Unfolding social hierarchies[*]

Fernando Vega-Redondo[†]

Facultad de Económicas
and
Instituto Valenciano de Investigaciones Economicas
Universidad de Alicante
03071 Alicante (Spain)

October, 1998
This version: May, 1999

Running title: Unfolding Social Hierarchies

---

[*] I am grateful to Subir Chattopadhyay, an anonymous referee and the associate editor for very helpful comments. Financial assistance by the Spanish Ministry of Education, CICYT project no. PB 97-0131 is also gratefully acknowledged.

[†] Phone no.: 34-965-90-36-16; fax no.: 34-965-90-36-85; e-mail: vega@merlin.fae.ua.es.

Unfolding social hierarchies

Fernando Vega-Redondo
University of Alicante

Abstract

Consider a large population of finitely-lived agents organized into n different hierarchical levels. Every period, all those placed at each level are randomly matched to play a given symmetric game. Based on the resulting outcome, a ρ-fraction of agents who (within their own level) attain the highest payoffs are promoted upwards. On the other hand, newcomers replacing those who die every period enter at the lowest level and choose irreversibly the strategy to be played for the rest of their life. This choice is made, with some noise, by imitating one of the strategies adopted at the highest level.

In this setup, the unique long-run behavior of the system is fully characterized for the whole class of 2×2 coordination games and two alternative variations of the model. The results crucially depend on the key "institutional" parameters ρ and n. In particular, it is shown that inefficient behavior prevails in the long run (even when risk-dominated) if promotion is only mildly selective – high ρ – and the social system is quite hierarchical – large n: In a stylized manner, these parameter conditions may be viewed as reflecting a sort of institutional deficiency that impairs economic performance.

Keywords: social learning, equilibrium selection, social hierarchy, coordination games.

1 Introduction

Social learning in population games has been an important topic of research in recent years. A significant part of this literature has pursued an evolutionary approach,[1] agents being assumed only boundedly rational when adjusting their behavior over time on the basis of relatively simple rules. These models display two important features:

(1) They assume (at least implicitly) that agents have access to relevant payoff information which is used to guide their behavioral adjustments. For example, it is typically supposed that agents can either compute a best response to the current situation, or they can imitate those actions with highest payoff performance.

(2) They formulate stylized theoretical frameworks which either display little "institutional" structure (e.g. a globally uniform matching scenario)[2] or have their potential institutional richness masked by a reduced-form specification and some abstract requirement of "monotonicity" (e.g. the Replicator Dynamics and some of its generalizations).

In the present paper, I propose a significant departure from traditional evolutionary analysis in each of two former respects.

First, in contrast with (1), it is assumed that the payoff achieved by the different strategies is not directly observed. Instead, only some indication of their relative performance is obtained through information regarding who has proven to be more successful in rising through the social hierarchy (see below). The motivation here is that, in many cases of interest, it is realistic to assume that the payoff achieved by some strategies can only be observed indirectly. Indeed, I would like to argue that the more complex and multifaceted a "strategy" is, the harder it usually becomes to attain a clear-cut assessment of its benefits. (For example, think of how difficult it is to evaluate the life-time benefits of a certain choice of education directly, i.e. through a clear observation of its "payoff".)

Regarding (2) above, the major novelty displayed by the present framework is embodied by the assumption that individuals are organized and interact within an evolving social hierarchy

---

[1] A very partial list includes Foster and Young (1990), Fudenberg and Harris (1992), Kandori, Mailath and Rob (1993), Young (1993), Samuelson (1994), and Bergin and Lipman (1996).

[2] There are some important exceptions where the (fixed) interaction pattern is assumed to be local (e.g. Ellison (1993)) or even flexible and endogenously determined through the agents' own adjustment of their, say, location decisions (e.g. Mailath, Samuelson and Shaked (1995), and Ely (1995)). Unlike these papers, our focus here is on a hierarchical segmentation in the interaction pattern, the "re-location" experienced by agents being imposed on them by the mechanism of social promotion. See below for details.

displaying n different levels. Specifically, it is postulated that agents, finitely lived, enter the population through level zero and select a certain strategy to be adopted for the rest of their life. In every period, all those agents who remain "eligible" within each level are internally matched to play a certain bilateral game. Then, on the basis of the resulting outcome, a given fraction $\rho \in (0;1)$ is chosen to rise one step in the hierarchy, the individuals so promoted being essentially chosen among those whose payoff ranked highest within that level.

