEG} \equiv \{\mathbf{M} : m_i = (E, H) \forall i \in A \,\&\, M_i = (U, L) \forall i \notin A | A \in \mathbf{A}\} \qquad (11)$$

In contrast, let $\mathbf{M}^{INT}$ be the set of states of perfect integration, where every other worker has a high level of education and is employed, and the locations in-between are occupied by low education workers who are unemployed.

$$\mathbf{M}^{INT} \equiv \left\{ \begin{array}{l} \mathbf{M} : (m_i = \{E, H\}, \forall i \text{ odd} \,\&\, m_i = \{U, L\}, \forall i \text{ even}) \\ \text{or} \,(m_i = \{E, H\}, \forall i \text{ even} \,\&\, m_i = \{U, L\}, \forall i \text{ odd}) | N \text{ even} \end{array} \right\} \qquad (12)$$

13

This set exists only when there are equally many agents with each type of education, which implies that the total number of agents must be even. Nevertheless, when these states do exist, they are a candidate to consider when calculating limit sets. Note that $\mathbf{M}^{INT}$ is a subset of $\mathbf{M}^{EU}$.

Rules of mobility

To know the properties of the process that governs the evolution of the state of play matrix, we must specify how decisions to move are made. To this end, we define a number of rules of mobility, each of which captures a different motive for moving. In the context of the above model, a worker revising his place of residence basically has two characteristics by which to evaluate how attractive a neighbourhood is to live in. One is the level of employment in the neighbourhood (or conversely the level of unemployment), the other is the proportion of high education workers in the neighbourhood (or conversely the proportion of low education workers). For any state $\mathbf{M}^t \in \Omega$ and any neighbourhood size k, let q_i^t denote the proportion of employed workers in the neighbourhood of the worker residing at i, and let h_i^t denote the proportion of high education workers in that neighbourhood. A revising worker can use one of these characteristics, or a combination of both, to calculate whether another location is better than his own.

Even if a worker desires to move to another location, the worker currently occupying that location might be unwilling to switch. In this case, the worker desiring to switch might compensate the other party, if he has the means to do so. Whether a switch is made thus depends on characteristics of the revising workers. One assumption is that employed workers have the means to compensate unemployed workers, and high education workers have the means to induce a switch with low education workers. The below rules of mobility capture variations of these ideas, depending on the neighbourhood characteristics by which workers evaluate the attractiveness of locations.

Let us start with the case where workers evaluate locations solely by the employment rate of their neighbourhoods. Here we make the basic assumption that workers are upwardly mobile, and thus desire to move to a location with an employment rate no lower than their current location. We also make the assumption that when one worker wants to move but not another, an employed worker can compensate an unemployed worker in order to induce a switch. For the mobility rules defined below, it would not make that much of a difference if we added the possibility that high education workers can compensate low education workers. Let us define three mobility rules based on upward mobility in terms of employment, starting with the one that requires the least in terms of an improvement in employment.

Mobility rule r^1 states that workers want to move to locations that have at least as many employed neighbours as their current location, where the employment rate of a neighbourhood is gauged by its level before a move is made. In other words, r^1 supposes a limited amount of rationality in workers, since a location that is as good as your current one before a move is made, might actually prove worse after the move is made.

DEFINITION 1:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations* $a,b \in [1,N]$ *are drawn to revise their locations.*

Then by rule r^1 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If i) $q_b^t \geq q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H,L\}$ *and iii)* $m_b^t = (U,L)$.

The definition of rule r^1 thus says that if two workers are drawn to revise their locations, they switch if one is employed and the other unemployed, and the latter is currently in a location with at least as many employed neighbours as the former. Note that this and the following definitions describe only the columns in which $\mathbf{M}^{t+1}$ differs from $\mathbf{M}^t$, i.e. the locations that are affected by workers revising their locations, for all locations i unaffected by such revisions $m_i^{t+1} = m_i^t$, as implicitly specified by the general description of the residence mode.

The next rule r^2 holds agents to be slightly more rational, comparing the rate of employment in their current location to what would be the rate of employment in a prospective location after they had moved there. As in the preceding rule, workers desire to move to locations with at least as many employed neighbours, and two workers switch locations if one is employed and desires to move and the other is unemployed.

DEFINITION 2:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations* $a, b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^2 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $q_b^{t+1} \geq q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H, L\}$ *and iii)* $m_b^t = (U, L)$.

A third rule r^3 states that workers want to move to locations where there are strictly more employed neighbours. Whether workers gauge employment by its level before or after a move is made, does not matter that much here, but we assume that they use the after-move level. As the below definition explains, an employed and an unemployed worker switch locations if the employed worker so desires.

DEFINITION 3:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations* $a, b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^3 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $q_b^{t+1} > q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H, L\}$ *and iii)* $m_b^t = (U, L)$.

Of course, workers need not be upwardly mobile. Frank (1985) suggests that it can be better to be a large fish in a small pond than vice versa. Let us include a rule reflecting this idea, where employed workers crave the status of being employed in a neighbourhood where few others are. Rule r^4 states that an employed and an unemployed worker switch locations if the employed worker gets strictly less employed neighbours this way. This rule is then in a sense the opposite of the preceding rule r^3.

DEFINITION 4:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations*

$a,b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^4 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $q_b^{t+1} < q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H, L\}$ *and iii)* $m_b^t = (U, L)$.

Having introduced a few rules based solely on neighbourhood levels of employment, let us now turn to rules that rely only on proportions of high and low education workers. In what follows, we will adopt the basic idea of Schelling (1971) that agents gravitate towards neighbourhoods that hold a greater number of agents of their own type, where the type of an agent is his level of education. High education workers thus prefer to live in neighbourhoods richer in high education workers, and low education workers prefer neighbourhoods poorer in high education workers. For the mobility rules to come, we will assume that a high education and a low education worker switch positions when the former so prefers, in order to keep definitions minimalistic. However, we could equivalently have assumed that a switch is made when both find it beneficial. Moreover, adding the possibility that low education employed workers compensate low education unemployed workers for making a switch the former finds desirable, would not significantly affect the results.

The first rule based on neighbourhood composition in terms of education, r^5, states that workers desire to move to locations where the number of neighbours currently sharing their level of education is at least as high as in their current locations. In the rationality awarded agents, this rule thus resembles r^1, since agents assess locations by their neighbourhood composition before a move is made.

DEFINITION 5:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations*

$a,b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^5 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If i) $h_b^t \geq h_a^t$, *ii)* $m_a^t \in \{E, U\} \times \{H\}$ *and iii)* $m_b^t \in \{E, U\} \times \{L\}$.

The next rule r^6 presupposes a higher degree of rationality, in stating that workers prefer to leave their current location if a prospective location holds more neighbours sharing their level of education, after the move is made.

DEFINITION 6:

Suppose that at time t we are in state $\mathbf{M^t} \in \Omega$, *and that two agents at locations* $a,b \in [1,N]$ *are drawn to revise their locations.*

Then by rule r^6 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $h_b^{t+1} \geq h_a^t$, *ii)* $m_a^t \in \{E,U\} \times \{H\}$ *and iii)* $m_b^t \in \{E,U\} \times \{L\}$.

A third rule based on neighbourhood composition in terms of education, r^7, states that workers move only to locations with strictly more of their own type. The definition assumes neighbourhood compositions are compared after a move is made, but comparisons being made before the move would not affect the results to come.

DEFINITION 7:

Suppose that at time t we are in state $\mathbf{M^t} \in \Omega$, *and that two agents at locations* $a,b \in [1,N]$ *are drawn to revise their locations.*

Then by rule r^7 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $h_b^{t+1} > h_a^t$, *ii)* $m_a^t \in \{E,U\} \times \{H\}$ *and iii)* $m_b^t \in \{E,U\} \times \{L\}$.

Once more, one might entertain the possibility that a worker would rather be a big fish in a small pond, than blend in with their neighbours. In the present context, this would mean that workers prefer locations poorer in neighbours sharing their level of education. Rule r^8 captures a variant of this idea, where locations with strictly fewer neighbours of their own type are preferred by workers.

DEFINITION 8:

Suppose that at time t we are in state $M^t \in \Omega$, *and that two agents at locations* $a,b \in [1,N]$ *are drawn to revise their locations.*

Then by rule r^8 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $h_b^{t+1} < h_a^t$, *ii)* $m_a^t \in \{E,U\} \times \{H\}$ *and iii)* $m_b^t \in \{E,U\} \times \{L\}$.

Thw two characteristics of a neighbourhood, the rate of employment and the proportion of high education workers, can also be combined in a variety of ways, to gauge how attractive locations are. Let us explore a few simple rules that combine the two. The first two of these rules are lexicographic orderings according to the two characteristics; workers prefer a location better to another according to a first characteristic, but if two locations are equally good according to the first characteristic, then workers prefer the location that is better according to the second characteristic. In this manner, rule r^9 states that an employed and an unemployed worker switch locations if the former worker gets a strictly higher number of employed neighbours that way; if he gets as many employed neighbours, a switch is made if he is a high education worker who gets strictly more high education neighbours if he moves.

DEFINITION 9:

Suppose that at time t we are in state $M^t \in \Omega$, *and that two agents at locations* $a,b \in [1,N]$ *are drawn to revise their locations.*

Then by rule r^9 $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $q_b^{t+1} > q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H,L\}$ *and iii)* $m_b^t = (U,L)$

or i) $q_b^{t+1} = q_a^t$ *and* $h_b^{t+1} > h_a^t$, *ii)* $m_a^t \in \{E,U\} \times \{H\}$ *and iii)* $m_b^t \in \{E,U\} \times \{L\}$.

Rule r^{10} is just the reverse, a high and low education worker switch locations if the former gets strictly more neighbours of his own type; if he gets at least as many, a switch is made if the low education worker is unemployed and the high education worker gets at least as many employed neighbours.

DEFINITION 10:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations* $a, b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^{10} $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $h_b^{t+1} > h_a^t$, *ii)* $m_a^t \in \{E, U\} \times \{H\}$ *and iii)* $m_b^t \in \{E, U\} \times \{L\}$

or i) $h_b^{t+1} = h_a^t$ *and* $q_b^{t+1} > q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H, L\}$ *and iii)* $m_b^t = (U, L)$.