The paper's main theoretical concern dwells on the question of whether the kind of payoff-responsive institutions described (as parametrized by ρ and n) may prove effective in compensating for the assumed lack of payoff observability and lead the population towards efficient social behavior. Of course, a key point here pertains to how incoming agents are assumed to choose their life-long strategies. In this respect, two alternative formulations will be considered, both reflecting some notion of upper-level imitation. The first one assumes that the continuum population of newcomers essentially reproduces (deterministically) the profile prevailing at the highest level of the hierarchy. In contrast, the second formulation is of a stochastic nature. It postulates that unbiased imitation of the upper level is independently channelled through a finite number of lower-level homogeneous groups that partition the newcomer population. As will be explained, these two alternative formulations reflect a different view of how the process of socialization is conducted.

Combining the different considerations explained, we are led to process of social learning that embodies inter-level imitation, intra-level interaction, and payoff-responsive promotion. These are the core components of the dynamic model studied in this paper. However, for both conceptual and technical reasons (e.g. to ensure the ergodicity of the process), it will also be assumed that imitation is perturbed occasionally (i.e. with very small probability) in a stochastic and time-independent manner.

The model outlined has been inspired by some recent work of Harrington (1995). Three crucial differences with it are as follows.

First, Harrington contemplates a model with no strategic interaction: the payoff achieved by each agent only depends on her own action. Therefore, in contrast with our emphasis here, his focus is not on games or issues of equilibrium selection.

Second, Harrington postulates that promotion to upper levels is the outcome of a pairwise contest, agents at each level being matched in pairs to this effect. (Hence, in particular, always half of the population at each level is promoted.)

Finally, the framework studied by Harrington does not allow for stochastic perturbations to play any role in the motion of the system. Therefore, one cannot not obtain the ergodicity conclusions (i.e. independence of initial conditions) which are the essence of our equilibrium-selection results.

The fact that stochastic perturbations may play a fruitful role in tackling issues of equilibrium selection is one of the main ideas underlying modern evolutionary literature. Thus, in this respect, the present paper borrows from this literature (in particular, Kandori et al (1993) and Young (1993)) some of its basic ideas and techniques. In a somewhat more distant vein, the approach pursued in this paper is also reminiscent of the recent literature that has fruitfully introduced considerations of social status into processes of accumulation and growth. The common starting point of the different papers in this line of research is the explicit postulate that agents display a certain (monotone) preference for status. However, the specific mechanism by which agents are supposed to attain status in each particular model is often quite different. For example, in Cole, Mailath and Postlewaite (1992),[3] status is the outcome of a social contest for "mates" that is resolved on the basis of relative wealth. In contrast, Fershtman, Murphy and Weiss (1996) postulate that the status associated to a certain action (in their case, the choice of occupation) is given by the average level of human capital prevailing among those who adopt this action. Naturally, when these diverse considerations are embedded in an intertemporal context, the way in which they affect agents' decisions is also quite diverse, e.g. they reinforce capital accumulation in Cole et al.(1992), or may introduce potentially harmful distortions that are detrimental to growth in Fershtman et al. (1996).

In the present framework, if the current "status" of an agent is identified with her position in the hierarchy, the role it plays in the dynamics of the model is quite different from that of the above mentioned literature. Here, status only acts as a signalling (or socialization) device that directs the adoption decisions of newcomers. As explained in more detail below, it may be natural to suppose that such notion of status has a tangible effect on payoffs (e.g. it could be postulated that the "game" played at higher levels is unambiguously better). However, even though this would be an appealing motivation for the imitation rule formulated for newcomers, it is not at all a required feature of the model.

I end this Introduction with a summary of the results and a brief discussion of how they compare with those of received evolutionary literature. As in much of this literature, the focus here is on symmetric 2×2-coordination games, i.e. games with two pure-strategy symmetric

---

[3] See also the later papers of Cole, Mailath, and Postlewaite (1998) and Corneo and Jeanne (1998).

equilibria.[4] In this context, it is first shown that a unique outcome is obtained in the long run as the stochastic perturbation becomes arbitrarily small in a suitable sense (i.e. a unique configuration is observed almost surely along any sample path). Moreover, such unique long-run outcome can be identified with one of the equilibria of the underlying game, every player in it choosing a common (equilibrium) strategy. Whether or not this strategy is the efficient one depends on a variety of different features of the environment: the selection rate ½ (recall above), the height of the social hierarchy n, and the off-equilibrium payoffs prevailing in the game.

The simplest case arises when the promotion conditions are quite demanding; specifically, when ½ < 1=2: Then, the efficient strategy (i.e. the strategy played in the efficient equilibrium) is always selected in the long run, independently of any other considerations – in particular, irrespectively of what are the off-equilibrium payoffs (e.g. whether or not the efficient strategy is risk dominant) or the height of the hierarchy.

However, in the polar case where ½ > 1=2 (that is, when social promotion is not quite as demanding), the long run prediction crucially depends both on off-equilibrium payoffs and on certain details of the socialization process. Concerning payoffs, the most interest situation arises when the off-equilibrium payoff for the efficient strategy is lower than for the inefficient one, a condition obviously implied by the (stronger) requirement of risk dominance for the inefficient strategy. In this case, this strategy is always found to be selected in the long run if the socialization scenario is of the first type, i.e. the profile of newcomers reproduces the situation prevailing at the highest level. When, alternatively, the socialization process is of the second kind, i.e. location-based and genuinely stochastic, the selection of the inefficient strategy is seen to occur only if, in addition to the payoff conditions mentioned and having ½ > 1=2; the social system is sufficiently hierarchic (i.e., it displays a large enough number of social strata).

These conclusions contrast sharply with those found in previous literature in at least two important respects:

(a) First and foremost, the issue of whether efficiency is achieved in the long run hinges upon certain "cultural" or "institutional" features of how society is organized (specifically, the strin-

---

[4] In still ongoing research, I have analyzed the implications of the present approach for the complementary family of 2£2-games that embody a problem of "asymmetric coordination" (e.g. the familiar Hawk-Dove game). As it turns out, one also obtains in these cases a unique long-run prediction that depends on interesting ways on its underlying parameters. However, in contrast with the present context, the nexus between efficiency and ½ may be quite complex, generally not leading to equilibrium behavior of the underlying game.

gency of its promotion requirements and the height of its hierarchic structure).

(b) Second, when such institutional features allow for the possibility of long-run inefficiency, the relevant criterion on payoffs which leads to its materialization is both qualitative in nature (i.e. ordinal) and significantly *weaker* than the customary notion of risk dominance.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 carries out the analysis. Section 4 concludes with a summary and further discussion of related literature. For the sake of smooth exposition, the formal proofs are contained in the Appendix.

2 The model

Consider a large population of agents, with the cardinality of the continuum, distributed in $n+1$ different social strata (or levels). At every $t = 1, 2, \ldots,$ those in each level $k = 0, 1, 2, \ldots, n$ are assumed randomly matched in pairs to play a bilateral symmetric game with strategies and payoffs as described in the following table:

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

Table 1

This game is assumed to be of the coordination type, i.e. $a > c$, $b > d$. Furthermore, for the sake of concreteness, we take $a > b$, i.e. strategy A is the efficient (equilibrium) strategy.

For analytical simplicity, it is assumed that, at every level $k = 0, 1, 2, \ldots, n$, agents are divided into a finite set of "locations". Furthermore, it is notationally convenient to posit that there are the same number r of locations at each level, although this is immaterial to the analysis.[5] Each location is assumed of equal relative size (i.e. includes the same fraction of players within their respective level) and homogeneous (every individual in it adopts the same action). Thus, by normalizing the measure of newcomers to one, the relative measure of agents present in each location at every level is given by $\nu \equiv 1/r$.

Locations may be conceived as clubs (or neighborhoods) where only individuals alike (or compatible) gather. The main motivation for this construct is technical rather than conceptual, i.e. it represents a convenient modelling simplification that permits a finite-state representation of the process. If the number of locations is large (an assumption that will be made throughout), the induced framework may be viewed as approximating a context with infinitesimal

---

[5] If, more generally, one allowed for a possibly varying r_k at each level k, the lower bounds on location numbers contemplated in our results would concern $r \equiv \min_k r_k$. The promotion rules would also have to be adapted accordingly, their indivisibility-induced adjustments tailored to the level in question.

locations/agents.[6] However, as explained below (cf. the end of this section), it may also be quite natural to interpret each of the $r \cdot (n+1)$ locations as the $n+1$ hierarchically arranged levels of r separate "organizations", e.g. firms. With this latter interpretation, the model could be understood as describing the joint evolution of a finite number of organizations where promotion at each level is linked to relative performance.