The final rule r^{11} does not rank characteristics, but states that a move is made whenever an employed workers can get strictly more employed neighbours by switching places with an unemployed worker, and whenever a high education worker can get strictly more neighbours with a high level of education by switching places with a low education worker.

DEFINITION 11:

Suppose that at time t we are in state $\mathbf{M}^t \in \Omega$, *and that two agents at locations* $a, b \in [1, N]$ *are drawn to revise their locations.*

Then by rule r^{11} $m_a^{t+1} = m_b^t$ *and* $m_b^{t+1} = m_a^t$

If thereby i) $q_b^{t+1} > q_a^t$, *ii)* $m_a^t \in \{E\} \times \{H, L\}$ *and iii)* $m_b^t = (U, L)$

or i) $h_b^{t+1} > h_a^t$, *ii)* $m_a^t \in \{E, U\} \times \{H\}$ *and iii)* $m_b^t \in \{E, U\} \times \{L\}$.

Denote by Θ the set containing all eleven rules of mobility.

$$\Theta = \left\{ r^j : j \in \{1, \ldots, 11\} \right\} \tag{13}$$

Given the way in which agents revise their employment status and place of residence, for any of the mobility rules $r^j \in \Theta$, the evolution of the state matrix $\mathbf{M}^t$ constitutes a Markov chain on the state space Ω. If we denote by $\mathbf{P}(r^j)$ the transition matrix of the process when rule r^j is in place, we can sum up the process as $\left(\Omega, \mathbf{P}(r^j) \right)$.

Long run behaviour when agents are mobile

For each of the mobility rules defined above, the limit sets can be computed. For each rule, there can be several limit sets, depending on the value of the parameter q_L^*. Rules that are stricter in their requirements for mobility, typically have more limit sets. Limit sets for a specific few of the above rules are explored below, but let us start by examining some regularities across mobility rules.

PROPOSITION 2

Consider the process $(\Omega, \mathbf{P}(r^j))$. *For* N_1 *and* N_2 *sufficiently large:*

i) *A state* **M** *is contained in a limit set for all rules* $r^j \in \Theta - \{r^8\}$, *if and only if* **M** $\in$ **M**$^{\text{EESEG}}$.

ii) *No state* **M** $\in \Omega$ *is contained in a limit set for all rules* $r^j \in \Theta$.

The first part of the proposition captures the fact that all states in the set **M**$^{\text{EESEG}}$, i.e. states of full employment and total spatial segregation according to education, are absorbing states or contained in a limit cycle for every mobility rule defined above except r^8. States of this kind are thus remarkably robust to variations in motives of mobility, in fact more so than the states of any other set. However, no set of states is contained in a limit set for all the previously defined mobility rules, as the second part of the proposition posits. Even for states in **M**$^{\text{EESEG}}$, there are thus bounds to robustness.

The intuitive reasons why states of full employment and full segregation remain in place almost whatever motive workers have for moving, are as follows. If everyone is employed, no worker has any unemployed neighbours, and thus no worker chooses to be unemployed. The local employment norm is everywhere too strong for unemployment to be an attractive option. No unemployed workers also means that there is no available location for an employed worker to move to, so no moves are made on the basis of neighbourhood employment rates.

With full segregation, the high education workers already occupy the locations with the most high education neighbours, and they therefore cannot gain more neighbours of their own type by switching locations with low education workers. They could get as many neighbours of their own type by moving, but since larger segments of high education workers are at least as attractive as smaller ones, there is always a chance that we return to a state of full segregation. The states of full employment and full segregation thus cannot be forever abandoned if moves are made on the basis of employment, or if workers prefer to live with their own kind. If, on the other hand, workers prefer to live with the other kind, they will move away from concentrations of their own kind and not return, in which case a fully segregated state can be forever abandoned.

In more technical terms, the reason states in $\mathbf{M}^{EESEG}$ are robust to all rules of mobility but one, can be explained in the following way. First, notice that when all players are employed, employment is the optimal choice for any worker drawn to revise his employment status. If we are in a state $\mathbf{M} \in \mathbf{M}^{EESEG}$, no worker thus ever changes his employment status, which means that in all later periods, we remain within the set of states where everyone is employed $\mathbf{M}^{EE}$. Second, for any of the above mobility rules, save rule r^8, either no location switches are possible by which we go from a state in $\mathbf{M}^{EESEG}$ to a state unsegregated according to education, or if such switches are possible there exists some series of switches which brings us back to a segregated state. For the four mobility rules based solely on neighbourhood employment rates, r^1, r^2, r^3 and r^4, this is a fairly trivial matter, since according to these rules one agent must be unemployed for a location switch to occur. As there is no unemployment in a state $\mathbf{M} \in \mathbf{M}^{EESEG}$, there is thus no possibility of a switch happening, and each state in the set is thus an absorbing state.

For the first two rules based on education type, r^5 and r^6, location switches are possible in any state $\mathbf{M} \in \mathbf{M}^{EESEG}$. Consider the following figure, where there are twelve locations, and six workers of each type forming contiguous segments.
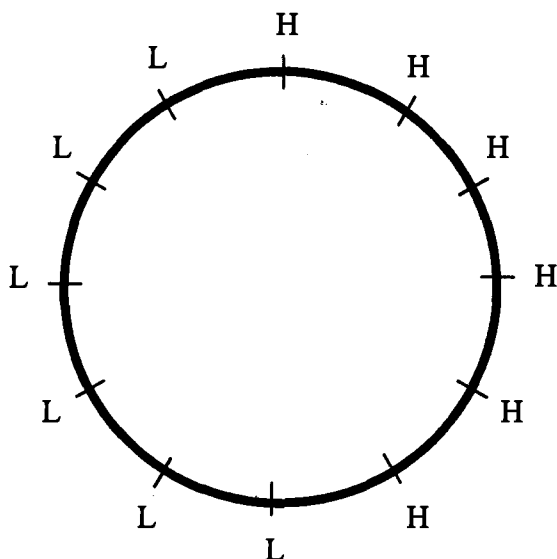
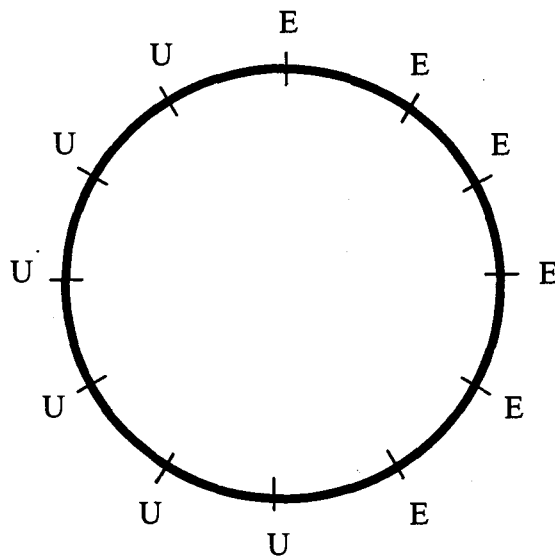*Figure 2. Illustration of a state of segregation.*

From the figure, we see that a high education agent at the edge of the high education segment has as many high education neighbours as his closest low education neighbour. Under rule r^5, this implies that two workers of this kind would exchange locations if called upon to consider this option. This particular rule thus allows us to leave a state $M \in M^{EESEG}$ for one without complete segregation. One can show, however, that from any state that is not segregated, one can reach any state $M \in M^{EESEG}$. Loosely, the reason why this happens is that any stray high education worker finds more high education neighbours in a contiguous high education segment than elsewhere. This means that the states in the set M^{EESEG} must be part of a limit set under rule r^5. Under rule r^6, adjacent high and low education workers in the above figure would not exchange locations, as they would get fewer neighbours of their own type after such an exchange. However, a high education worker at one edge of the high education segment could exchange locations with a low education worker at the other edge of that segment. Thus, location switches can rotate the high education segment around the circle, which implies that under rule r^6, the states $M \in M^{EESEG}$ form a limit cycle.

As figure 2 reveals, there are no locations to which a high education worker can move and get strictly more high education neighbours. In a state $M \in M^{EESEG}$, no moves are thus possible under rule r^7, which makes each state in the set an absorbing state.

Rules r^9, r^{10} and r^{11} just combine strict requirements of employment and high education neighbours in various ways, and thus do not allow any location switches, making any state $M \in M^{EESEG}$ an absorbing state under any of these rules. In sum, then, for rules r^1 through r^7 and r^9 through r^{11}, any state $M \in M^{EESEG}$ is contained in a limit set.

The reason why no other state is contained in a limit set for all these ten rules, can be explained in two steps. First, we can show that for rule r^1, only states of full employment $M \in M^{EE}$ are contained in limit sets. Consider the following figure in which the Ls and Hs of figure 2 have been replaced by Us and Es, respectively.

Figur 3. Illustration of a state containing employment and unemployment.

For similar reasons as in the above discussion of rule r^5, through location switches under rule r^1, we can reach a state in which there is total segregation according to employment status. Note that an employed worker at the top of the circle would now want to change locations with the unemployed worker to his left. If this happens, the next employed worker to the right on the circle would also want to switch positions with the unemployed worker. Thus we can continue until the unemployed worker has only employed neighbours, and chooses employment if called upon to revise his employment status. Every unemployed worker can be brought into the employed fold in this manner, and made to choose employment. Once everyone is employed, noone

wants to switch back to unemployment. For the rule r^1, a state is contained in a limit set only if $M \in M^{EE}$.

Second, under rule r^6, only states segregated according to education are contained in limit sets. Through a slightly more complicated argument than in the above case of rule r^5, one can show that under r^6 any state in which there is incomplete segregation according to education can be transformed into one of complete segregation through a series of location switches. Again the main reason is that high education workers prefer to move to locations where the concentration of high education neighbours is greater, which it is in contiguous segments. Once a state of segregation is reached, r^6 does not permit segregation to be abandoned. In sum, then, since under r^1 only states of employment are contained in limit sets, and under r^6 only states of segregation according to education are contained in limit sets, no state $M \notin M^{EESEG}$ can be contained in a limit set for all rules r^1 through r^7 and r^9 through r^{11}.