Agents are assumed to live for n periods.[7] At every t, a continuum of them (of measure one) start their life by being evenly distributed among the r locations at level zero. Once there, they adopt a certain strategy that, as indicated, is common to all individuals in that location. This strategy is conceived as a life-long behavioral trait for every newcomer, e.g. the learning a profession or the adoption a set of "values". It determines (together with the corresponding opponents' strategies) the payoffs earned throughout her whole life, thus underlying her success or failure in escalating along the social hierarchy.

For every t and k, denote by $\omega_k(t) \in \Gamma_r \equiv \{0, \nu, 2\nu, \ldots, r\nu (= 1)\}$ the frequency of locations at level k which adopt strategy A at t (thus, $1 - \omega_k(t)$ adopt strategy B). The dynamics of the process (i.e. the law of motion for each ω_k) can be decomposed into two components: (i) the socialization mechanism by which newcomers choose their strategy; (ii) the interaction and promotion mechanisms prevailing at each level. They are addressed in turn.

(i) Socialization mechanism

Concerning socialization, it will be postulated that newcomers adopt their action by mimicking one of those chosen at the highest level in the hierarchy. This is the stylized approach typically proposed by socio-biological models of cultural evolution (see, for example, the classical work by Boyd and Richerson (1985)). In their language, it amounts to identifying those agents who occupy the highest position in the social hierarchy as the "role models" shaping

---

[6] As r grows unboundedly, the induced framework becomes one where each agent's location is of almost negligible size $1/r$. In Vega-Redondo (1997), I study directly the limit scenario with a continuum of agents where each individual defines by herself a "location" of infinitesimal size. Even though the mathematics are then more involved, the gist of the analysis is as in the present paper for the case where socialization is modelled deterministically – cf. (S.1) below. The key point to note in this respect is that, under (S.1), the order of limits in r and ε (the perturbation probability) can be interchanged without affecting any essential feature of the model. Thus, the cardinality of locations (finite or infinite) is immaterial for the analysis. Clearly, such a limit interchange cannot be performed inconsequentially under (S.2) below (e.g. $r \to \infty$ must be performed after $\varepsilon \to \infty$ if socialization is to remain inherently stochastic), which explains the different results obtained in this alternative case.

[7] In fact, the only essential requirement is that agents live long enough to be able to reach the nth level of the hierarchy on which socialization is based. Interaction and promotion could proceed beyond this point but it would be irrelevant for the dynamics of model.

social learning.

Since individuals are taken to occupy homogeneous locations, it is convenient to describe socialization as a location-based phenomenon.[8] Two alternative formalizations of it will be considered here, with interestingly different implications. The first one postulates that the action frequencies materialized among lower-level locations exactly reproduce those prevailing at the highest level. That is:

$$\text{(S.1)}\quad \forall t = 1, 2, \ldots, \qquad \omega_0(t) = \omega_n(t).$$

The formulation embodied by (S.1) rigidly attributes to each location at the highest level an identical ex post influence on the socialization at the lowest level. For want of a better name, it will be labelled rigid socialization.

Alternatively, one could insist that inter-location symmetry should prevail only ex ante for some unbiased (stochastic) mechanism, whose realized outcomes might possibly be quite asymmetric ex post. This alternative approach will be formulated as follows: at every t, any particular lowest-level location simply mimics the action adopted by one at the highest-level, each of the latter locations selected with equal probability. This gives rise to what will be labelled flexible socialization. Formally, it amounts to assuming that, at each t, the profile $\omega_0(t)$ is determined stochastically with the following binomial probabilities:

$$\text{(S.2)}\quad \forall t = 1, 2, \ldots, \; \forall s = 0, 1, 2, \ldots, r, \quad \Pr\{\omega_0(t) = s\gamma\} = \binom{r}{s} (\omega_n(t))^s (1 - \omega_n(t))^{r-s}.$$

The alternative formulations (S.1) and (S.2) may be interpreted as reflecting different socialization environments. For example, (S.1) could represent a context where socialization is a "responsibility" equally shared by the different "senior locations" at the highest-level. In contrast, (S.2) reflects a situation where socialization is subject to a much less rigid structure and therefore displays some room for randomness and unpredictability. This contrast will be further clarified below in terms of the simple illustration closing this section. As advanced, their implications will be found to be quite different. Whereas the relative socialization rigidity embodied by (S.1) may lead to the consolidation of long-run inefficiencies (cf. Theorem 1 below), the enhanced flexibility displayed by (S.2) guarantees long-run efficiency, at least if the social hierarchy is not too tall (Theorem 2).

---

[8] Equivalently, one could describe socialization as an individual-based phenomenon (i.e. independently operating on each newcomer) provided that once such a socialization is complete, alike individuals gather in homogeneous locations.

(ii) Promotion mechanism

Concerning promotion, I shall follow Harrington (1995) in making the simplifying assumption that, at each $t,$ only those agents who have not yet suffered a set-back in their promotion at some earlier time are still "in the race". This is equivalent to the rather realistic notion that promotion remains a possibility only for those agents who are of the "right" (i.e. youngest) age for the level in question. The rest of agents are taken out of the system, in that they are never re-considered for promotion again or they interact with those still eligible for it.

Consider the set of agents who, at each level $k = 0, 1, \dots, n-1,$ are still part of the hierarchical system (i.e. have always been promoted thus far). Their contest for further promotion is structured as follows. First, agents are randomly matched within their respective level to play the game. For simplicity, it will be assumed that the matching mechanism ignores location boundaries, even though in some cases (cf. the illustration of the model outlined below) it could be reasonable to posit instead that agents in each location are only matched with those belonging to other locations at the same level. Since location size ν will be assumed "small enough", either of these two alternative formulations turn out to display identical implications. The really key assumption made on the matching mechanism is that no aggregate uncertainty exists on its induced outcome. Heuristically motivated by the large numbers involved,[9] this assumption conveniently limits all stochastic features of the model to its socialization component. More precisely, we postulate:

(M) Given the strategy profile played by the population, the induced matching frequencies associated to each strategy pair at every level coincide with the respective ex ante probabilities induced by uniform matching.[10]

After play has been conducted and the payoff profile resulting from it realized, the promotion mechanism operates. This mechanism selects, among the agents currently placed at each level $k < n,$ a ρ fraction of them $(0 < \rho < 1)$ to be promoted to the next upper level $k + 1.$ Heuristically, one would like to postulate that the set of promoted players is selected so that the following condition is satisfied:

---

[9] Here, we abstract from the well-known technical problems raised by this approach in a continuum context. See, for example, Feldman & Gilles (1985) or Judd (1985) for an elaboration on these issues, and Alós-Ferrer (1999) for a recent proposal on how to address them.

[10] Thus, suppose that the A-frequency at some level k is ω_k. Then, one would have a frequency $(\omega_k)^2$ of A players matched among themselves, $(1-\omega_k)^2$ of B players matched among themselves, $\omega_k(1-\omega_k)$ of A players matched with B players, and $(1-\omega_k)\omega_k$ of B players matched with A players. Of course, all of these numbers add up to one.

$(P)^0$ If an agent with payoff ¼ is promoted from level k to k + 1; every other agent with a strictly higher payoff must be promoted as well.

However, due to the indivisibilities resulting from our assumption of homogeneous locations, the above condition must be relaxed as follows:

(P) If a certain agent with payoff ¼ is promoted from level k to k + 1, there cannot be a set of agents at level k whose relative measure (i.e. within k) is ½° and satisfy:

(a) none of them has been promoted;

(b) they all adopt the same action;

(c) they all have obtained payoffs (possibly different among themselves) higher than ¼:

To understand this condition, note that the fraction (i.e. relative measure) of agents present at any given level k < n who may jointly occupy some location at k + 1 is simply ½° – i.e. the absolute sizes of locations across consecutive levels are scaled down by the same factor as the corresponding populations. Thus, in this light, (P) simply amounts to the requirement that no player should be promoted to an upper level location if it is possible to fill that location homogeneously with players who have obtained a higher payoff. Contrasting $(P)^0$ and (P), note that the fraction of population which could be promoted under one rule but not the other is of order ½°: Thus, for large r(= 1=°); both (P) and $(P)^0$ would behave in an almost identical manner.

Formally, the process described by (S.1) or (S.2), (M), and (P), can be represented as a Markov chain on the finite state space -; whose typical element $! = (!_1; !_2; :::; !_n)$ specifies the fraction of locations $!_k$ 2 $¡_r$ where strategy A is played at each level k (note that $!_0$ is a random variable only dependent on $!_n$). Clearly, the long-run behavior of this process is crucially dependent on initial conditions. (For example, both of the states where only one fixed strategy is played at every level is stationary.) To tackle such multiplicity of possible long-run predictions, we shall pursue the customary approach of focusing on a "slight" perturbation of the original process.