Finally, no state $M \in \Omega$ is contained in a limit set across all mobility rules $r^j \in \Theta$, due to the fact that no state $M \in M^{EESEG}$ is contained in a limit set under rule r^8. From figure 2, it is obvious that any high education agent would want to exchange locations with any low education agent under rule r^8, since the former agent would thereby reduce his number of high education neighbours. However, any high education agent thus separated from a high education segment would not want to rejoin that segment, since his number of high education neighbours would then rise. In a sense, high education workers want to avoid congregations of their own kind. Thus any state $M \in M^{EESEG}$ can be left for a state outside that set, but since the final switch that would lead us back to a state in that set from any other set cannot be made, the states $M \in M^{EESEG}$ cannot be part of any limit set.

As noted earlier, though states in M^{EESEG} are particularly robust to variations in mobility rules, specific mobilty rules permit a variety of limit sets, sometimes including states that are not in M^{EESEG}. The mobility rules requiring prospective neighbourhoods to be weakly better are what drives the restriction of limit sets seen in

25

proposition 2. A greater range of limit sets exists under the rules that require a prospective neighbourhood to be strictly better for a worker to want to move there. A closer examination of long run outcomes under these rules is therefore warranted. A full characterization of limit sets is difficult for the rules in question, yet the following propositions adequately capture the variety in possible long run outcomes.

PROPOSITION 3

*Consider the process* $\left(\Omega, \mathbf{P}(r^{j})\right).$

For $j = 3$, *and* N_1 *and* N_2 *sufficiently large:*

i) Any state $\mathbf{M} \in \mathbf{M}^{\mathbf{EE}}$ *is an absorbing state.*

ii) Any state $\mathbf{M} \in \mathbf{M}^{\mathbf{EUSEG}}$ *is an absorbing state if and only if* $q_L^* \in \langle 0.5, 1 \rangle.$

iii) Any state $\mathbf{M} \in \mathbf{M}^{\mathbf{INT}}$ *is an absorbing state if* $q_L^* \in \langle 0.5, 1 \rangle$ *and k is even.*

This proposition addresses rule r^3, by which employed workers move to locations that are strictly better in terms of employed neighbours. As the first part of the proposition indicates, any state of full employment, regardless of the spatial location of high and low education agents, is an absorbing state. The reasons for this are that when everyone is employed, no worker ever chooses unemployment, and since there are no unemployed workers to switch locations with, no two workers ever exchange locations.

According to the second part of the proposition, any state in which there is total segregation according to education and every low education worker is unemployed, is an absorbing state provided low education workers are more easily persuaded to choose unemployment than employment, $q_L^* > 0.5$. As figure 3 tells us, no employed worker in such a state would get more employed neighbours by switching places with an unemployed worker. And the proportion of employed neighbours for any low education worker in such a state is one half or less, which implies that if $q_L^* > 0.5$, all low education workers stay unemployed. If, on the other hand, $q_L^* \le 0.5$ the low education workers at the edges of the unemployed segment could switch to employment, and a succession of such switches would make the unemployed segment crumble from its edges.

The third part of the proposition claims that a state in which all low education workers are unemployed, and there is full integration in the sense that employed high education workers and unemployed low education workers occupy alternate locations on the circle, is an absorbing state provided $q_L^* > 0.5$ and the number of neighbours to each side k is an even number. To appreciate why this is, consider the following figure.
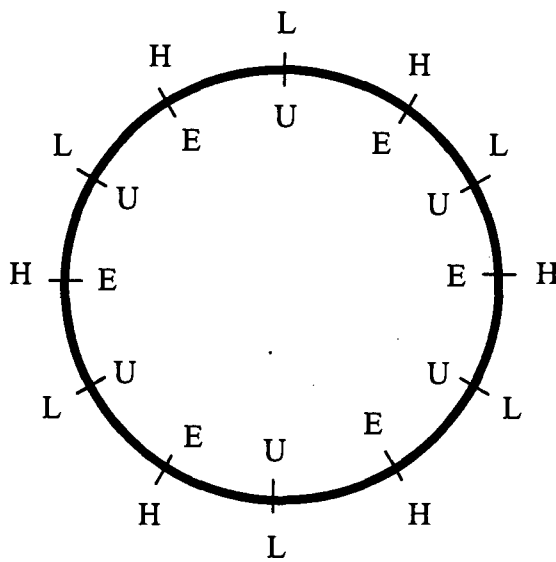
Figur 4. Illustration of integrated state

Imagine that $k = 2$, so each player has four neighbours, two to each side. Exactly half the neighbours of every worker are then employed. Thus if $q_L^* > 0.5$, low education workers remain unemployed. Every employed worker has two employed neighbours, and would get two or less by switching locations with an unemployed worker, so no location switches will occur. This line of reasoning extends to any case in which k is even. If, on the other hand, k were odd, less than half the neighbours of an employed person would be employed, whereas half or more than half his neighbours would be employed if he switched locations with an unemployed worker. For k odd, then, a state of total integration would crumble.

Very similar results to those of proposition 3 can be derived when mobility decisions are motivated by neighbourhood levels of education, or by a combination of employment rates and levels of education.

PROPOSITION 4

Consider the process $\left(\Omega, \mathbf{P}(r^j)\right)$.

For $j \in \{7,9,10,11\}$, *and* N_1 *and* N_2 *sufficiently large:*

i) Any state $\mathbf{M} \in \mathbf{M}^{EESEG}$ *is an absorbing state.*

ii) Any state $\mathbf{M} \in \mathbf{M}^{EUSEG}$ *is an absorbing state if and only if* $q_L^* \in \langle 0.5,1 \rangle$.

iii) Any state $\mathbf{M} \in \mathbf{M}^{INT}$ *is an absorbing state if* $q_L^* \in \langle 0.5,1 \rangle$ *and k is even.*

The second and third parts of proposition 4 mirror those of the preceding proposition. When we allow for the fact that here mobility is (also) based on levels of education, the reasons why segregated and integrated states of full unemployment among those with a low level of education are absorbing states, are very similar to those recounted in the above justification of proposition 3. Let us instead focus on where mobility based on education levels produces a different result from mobility based on education, as captured by the first part of the two propositions. When mobility happens due to differences in levels of education, or such differences provide an added reason to move, all states of full employment need not be contained in limit sets. Intuitively, this can be explained by imagining a state of full employment in which all high education workers but one form a contiguous segment. The high education worker isolated from his peers would then gain high education neighbours by switching locations with a low education worker at the edge of the contiguous segment of high education workers. Once such a move is made, we reach a state of full segregation, in which no high education worker can get more high education neighbours through further moves. In other words, while segregated states of full employment are absorbing states, not all non-segregated states of full employment are absorbing states.

In conjunction, propositions 3 and 4 reveal that if moves are made to locations that are strictly better on some characteristic, a wide range of long run outcomes can be observed. Both states of full employment and of full unemployment among those with

a low level of education can be absorbing states, as can both fully segregated and fully integrated states. Moving processes of this kind thus permit a wide variety of worlds to exist in the long run. However, variety is greater where decisions to move are based solely on employment levels, than where these decisions feature a comparison of neighbourhood education levels. In a sense, then, mobility based on education generates a bias towards more segregated long run outcomes.

Concluding remarks

The results obtained in this paper show that complete segregation and full employment is a long run outcome robust to variations in sample size, payoffs to workers and mobility motives. Though sample size was studied only in the initial model with fixed locations, where reduced samples were heralded as a means of selection among limit sets, a similar point could be made if employment decisions in the models including mobility were based on limited samples. However, then we would also have to tackle the question of whether only employment decisions should reflect limited samples, or if samples ought also to be assumed limited in mobility decisions. If so, more moves would be permitted under each of the above mobility rules, which on the one hand could mean that segregated states would be easier to reach, while on the other hand segregated states could also be easier to leave. Though limited samples might in this context reduce the number of limit sets, the number of states included in each limit set might rise, which makes the selection effect more dubious.

As noted initially, the purpose of the above framework is to study the joint impact of employment and mobility decisions in an urban context. A few of the assumptions on which the analysis is based are of course highly stylized, in particular the idea that workers inhabit a circular space, and that their payoffs from employment and unemployment are restricted they way they are. A more general model would let workers inhabit a more general social grid, and one way of analyzing such a model would be by means of the concept of contagion thresholds introduced by Morris (2000). Assuming greater variation in the characteristics that determine the payoffs to individuals would make the analysis richer, but also more complex as long run

outcomes would vary according to how the population is distributed across these characteristics. Finally, as matters of education or productivity are influenced by the choices, opportunities and social situation of workers, making these characteristics endogenous would also constitute an improvement to the framework proposed here.

Appendix: Proof of propositions 1 through 4

The two processes $(\overline{\Omega}, \overline{P}(s))$ and $(\Omega, P(r^j))$ are discrete time Markov processes on finite state spaces, 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. A state m' (or M') of such a process is accessible from another state m (or M), if there is a positive probability of reaching m' (or M') from m (or M) in a finite number of periods. Two states communicate if each is accessible from the other. A limit set is defined as a set of states such that all states in the set communicate, and no state outside the set is accessible from any state in 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 single state, whereas we call a limit set consisting of several states a limit cycle.

For the process $(\overline{\Omega}, \overline{P}(s))$, an absorbing state is a state in which no worker would alter his employment status, for any sample he could draw of his neighbours. For the process $(\Omega, P(r^j))$, an absorbing state is a state in which no worker would alter his employment status, and no two workers would switch locations by rule r^j. In the below proofs of the propositions, we typically establish some absorbing states (or limit sets), and then proceed to rule out further limit sets by showing that an absorbing state (or a limit set) is accessible from the remaining states.

PROOF OF PROPOSITION 1:

i) In state m^{EE}, for any revising worker, the proportion of employed neighbours observed is one, $\overline{q}_i = 1 > q_L^*$. No worker ever changes his employment status, which means that no other state is accessible from m^{EE}. The state m^{EE} is thus an absorbing state.

ii) If at time t we are in state m^{EU}, then

$$q_{N_1+1}^t \geq q_i^t \text{ for all } i \in [N_1+2, N] \qquad (A1)$$

The sequence of revisions in which agents at positions $N_1 + 1$ through N successively revise their employment status has positive probability. If with positive probability $\bar{q}_{N_1+1}^t \geq q_L^*$, then with positive probability we get $m_{N_1+1}^{t+1} = E$, which implies that with positive probability $\bar{q}_{N_1+2}^{t+1} = \bar{q}_{N_1+1}^t$ and $m_{N_1+2}^{t+2} = E$, which by repeated application implies that with positive probability $m_{N_1+i}^{t+i} = E$ for all $i \in [1, N - N_1]$. With positive probability we thus reach m^{EE} in a finite number of periods, which implies that m^{EE} is accessible from m^{EU}, and since m^{EE} is an absorbing state, m^{EU} can therefore not be contained in any limit set.

If on the other hand the probability that $\bar{q}_{N_1+1}^t \geq q_L^*$ is zero, then by virtue of (A1), $\bar{q}_i^t < q_L^*$ for any revising player at location $i \in [N_1 + 2, N]$. No sequence of revisions thus exists, for which $m_{N_1+i}^\tau = E$ for any $\tau > t$ and $i \in [1, N - N_1]$. In this case, no other state is accessible from m^{EU}, it is consequently an absorbing state.

The state m^{EU} is thus an absorbing state if and only if $\bar{q}_{N_1+1}^t < q_L^*$ for all possible samples the player at position $N_1 + 1$ could draw. For any sample size $s \in [1, 2k]$, $\bar{q}_{N_1+1}^t \geq 0.5$ with positive probability, since it is always possible that the sample the player at the edge of the employed segment draws contains all employed neighbours or only employed neighbours. For m^{EU} to be an absorbing state, we must therefore have $q_L^* > 0.5$. Moreover, since at most $\bar{q}_{N_1+1}^t = 1$ if $s < k$ and $\bar{q}_{N_1+1}^t = \frac{k}{s}$ at most if $s \geq k$, we must have $\frac{k}{s} < q_L^*$ for $\bar{q}_{N_1+1}^t \geq q_L^*$ to have zero probability. Thus, for m^{EU} to be an absorbing state, we must have $s > \frac{k}{q_L^*}$.

32

iii) Imagine that at time t we are at some state $m^* \notin \{m^{EE}, m^{EU}\}$. Starting at location $N_1 + 1$ and moving clockwise, find the first two locations for which $m'_i \neq m'_{i+1}$ where $i \in [N_1 + 1, N - 1]$. Two adjacent agents have $2k - 2$ neighbours in common, they have each other as neighbours, and their final neighbour they do not have in common. For two adjacent agents with different employment status, the employed agent then has at least as many unemployed neighbours as the unemployed agent, and the unemployed agent has at least as many employed neighbours as the employed agent. Thus, if $m'_i = E$ and $m'_{i+1} = U$, then $q^t_{i+1} \geq q^t_i$, and vice versa. For at least one of the two agents there must then exist some sample which would make him alter his strategy upon revision. If the player at location i alters his strategy, then by implication all players from $i - 1$ counter-clockwise to $N_1 + 1$ might successively alter their strategies. If the player at location $i + 1$ alters his strategy, we proceed clockwise to the next pair of adjacent agents with different employment status. By repeated applications of this procedure, we eventually end up in a state where all low-education agents have the same employment status, i.e. in m^{EE} or m^{EU}. No state $m^* \notin \{m^{EE}, m^{EU}\}$ can thus be contained in a limit set.□

PROOF OF PROPOSITION 2:

i) First we prove that any state $M \in M^{EESEG}$ is contained in a limit set under all mobility rules $r^j \in \Theta - \{r^8\}$. Note that in any state $M \in M^{EE}$, $q_i = 1 > q^*_L$ for any location $i \in [1, N]$, so no state outside M^{EE} is accessible from a state in M^{EE}. In words, no worker ever switches to unemployment since all his neighbours are employed. Since $M^{EESEG} \subset M^{EE}$, no state outside M^{EE} is thus accessible from a state in M^{EESEG}.

For the mobility rules based on employment, r^1, r^2, r^3 and r^4, location switches occur only between unemployed and employed workers, and since there are no unemployed workers in any state $M \in M^{EESEG}$, these states must be absorbing states.

For rules r^5 and r^6, location switches are possible. However, we can prove that from any state in M^{EE} that is not in M^{EESEG}, i.e. that is not segregated, we can transit to a state in M^{EESEG}. Hence, states in M^{EESEG} must be contained in some limit cycle. Consider any state $M' \in M^{EE} - M^{EESEG}$, and note that it has at least two segments of adjacent high education agents, and two segments of low education agents, otherwise it would be segregated. By implication, there are at least four pairs of high and low education workers residing at adjacent locations. By virtue of an argument similar to that used in the proof of proposition 1iii), the low education worker of such a pair must have at least as many high education neighbours as the high education worker of that pair. This due to the fact that they have $2k - 2$ neighbours in common, they have each other as neighbours, and only one neighbour that they do not share.

From the pairs of adjacent high and low education workers, take the pair with the highest number of high education neighbours (if there are several such pairs, pick any one of them). Let us say their proportion of high education neighbours is $\hat{h}$. If the number of workers with each type of education, N_1 and N_2, are large, there now exists some other pair of adjacent high and low education workers that have $\hat{h}$ or less high education workers, and that do not have the former pair in their neighbourhood. Both under rule r^5 and r^6, the high education worker of the latter pair and the low education worker of the former would switch locations.

We thus reach a new state M'', where the high education worker that just moved has $\hat{h}$ high education neighbours, and any low education worker living next to him also has $\hat{h}$ high education neighbours. Furthermore, for any worker that does not have this high education worker as a neighbour, the proportion of high education workers is equal to or less than what he had in state M'. Among these, there thus exists some pair of high and low education adjacent agents, where the high education worker would switch locations with the low education worker adjacent to the high education worker that just moved.

Thus we reach a new state M''', from which we can repeat the argument a finite number of times until a low education segment is eradicated. Then we start all over

again by finding the pair of high and low education workers with the highest number of high education neighbours, and gradually eradicate the segment associated with this low education worker as well. A finite number of repetitions of this procedure eradicates all low education segments but one, and we have reached a state in M^{EESEG}.

For rules r^7, r^9, r^{10} and r^{11}, in a state $M \in M^{EESEG}$ no location switch is ever made on the basis of employment, since under rule r^7 it is not permitted, and under rules r^9, r^{10} and r^{11} there are no unemployed workers with whom an employed worker can switch positions. Any switches would have to be made on the basis of education. However, if at time t we are in a state $M \in M^{EESEG}$, then

$$h_i^t \leq h_j^t \text{ if } m_i^t = (\cdot, L) \text{ and } m_j^t = (\cdot, H) \tag{A2}$$

By implication, no high education worker can get strictly more neighbours with a high level of education by switching locations with a low education worker, and no switches are thus ever made. Under rules r^7, r^9, r^{10} and r^{11}, any state $M \in M^{EESEG}$ is thus an absorbing state.

Next we prove that only states $M \in M^{EESEG}$ are contained in limit sets for all rules $r^j \in \Theta - \{r^8\}$. This is done in two steps, first by showing that only states in M^{EE} are contained in limit sets under rule r^1, and second by showing that of the states in M^{EE}, only those in M^{EESEG} are contained in limit sets under rule r^6.

Under mobility rule r^1, any state $M' \in \Omega$ that is not segregated according to employment, can be transformed into one that is thus segregated by a series of location switches. Start with any employed agent, and number his position on the circle 1. Then move clockwise to the first location occupied by an unemployed agent, say location a. Then proceed clockwise to the first subsequent location occupied by an employed agent, say at location b. The employed agent at location b has at least as few employed neighbours as the unemployed worker at location $b-1$, and would thus want to exchange locations with him. Having moved to location $b-1$, the employed

35

worker has at least as few employed neighbours as any unemployed worker immediately preceding him on the circle, and would want to move once more. Thus we can continue until the employed worker reaches location a. By repeating this process for each employed worker, a contiguous segment of employed workers forms from location 1 onwards, a segment that eventually holds all employed workers, which means that we are in a state of full segregation according to employment.

From a state of full segregation according to employment, we can proceed to eradicate unemployment through further location switches. Imagine that the employed segment stretches from location 1 to location c. The unemployed worker at location $c+1$ has at least as many employed neighbours as the employed worker at location c, and the two might therefore exchange locations. The unemployed worker now at location c has at least as many employed neighbours as the employed worker at location $c-1$, and the two might exchange locations. Thus we can keep moving the unemployed worker into the employed segment. If the number of high education workers is sufficiently high, the unemployed worker in question eventually has only employed workers in his neighbourhood. If selected to revise his employment status, he would then choose employment. In a similar manner we can move every single unemployed worker at locations $c+2$ through N into the employed segment one at a time, and make them choose employment, which means that we eventually reach some state of full employment $\mathbf{M} \in \mathbf{M}^{EE}$.

Under rule r^1, from any state that is not segregated according to employment we can move to one that is segregated, and from any segregated state we can move to one of full employment. A state in $\mathbf{M}^{EE}$ is thus accessible from any state outside that set. But as argued above, no state outside $\mathbf{M}^{EE}$ is accessible from a state in $\mathbf{M}^{EE}$, which implies that no state outside $\mathbf{M}^{EE}$ is contained in a limit set under rule r^1.

For rule r^6, we have already proved that from any state in $\mathbf{M}^{EE}$, we can transit to a state in $\mathbf{M}^{EESEG}$. We now add a proof of the fact that no state outside $\mathbf{M}^{EESEG}$ is accessible from any state in $\mathbf{M}^{EESEG}$ under rule r^6. Consider the high education workers in a state $\mathbf{M} \in \mathbf{M}^{EESEG}$. If not on the boundary of the high education segment, a high education worker has at least $k+1$ high education neighbours, and switching

locations with any low education worker would leave him with at most k high education neighbours. If on the boundary of the high education segment, a high education worker has exactly k high education workers, and would get as many high education neighbours if he exchanged locations with the low education worker at the opposite boundary of his segment, otherwise he would get strictly fewer. The only location switches permitted by rule r^6 in a state $\mathbf{M} \in \mathbf{M}^{EESEG}$ are thus between high and low education workers at opposite boundaries of the high education segment. This implies that from any state in $\mathbf{M}^{EESEG}$ we can move only to other states in $\mathbf{M}^{EESEG}$, which implies that the states in $\mathbf{M}^{EESEG}$ form a limit cycle. Moreover, since a state in $\mathbf{M}^{EESEG}$ is accessible from any state in $\mathbf{M}^{EE}$, no state in $\mathbf{M}^{EE}$ is contained in a limit set under rule r^6.

Summing up, only states in $\mathbf{M}^{EE}$ are contained in limit sets under rule r^1, and of the states in $\mathbf{M}^{EE}$ only those in $\mathbf{M}^{EESEG}$ are contained in limit sets under rule r^6, which implies that no state not in $\mathbf{M}^{EESEG}$ can be contained in a limit set for all rules $r^j \in \Theta - \left\{ r^8 \right\}$.

ii) Here we need only prove that no state in $\mathbf{M}^{EESEG}$ is contained in a limit set for rule r^8. Note that any high education worker in a state $\mathbf{M'} \in \mathbf{M}^{EESEG}$ has k or more high education neighbours, where the workers at the boundary of the high education segment have exactly k and those not at the boundary more than k. By switching locations with any low education worker, they would get k or less high education neighbours. Thus any high education worker not at the edge of the high education segment would switch locations with any low education worker. Moreover, a high education worker on the boundary would want to exchange locations with the low education neighbour next to him on the circle, since he would then get $k-1$ high education neighbours. (The only location switch between high and low education workers that is not permitted under rule r^8 is between a low and high education worker at opposite edges of the high education segment.) Thus from a state $\mathbf{M'} \in \mathbf{M}^{EESEG}$, we can transit to a state $\mathbf{M''} \notin \mathbf{M}^{EESEG}$.

However, any switch that caused such a transition cannot be undone, since the high education worker who moved now has at most k high education neighbours if he moved from the interior of the high education segment or at most $k-1$ high education neighbours if he moved from the boundary. To make the high education segment complete again, he would have to move to a location where he would get at least $k+1$ high education neighbours in the former case, and k high education neighbours in the latter. Under rule r^8, such a move would not be made. And since location switches happen sequentially, one at a time, such a move is needed as the final switch in a series through which an unsegregated state is supplanted by a segregated state. Under rule r^8, then, from a state in $\mathbf{M^{EESEG}}$ we can transit to a state in $\mathbf{M^{EE}} - \mathbf{M^{EESEG}}$, but no state in $\mathbf{M^{EESEG}}$ is accessible from a state in $\mathbf{M^{EE}} - \mathbf{M^{EESEG}}$, which means that no state is contained in a limit set for all rules $r^j \in \Theta$. $\square$

PROOF OF PROPOSITION 3:

i) In any state $\mathbf{M} \in \mathbf{M^{EE}}$, $q_i = 1 > q_L^*$ for a worker at any location $i \in [1, N]$. Moreover, no worker is unemployed, so no two workers ever exchange locations by rule r^3. Any state $\mathbf{M} \in \mathbf{M^{EE}}$ is therefore an absorbing state.

ii) If at time t we are in any state $\mathbf{M^t} \in \mathbf{M^{EUSEG}}$, and we let $N_1 + 1$ be the location of the unemployed worker who has an employed worker before him and an unemployed worker after him on the circle, then (A1) holds. By implication, since $q_{N_1+1}^t = 0.5$, then no worker would ever alter his employment status if $q_L^* \in \langle 0.5, 1 \rangle$. If on the other hand, $q_L^* \in \langle 0, 0.5]$, then upon revision the player at position $N_1 + 1$ could choose employment, $m_{N_1+1}^{t+1} = (E, L)$, which implies $q_{N_1+2}^{t+1} = 0.5$, which could mean $m_{N_1+2}^{t+2} = (E, L)$, and so on until $m_{N_1+i}^{t+i} = (E, L)$ for all $i \in [1, N - N_1]$. We have thus reached some state $\mathbf{M^{t+i}} \in \mathbf{M^{EESEG}} \subset \mathbf{M^{EE}}$, and no state $\mathbf{M^t} \in \mathbf{M^{EUSEG}}$ is therefore contained in a limit set.

38

In any state $\mathbf{M} \in \mathbf{M^{EUSEG}}$, any employed worker has at least k employed neighbours. By switching locations with an unemployed worker, the most employed neighbours he could get is k. No two workers would therefore exchange locations under rule r^3. In conlusion, any state $\mathbf{M} \in \mathbf{M^{EUSEG}}$ is absorbing if and only if $q_L^* \in \langle 0.5,1 \rangle$.

iii) In any state $\mathbf{M} \in \mathbf{M^{INT}}$, if k is even then $q_i = 0.5$ for all locations $i \in [1,N]$. By implication, no revising worker changes his employment status if $q_L^* \in \langle 0.5,1 \rangle$. Any employed worker in a state $\mathbf{M} \in \mathbf{M^{INT}}$ has k employed neighbours, and would get k or less by switching locations with an unemployed worker, so under rule r^3 no location switches occur in such a state. Any state $\mathbf{M} \in \mathbf{M^{INT}}$ is therefore an absorbing state for k even and $q_L^* \in \langle 0.5,1 \rangle$.□

PROOF OF PROPOSITION 4:

i) See the proof of proposition 2.

ii) From the proof of proposition 3 we know that in any state $\mathbf{M} \in \mathbf{M^{EUSEG}}$, no worker changes his employment status if $q_L^* \in \langle 0.5,1 \rangle$, whereas all unemployed workers could change sequentially to employment if $q_L^* \in \langle 0,0.5]$. Any high education worker has at least k high education neighbours in a state $\mathbf{M} \in \mathbf{M^{EUSEG}}$, and would get k or less by switching locations with a low education worker. From the proof of proposition 3ii) we know that no employed worker can get more employed neighbours by switching locations with an unemployed worker. Under rules r^7, r^9, r^{10} and r^{11}, no location switches thus occur in a state $\mathbf{M} \in \mathbf{M^{EUSEG}}$. In conclusion, any state $\mathbf{M} \in \mathbf{M^{EUSEG}}$ is absorbing if and only if $q_L^* \in \langle 0.5,1 \rangle$.

iii) The proof of proposition 3iii) tells us that in a state $\mathbf{M} \in \mathbf{M^{INT}}$, no worker changes his employment status if k is even and $q_L^* \in \langle 0.5,1 \rangle$, and no employed worker could get more employed neighbours by switching locations with an unemployed worker. If

39

k is even, in any state $\mathbf{M} \in \mathbf{M}^{\mathbf{INT}}$ a high education worker has k high education neighbours, and would get k or less if by switching with a low education worker. Under rules r^7, r^9, r^{10} and r^{11}, then, any state $\mathbf{M} \in \mathbf{M}^{\mathbf{INT}}$ is absorbing if k is even and $q_L^* \in \langle 0.5,1 \rangle$. $\square$

References

Bala, Venkatesh and Sanjeev Goyal (2001), "Conformism and diversity under social learning", *Economic theory*, 17, 101-120

Bertrand, Marianne, Erzo F.P. Luttmer, Sendhil Mullainathan (1999), "Network effects and welfare cultures", mimeo, NBER

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

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

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

Fieldhouse, Edward A. (1999), "Ethnic minority unemployment and spatial mismatch: The case of London", *Urban Studies*, vol. 36, no. 9, 1569-1596

Frank, Robert (1985), *Choosing the right pond: Human behavior and the quest for status*, Oxford University Press, New York

Glaeser, Edward L., B. Sacerdote and J. Scheinkman (1996), "Crime and social interactions", *Quarterly Journal of Economics*, 111:2, 508-548

Glaeser, Edward L. (1999), "The future of urban research: Non-market interactions", Mimeo, Harvard University/NBER

Glaeser, Edward L., Matthew E. Kahn and Jordan Rappaport (2000), "Why do the poor live in cities", Discussion Paper No 1891, Harvard Institute of Economic Research

Glaeser, Edward L., José A. Scheinkman (2000), "Non-market interactions", Working Paper 8053, NBER

Immergluck, Daniel (1998) , "Job proximity and the urban employment problem: Do suitable nearby jobs improve neighbourhood employment rates?", *Urban Studies*, vol. 35, no. 1, 7-23

Kain, J. F. (1968), "Housing segregation, negro employment, and metropolitan decentralization", *Quarterly Journal of Economics*, 82, 175-197

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

Lindbeck, Assar, Sten Nyberg and Jörgen W. Weibull (1999), "Social norms and economic incentives in the welfare state", *Quarterly Journal of Economics*, vol. CXIV, issue 1, 1-35

Morris, Stephen (2000), "Contagion", *Review of economic studies*, 67, 57-78

O'Regan, Katherine M. and John M. Quigley (1998), "Where youth live: Economic effects of urban space on employment prospects", *Urban Studies*, vol. 35, no. 7, 1187-1205

Raphael, Steven (1998), "The spatial mismatch hypethesis and black youth joblessness – Evidence from the San Francisco Bay Area", Journal of Urban Economics, vol. 43, iss. 1, 79-111

Reingold, David A. (1999), "Social networks and the employment problem of the urban poor", *Urban Studies*, vol. 36. no. 11, 1907-1932

Schelling, Thomas C. (1971), "Dynamic models of segregation", *Journal of mathematical sociology*, 1, 143-186

Wilson, William Julius (1987), *The truly disadvantaged: The inner city, the underclass and public policy*, University of Chicago Press, Chicago

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

Young, H. Peyton (1996), "The economics of convention", *Journal of economic perspectives*, vol. 10, no. 2, 105-122

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

Evolution with endogeneous mutations[*]

Ivar Kolstad[**]

May 2002

Abstract

Bergin and Lipman (1996) prove that equilibrium selection in the evolutionary dynamics of Kandori, Mailath and Rob (1993) and Young (1993), is not robust to variations in mutation rates across states. Specifically, a risk dominant equilibrium can be selected against if the mutation rates are higher in its basin of attraction than elsewhere. Van Damme and Weibull (1998) model mutations as a compromise between payoff losses and control costs, finding that mutation rates are lower in the risk dominant equilibrium. This paper argues that this result in not driven by control costs, but by the assumption that players focus on payoff losses when choosing mutation rates.

Keywords: Evolution, mutation, equilibrium selection

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

1

Introduction

One of the more promising approaches to selection among multiple equilibria in games, has been the one proposed by Kandori, Mailath and Rob (1993) and Young (1993). Both these contributions study a population of agents who are matched repeatedly to play a game, acting adaptively by choosing a best reply to play in previous periods. In the model of Kandori et al, agents adapt to play in the preceding period, whereas in the model of Young, agents adapt to an observation of a limited sample of play in the most recent periods. In a simple 2x2 coordination game, these adaptive dynamics lead us to one of the two pure Nash equilibria. What makes selection among the equilibria possible is the introduction of a mutation probability on top of the adaptive dynamics. Every time an agent decides on a strategy, there is a small probability that he implements a strategy different from his preferred one. As this mutation probability becomes arbitrarily small, the relative frequency with which we observe one of the equilibria in the very long run tends to one. Specifically, Kandori et al and Young show that the equilibrium thus selected is what Harsanyi and Selten (1988) term the risk dominant equilibrium.

Bergin and Lipman (1996), however, challenge the notion that small mutation rates select a unique equilibrium. Their basic point is that Kandori et al and Young base their results on the restrictive idea that mutation rates are the same in all states. If we allow any variation of mutation rates across states, Bergin and Lipman show that any strict Nash equilibrium can be selected when mutation rates are taken to zero. Specifically, the risk dominant equilibrium need not be selected if mutation rates are higher in its basin of attraction than in that of another equilibrium.

To say anything precise on equilibrium selection, then, we need to specify the element of mutations more precisely. What is it we mean to capture by mutations, and what implications does this have for variations in mutation rates across states? Van Damme and Weibull (1998) suggest one way in which to understand mutations; as mistakes that agents can control at a cost. By assuming that agents want to avoid mistakes more that are more costly in terms of payoff losses, van Damme and Weibull reach the conclusion that agents have lower mistake probabilities in the risk dominant equilibrium. Selection of the risk dominant equilibrium is thus upheld under this interpretation of mutations.

This paper takes a closer look at mutations in evolutionary models, and in particular the mechanism behind the latter result. In the next section, we recount the rudiments of the evolutionary models of Kandori et al and Young, explaining the intuition behind equilibrium selection when mutation rates are fixed and when they vary. In section three, we consider different ideas of what mutations can be taken to represent, and explore specifically the implications of viewing mutations as experiments or mistakes. In section four, we review the results of van Damme and Weibull, focusing on why mutation rates in their model are lower in the risk dominant equilibrium. It is suggested that what drives their result is not the cost associated with lowering mutation probabilities, but their assumption that an agent wants to reduce probabilities more when these are more costly to that agent. If agents care about the possibility of reaching a better state through mutations, the opposite result can be proven, where the risk dominant equilibrium has a higher mutation rate. A final section argues that if the equilibrium selection results are to become more precise and applicable to a human decision making context, the basic evolutionary framework must reflect empirical regularities in the way humans make decisions.

The dichotomy of adaptive dynamics and mutations

The evolutionary models of Kandori, Mailath and Rob (1993) and of Young (1993) basically view actions as the outcome of two distinct processes, an adaptive dynamic and a random process of mutation. The two models allow some variation in the way the adaptive process is modelled, but the following framework captures their general gist. From a finite population, agents are matched repeatedly to play a game. In each period, agents update their strategies, assuming that the frequency with which strategies were played in the preceding period reflect the probability with which they will be played in the current period, and choosing a best reply to this probability distribution. If we let the number of players using each strategy represent the state of play in any period, the process of adaptation is then a Markov process on a finite state space.

Let Ω be the state space, and $\mathbf{P}$ the transition matrix defined by the adaptive process, where element ij of $\mathbf{P}$ is the probability of going from state i to j from one period to the next. Let μ be an invariant probability distribution over Ω, which has the property $\mu = \mu\mathbf{P}$. Since the

process **P** on Ω can have several such invariant distributions, the long run outcome of the process need not be unique. However, if we add mutations, in the sense that in each period each agent has a small probability ε of choosing a strategy at random, we get a new process **P**$^\varepsilon$ on Ω which is irreducible since every state has a positive probability of being reached from every other state. And an irreducible Markov process has a unique invariant distribution μ^ε. Young (1993) shows that if we let the probability of mutation ε get arbitrarily small, then $\lim_{\varepsilon \to 0} \mu^\varepsilon$ exists and equals an invariant distribution of the process without mutations, **P** on Ω. Arbitrarily small mutation rates thus permit selection among the states that can be observed in the long run without mutations. The states that have a positive probability in $\lim_{\varepsilon \to 0} \mu^\varepsilon$, Young calls stochastically stable. They can be identified as the states with minimum stochastic potential, which basically means that they are more costly to leave in terms of mutations than they are to reach from other states.

Let us briefly review why the above process of adaptation and mutations selects the risk dominant equilibrium in a 2x2 coordination game. For simplicitly, assume the game is symmetric.

Player 2

		A	B
		A	B
Player 1	A	a,a	c,d
	B	d,c	b,b

(G1)

We assume $a > d$ and $b > c$, which makes (G1) a coordination game with two Nash equilibria in pure strategies, (A,A) and (B,B). We also assume that $a - d > b - c$, which implies that (A,A) is the risk dominant equilibrium as defined by Harsanyi and Selten (1988).[1] Finally, assume that $b > a$, which makes (B,B) the Pareto dominant equilibrium.

---

[1] 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.

A player is indifferent between A and B if the probability with which his opponent chooses A is α, where

$$\alpha \equiv \frac{b-c}{(a-d)+(b-c)}$$ (1)

If the probability is higher than α, a player prefers strategy A, if it is lower, he prefers B. The assumption that (A,A) is risk dominant implies that $\alpha < 0.5$.

According to the adaptive dynamics, the probability that your opponent chooses A equals the frequency with which A was played in the population in the previous period. If the proportion playing A is above α, players thus switch to A, increasing the proportion playing A. If the proportion playing A is below α, it decreases over time. If the population constists of N agents, and $z \in [0, N]$ denotes the number playing strategy A, we can depict the adaptive dynamics as follows.
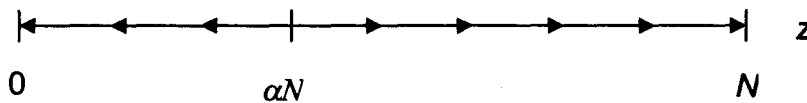

Figure 1. Adaptive dynamics

The adaptive dynamics thus leads us to a state in which everyone plays A, i.e. where equilibrium (A,A) is observed, or to a state where everyone plays B, and (B,B) is observed. As the figure shows, there are more states leading to (A,A) than to (B,B), the former state thus has the larger basin of attraction.

By adding mutations, we can select between the two equilibria. The greater size of the basin of attraction of (A,A) implies that it takes more mutations to upset this equilibrium than equilibrium (B,B). More than half the population must mutate for us to leave a state in which everyone plays A for one where everyone plays B, whereas less than half the population must mutate for the opposite transition to occur. If mutations are equally probable in all states, it is

5

thus less likely that equilibrium (A,A) is upset by mutations than (B,B), and infinitely less likely as the probability of mutation ε tends to zero. In having the larger basin of attraction, the risk dominant equilibrium (A,A) is thus the more costly to leave in terms of mutations, and it is therefore the one selected by the process of adaptation and mutation.

Bergin and Lipman (1996) point out, however, that the above result is based on the restrictive assumption that mutation rates are the same in all states. They show that if we allow any variation in mutation rates across states, we can make the invariant distribution of the process with mutations equal any invariant distribution of the mutationless process, as $\varepsilon \to 0$. In other words, any strict Nash equilibrium can be selected by the process of adaptation and mutations, if mutation rates are not the same in all states.

By implication, we can select the Pareto dominant equilibrium (B,B) of (G1) if mutation rates are higher in the basin of attraction of (A,A) than in the basin of attraction of (B,B). To be more specific, the mutation rates must converge to zero at different rates, if ε is the rate of mutation in the former basin of attraction and ε^x in the latter, then we must have $x > 1$ for (B,B) to be selected. By reference to figure 1, we can explain the selection of (B,B) as follows. Though the risk dominant equilibrium (A,A) has the larger basin of attraction, in the sense that it takes more mutations to leave, the basin of the Pareto dominant equilibrium (B,B) is deeper or steeper in that each mutation is less probable. If the basin of (B,B) is sufficiently deep, it is harder to upset (B,B) than (A,A) through mutations, which means that the process of adaptation and mutations selects the Pareto dominant equilibrium (B,B).

The conclusion we can draw from Bergin and Lipman's result is that if we know nothing of how mutation rates vary, we cannot determine which equilibrium will be selected. Conversely, to argue for the selection of one equilibrium over another, we need a closer specification of mutation rates. The contents of such a model of mutations depends on what we take mutations to represent.

Mutations as experiments or mistakes

The idea of a dynamic process consisting of adaptation and mutation is taken from biology, where it has a firm basis in Darwin's theory of evolution and in genetics. The adaptive part here reflects the idea that genes that are more successful in reproducing themselves, will over time grow more prominent in the total gene pool. Mutations capture the idea that genes change spontaneously on rare occasions. Whether it is reasonable to base a theory of decisions on a similar dichotomy, depends on what the theory is meant to represent. Is it a theory of how *rational* beings make decisions, or is it a theory of how *human* beings make decisions?

If the evolutionary models of Kandori, Mailath and Rob (1993) and of Young (1993) are a continuation of the endeavours of game theory in studying the behaviour of rational beings, then equilibrium selection through mutations seems to come at the expense of internal consistency. By allowing deviations from a clearly defined best reply relation, mutations seem to introduce an element of bounded rationality into a theory of rationality. Whether rationality is indeed sacrificed through the introduction of mutations, of course depends on what definition of rationality you use. Mutations and even the idea of gradual adaptation seem hard to reconcile with a theory of ideally rational players, who have perfect foresight and are infallible. Sugden (2001) thus suggests that the evolutionary approach marks a sharp departure from the fundamentals upon which game theory is based. Yet on a weaker conception of rationality, it is possible to imagine players who are rational in pursuing their ends, yet fallible in sometimes not being able to carry out their preferred strategies. Superimposing adaptation and mutations on a theory of decisions thus seems to come at the expense of the degree of rationality rather than at the expense of internal consistency.

If, one the other hand, the evolutionary models are to be models of human decision making, then the idea that decisions can be divided into adaptation and mutations must be empirically justified. Thus far, the empirical basis for the analogy of human decisions to biological evolution is anecdotal at best. To explore the empirical justification of using adaptation and mutations to model decisions, the field of cognitive psychology and in particular social psychology should be consulted. Sugden (2001) suggests that behavioural economics sets a useful standard for how a theoretical approach to human decisions ought to evolve.

The distinction between evolutionary models as models of rational decision making versus human decision making, has some implications for our understanding of what mutations represent. The adaptive component is thought to represent the way agents adjust to or learn from the actions of others, and typically takes the form of strategy changes that leave agents better off. Mutations, on the other hand, are idiosyncrasies in behaviour, random actions taken though they may leave an agent worse off. These random idiosyncrasies are typically seen as *experiments* on the part of agents, or as *mistakes* in implementing strategies. Whether we view mutations as experiments or as mistakes, should somehow be reflected in the way in which we model mutations, in particular in the way we let mutation rates vary across states.

Van Damme and Weibull (1998, p. 1) suggest that there are natural or reasonable implications of viewing mutations as experiments or as mistakes. If mutations represent experiments, "individuals may be expected to experiment less in states with higher payoffs". In other words, the rate of experimentation can be expected to be higher in the risk-dominant equilibrium (A,A) of (G1) than in the Pareto-dominant equilibrium (B,B). On the other hand, if mutations are mistakes, "one might argue that mistakes associated with larger payoff losses are less likely". Which means that the rate of mistakes is lower in the risk dominant equilibrium (A,A) than in the Pareto-dominant equilibrium (B,B), since (loosely) the loss from being the sole agent playing B in a population of A-players is larger than that of being the only A-player in a B-playing population.

These implications are not as self-evident as they may seem. They are made on the basis of appeals to intuition, yet it does not seem less intuitively appealing to make the reverse claim: Individuals make fewer mistakes in states with higher payoffs, and experiments associated with larger payoff losses are less likely. What is missing is a definition of experimentation and of mistakes, which clearly delineates what each is and how it is different from the other. From such a distinction one could then attempt to draw some implications about variations in mutation rates.

In trying to make a distinction between the two notions of mutation, the basis on which such a distinction is to be made depends on what the evolutionary model is trying to capture. If the evolutionary model is a model of how rational beings make decisions, the two terms need only be defined in a way consistent with rationality in the sense assumed by the model. Since a world of rational beings is an abstract construct, there is no natural way in which to

distinguish experimentation from mistakes. Thus one might choose to define experimentation as the kind of random idiosyncrasy which has a higher probability when payoffs are lower, and mistakes as the kind more likely when payoff losses are lower. But any insights about equilibrium selection in each case are then a matter of tautology.

On the other hand, if the evolutionary model is one of how human beings make decisions, experimentation and mistakes can perhaps be viewed as empirically distinct and distinguishable aspects of human conduct. There might thus be a natural way of separating the two, and hence each might give rise to a different probability structure across states. However, a forceful empirical argument for the idea that experimentation and mistake probabilities differ in the particular way suggested by van Damme and Weibull, has yet to be made.

Evolution with mutation control costs

Based on their idea of what mistakes are, van Damme and Weibull (1998) formulate a model of how mistake probabilities are determined. Using an idea from van Damme (1987), they assume players can control the probability with which they make mistakes, but at a cost. These control costs are assumed to grow infinitely large as mistake probabilities are eliminated completely, which implies that players will not invest in completely avoiding mistakes. The gains from reducing mistake probabilities, van Damme and Weibull assume to be the reduced expected loss to an agent from deviating from his best reply. In other words, agents want to avoid mistakes more that are more costly to make.

If we limit ourselves to game (G1), the basic structure of this model of endogenous mistake probabilities is as follows. Let $z \in [0, N]$ denote the number of A-players in a population of N agents. The expected payoffs to an agent from choosing strategy $s \in \{A, B\}$ is denoted by $\pi(s, z)$. Define $b(z)$ as the payoffs to a player from choosing his best reply to the strategy profile z of the population

$$b(z) \equiv \max\{\pi(A, z), \pi(B, z)\} \qquad (2)$$

Similarly, let $w(z)$ be the payoffs from choosing his worst reply to z

$$w(z) \equiv \min\{\pi(A,z), \pi(B,z)\} \tag{3}$$

For mistake probability $\varepsilon \in (0,1)$, let the function $v(\varepsilon)$ denote the control costs of keeping mistakes at probability ε. The control cost function $v(\varepsilon)$ is assumed to be positive, strictly convex, symmetric around 0.5, twice differentiable, and has $\lim_{\varepsilon \to 0} v(\varepsilon) = \lim_{\varepsilon \to 1} v(\varepsilon) = +\infty$. In other words, reducing ε increases costs, and more from a lower level of probabilities ε, until costs increase infinitely as ε approaches zero. The below figure illustrates a control cost function which meets these requirements.
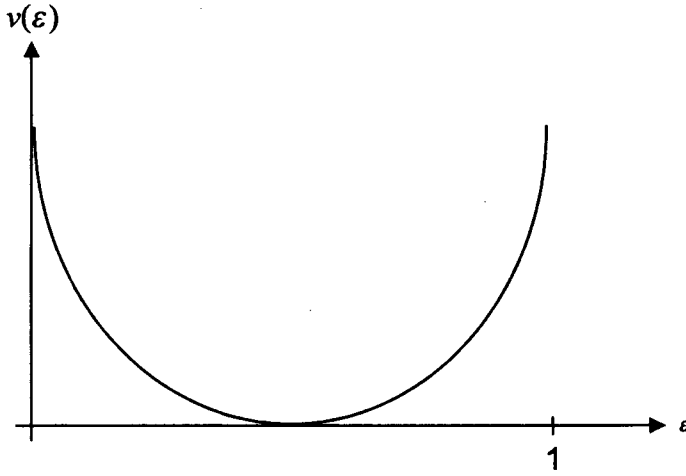

Figure 2. Illustration of control cost function $v(\varepsilon)$

Each agent is assumed to choose the mistake probability ε which maximizes his expected payoff, wherein the control costs $v(\varepsilon)$ are scaled by a parameter δ which measures the relative size of control costs to payoffs

$$\max_{\varepsilon} \quad (1-\varepsilon)b(z) + \varepsilon w(z) - \delta v(\varepsilon) \tag{4}$$

The first-order condition is

$$-v'(\varepsilon) = \frac{b(z) - w(z)}{\delta} \tag{5}$$

10

The left-hand side of (5) is the marginal cost of reducing the mistake probability ε. Given the assumptions about $v(\varepsilon)$, $-v'(\varepsilon)$ decreases from infinity to zero on the interval $\varepsilon \in (0, \frac{1}{2})$. The right-hand side of (5) is the payoff loss from mistakes in state z, divided by the scaling parameter δ. Since by assumption $a - d > b - c$, the payoff loss is greater when everyone plays according to the risk dominant equilibrium (A,A), i.e. $z = N$, than in the Pareto dominant equilibrium (B,B), $z = 0$. As shown in the below figure, the first order condition thus tells us that the risk dominant equilibrium (A,A) is associated with lower mistake probabilities $\varepsilon(N)$, than that of the Pareto dominant equilibrium, $\varepsilon(0)$.
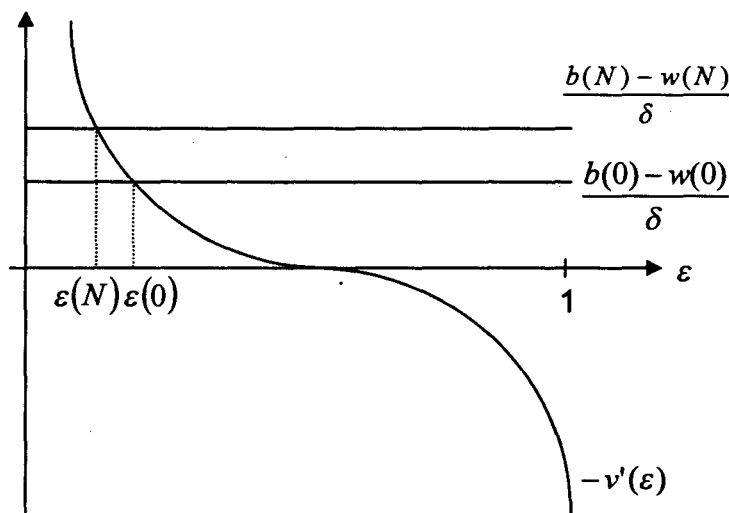

Figure 3. Mutation rates in the van Damme and Weibull model

With the endogenously determined mutation rates of the above model, the equilibrium selection result of Kandori, Mailath and Rob (1993) and Young (1993) thus holds; the evolutionary process selects the risk dominant equilibrium. Van Damme and Weibull suggest that what drives the result, are the control costs. Though control costs are important to the result, it is more reasonable to claim that what drives the result is the assumption that agents want to avoid mistakes that lead to greater payoff losses. To make this point more fully, let us consider another motive agents might have, one that can lead us to the opposite result even in the presence of control costs.

Besides payoff losses from mistakes, agents might consider the possibility that mistakes bring them to a better or worse state than the current one. If enough agents make a mistake in either of the two equilibria, the adaptive dynamics lead us to the other equilibria. An agent might

thus perceive choosing a higher mistake probability in one state as increasing the chances that the other equilibrium is reached. He might therefore want to have a high probability of mistakes in the risk dominant equilibrium, to increase the probability of jumping to the Pareto dominant equilibrium where payoffs are higher. Similarly, he might want to have a low probability of mistakes in the Pareto dominant equilibrium, since a jump to the other equilibrium would entail lower payoffs.

Let us consider one specific version of this idea. Assume that in choosing mistake probabilities, agents want the highest possible expected payoffs given a notion that all other agents choose the same mistake probability. Agents are in a sense Kantian, in seeing where mistakes would lead them if all agents shared their choice of mistake probabilities. The idea that the adaptive process is governed by selfish motives, and mutations by Kantian motives might seem contradictory. However, since the adaptive process and mutations are meant to capture distinct and different phenomena, there is nothing inherently inconsistent in assigning different governing motives to the two.

Let $p(z)$ be the payoffs from the equilibrium in whose basin of attraction is the present state z

$$p(z) \equiv \begin{cases} \pi(A,N) & \text{if } z > \alpha N \\ \pi(B,0) & \text{if } z < \alpha N \end{cases} \tag{6}$$

Similarly, let $r(z)$ be the payoffs from the equilibrium in whose basin of attraction the present state z is not

$$r(z) \equiv \begin{cases} \pi(A,N) & \text{if } z < \alpha N \\ \pi(B,0) & \text{if } z > \alpha N \end{cases} \tag{7}$$

From the equilibrium in whose basin of attraction is the current state z, let $f_z(\varepsilon)$ be the probability that enough mistakes occur to reach the other basin of attraction. In other words

$$f_z(\varepsilon) \equiv \begin{cases} \sum_{i \in ((1-\alpha)N, N]} const_i \varepsilon^i (1-\varepsilon)^{(N-i)} & \text{if } z > \alpha N \\ \sum_{i \in (\alpha N, N]} const_i \varepsilon^i (1-\varepsilon)^{(N-i)} & \text{if } z < \alpha N \end{cases} \tag{8}$$

12

This implies that

$$
f_z'(\varepsilon) = \begin{cases} \sum\limits_{i \in ((1-\alpha)N, N]} const_i \left[\varepsilon^{(i-1)} (1-\varepsilon)^{(N-i-1)} \right] (i - \varepsilon N) & \text{if } z > \alpha N \\ \sum\limits_{i \in (\alpha N, N]} const_i \left[\varepsilon^{(i-1)} (1-\varepsilon)^{(N-i-1)} \right] (i - \varepsilon N) & \text{if } z < \alpha N \end{cases}
\tag{9}
$$

Note that $f_z'(\varepsilon) > 0$ if $\alpha > \varepsilon$, which holds if mistake probabilities are sufficiently low.

In choosing mistake probabilities, we assume that agents maximize a weighted average of payoff losses as in (4), and expected equilibrium payoffs if everyone chooses the same mistake probability, from which we subtract control costs $v(\varepsilon)$ defined as above.

$$
\max_\varepsilon \ \beta \left[(1-\varepsilon)b(z) + \varepsilon w(z) \right] + (1-\beta) \left[((1 - f_z(\varepsilon))p(z) + f_z(\varepsilon)r(z) \right] - \delta v(\varepsilon)
\tag{10}
$$

The first-order condition thus becomes

$$
-v'(\varepsilon) = \frac{\beta}{\delta} \left[b(z) - w(z) \right] + \frac{(1-\beta)}{\delta} \left[p(z) - r(z) \right] f_z'(\varepsilon)
\tag{11}
$$

We can now prove the following relation between mistake probabilities when everyone plays A, $\varepsilon(N)$, and mistake probabilities when everyone plays B, $\varepsilon(0)$.

PROPOSITION 1

For agents maximizing (10), there exists some $\beta \in (0,1)$ and some control cost function $v(\varepsilon)$ such that $\varepsilon(N) > \varepsilon(0)$.

PROOF OF PROPOSITION 1:

For $z = N$, $p(z) < r(z)$ and for $z = 0$, $p(z) > r(z)$. Given α, from (9) we know there is some region $\varepsilon \in (0, \hat{\varepsilon})$ for which $f_z'(\varepsilon) > 0$ for any z.

Let $\beta \to 0$. The right hand side of the first-order condition (11) starts in the origo and in the interval $\varepsilon \in (0, \hat{\varepsilon})$ slopes upward for $z = 0$ and downward for $z = N$.

The left hand side of (11) slopes downward from plus infinity to zero in the interval $\varepsilon \in (0, \frac{1}{2})$, and from zero to minus infinity in the interval $\varepsilon \in (\frac{1}{2}, 1)$. We let $\varepsilon(0)$ denote the intersection of $-v'(\varepsilon)$ and $\frac{\beta}{\delta}[b(0) - w(0)] + \frac{(1-\beta)}{\delta}[p(0) - r(0)]f_0'(\varepsilon)$. If $\hat{\varepsilon} \geq \frac{1}{2}$, or if the slope of $-v'(\varepsilon)$ is sufficiently flat in the interval $\varepsilon \in \langle \hat{\varepsilon}, \frac{1}{2} \rangle$, then $\frac{\beta}{\delta}[b(N) - w(N)] + \frac{(1-\beta)}{\delta}[p(N) - r(N)]f_N'(\varepsilon)$ is downward sloping in the interval $\varepsilon \in [0, \varepsilon(0))$, and consequently intersects $-v'(\varepsilon)$ at a point $\varepsilon(N) > \varepsilon(0)$. □

The intuitive argument for proposition 1 is perhaps best understood by means of a figure of the first order condition (11). The right hand side of (11) is now not a constant, since ε features in the term $f_z'(\varepsilon)$. Note that in the risk dominant equilibrium, $z = N$, the payoffs in the present state is lower than that of the other equilibrium, $p(N) - r(N) < 0$. In the Pareto dominant equilibrium, $z = 0$, on the other hand, $p(0) - r(0) > 0$. For small ε, equation (9) tells us that $f_z'(\varepsilon) > 0$. The right hand side of (11) is thus upward sloping for the Pareto dominant equilibrium and downward sloping for the risk dominant one.
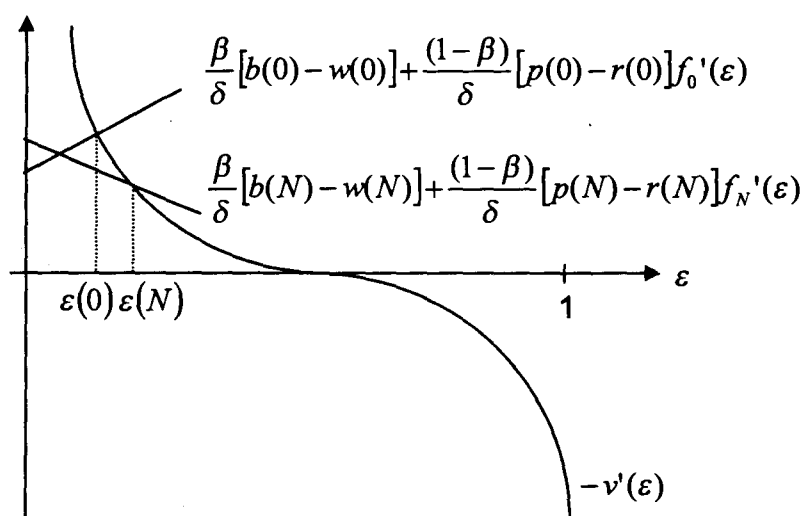

$$\frac{\beta}{\delta}[b(0) - w(0)] + \frac{(1-\beta)}{\delta}[p(0) - r(0)]f_0'(\varepsilon)$$

$$\frac{\beta}{\delta}[b(N) - w(N)] + \frac{(1-\beta)}{\delta}[p(N) - r(N)]f_N'(\varepsilon)$$

$\varepsilon(0)\,\varepsilon(N)$

1

ε

$-v'(\varepsilon)$

Figure 4. Mutation rates in the augmented model

14

As the figure illustrates, we can thus have a situation where the the right and left hand side of (11) intersect at a lower mistake probability ε for the Pareto dominant equilibrium than for the risk dominant one. In other words, we can get a higher rate of mutations in the risk dominant equilibrium than in the Pareto dominant equilibrium. Augmenting the objective function of the agents we thus get the opposite result of van Damme and Weibull, even if the control costs have the same structure as in their model. Control costs alone are therefore not sufficient to uphold the equilibrium selection results of Kandori, Mailath and Rob (1993) and Young (1993).

Concluding remarks

Bergin and Lipman (1996) question the robustness of equilibrium selection in the evolutionary models of Kandori, Mailath and Rob (1993) and Young (1993), by pointing out that if mutation rates are allowed to vary in any way across states, any strict Nash equilibrium can be selected by the evolutionary process. Equilibrium selection thus becomes a matter of determining the structure of mutations in specific cases, which comes down to a question of what we want mutations to represent. Van Damme and Weibull (1998) endogenize mutation rates by assuming that mutations represent mistakes that agents can control at a cost, and suggest that this corroborates the equilibrium selection results of Kandori et al and Young. In this paper, I show that the equilibrium selection result obtained by van Damme and Weibull follows from the assumption that agents seek to avoid mistakes that are more costly, more than from the introduction of control costs. Specifically, if agents see mistakes as a way to get to a better (or worse) equilibrium, we can get the opposite result even in the presence of control costs.

The discussion of the technicalities of equilibrium selection obscures what is perhaps a more fundamental challenge to the evolutionary perspective: Does a biological model of adaptation and mutation provide a good or even useful representation of decision making. As stated earlier, this depends on what such a model of decision making is supposed to capture. If it is a model of how rational beings make decisions, spontaneous idiosynchrasies in the form of mutations must somehow be reconciled with the concept of rationality which underpins the approach. If it is a model of how human beings make decisions, the specifics of the model

15

must be reconciled with empirical regularities in human decision making. If the dichotomy of adaptation and mutation is at all useful in the latter sense, a closer examination of the forces behind human experimentation and error is needed to increase our understanding of how humans come to focus on one out of several possible equilibria.

References

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

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

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

Sugden, Robert (2001), "The evolutionary turn in game theory", *Journal of economic methodology*, 8:1, 113-130

van Damme, Eric (1987), *Stability and perfection of Nash equilibria*, Springer Verlag, Berlin

van Damme, Eric and Jörgen W. Weibull (1998), "Evolution with mutations driven by control costs", mimeo, Center for economic research, Tilburg University

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