For simplicity, it will be assumed that only the socialization component of the process is perturbed. More precisely, we shall postulate that, at every t; the newcomer profile induced by socialization (i.e. by (S.1) or (S.2) in each of the two scenarios considered) is subject to change at each separate location with some (time- and location-) independent probability " > 0: If such a perturbation materializes at any given location, it has the consequence of changing the strategy $s \in \{A, B\}$ originally prevailing at it to the alternative strategy $s' \neq s$, which is then adopted homogeneously by every individual at that location. Heuristically, we may think of this perturbation as the result of unmodelled background noise that affects the fidelity of the socialization process. From an analytical point of view, it will play the useful role of making the process ergodic, thus removing any long-run dependence of initial conditions.

Analogously, one could allow for additional noise to operate in the promotion mechanism, some agents being promoted with small positive probability to the upper level even if they are not "entitled" to it given their current relative performance. More precisely, it could be postulated that with (time-, location-, and level-independent) probability ε; locations may be occupied at each level $k \geq 1$ with players originating from the lower level whose payoffs were not among the ρ best (allowing for indivisibilities). As long as this additional source of noise were restricted to the promotion mechanism alone,[11] it may be easily verified that our analysis of the process would remain essentially unaffected. Thus, for the sake of modelling simplicity, we shall restrict perturbations to operate only on its socialization component, as described above.

The model described may be conceived as a stylized representation of a population strategic context where agents' interaction is segmented by age, individuals earn their payoffs through "small-scale" interaction with a few others, and the promotion mechanism ignores (or is unable to depend upon) the specific circumstances faced by each agent. Very schematically, these features appear to be relevant features of many social systems, e.g. an inter-firm labor market for managers, the internal dynamics of a large firm, or even some academic environments. To illustrate matters, I close this section with a brief elaboration on the first of these contexts: a labor market for firm managers.

---

[11] In principle, action flexibility – via socialization or "mutation"– could persist for more than one period, but then the height of the hierarchy should be evaluated from the last period where such flexibility is possible. What is essential for the analysis is that the model displays a clear-cut dichotomy between early-life (individual-based) choice and later-life (institution-based) selection. For otherwise, if either socialization or mutation can introduce fully new behavior at any level (i.e. behavior not formerly present at the immediately lower level), the system would behave in effect as a two-level system and the height of the hierarchy would lose much of its relevance under (S.2). In particular, Part (b) of Theorem 3 would fail to hold under these circumstances. However, it is worth noticing all our other results – in particular, the selection of efficient behavior under (S.2) in Theorem 2 and the conclusions established by Theorem 1 under (S.1) – still apply if perturbations of comparable magnitude may operate unrestricted (i.e. possibly introducing totally new behavior) at any level.

A stylized illustration of the model

Consider a population of r equal-sized firms, whose "managers" are pyramidally organized into $n+1$ levels. Specifically, suppose that, within any given firm, managers of seniority $k = 1, 2, \ldots, n$ are placed at level k; the total number of these being a ρ fraction of those at $k-1$. Junior managers (i.e. those with "no seniority") enter the firm at $k = 0$.

At each t, managers of any seniority are involved in the development of an inter-firm project to be performed in collaboration with one other manager randomly chosen from another different firm. To carry out this project successfully, each manager must rely on, say, the technical support afforded by her own firm. However, this support can only be used effectively if her "skills" (here identified with either of the two actions in the abstract model, A or B) are compatible with (for simplicity, the same) as those displayed by the other managers of her own firm at the same level.[12] This motivates the assumption that, when a firm hires managers at any given level (see below), it always chooses a collection of individuals with identical skills (i.e. all A or all B).

Out of the collaboration between any two managers of different firms, the payoffs derived are as given by Table 1. Thus, this collaboration is taken to be hampered when the two managers are not coordinated on the same action. By way of illustration, we may think of their joint project as exploiting some potential complementarities across the two firms/managers involved, its eventual success hinging upon how similar the skills (technologies or cultures) of the two firms are at the relevant level.

Over time, managers are promoted on the basis of economy-wide relative performance. Indeed, since the labor market for managers is assumed fully flexible, every firm is allowed to find its new managers at any level among those who have been successfully promoted in the economy at large (thus, even though promoted managers could have been employed in the same firm before, this is not necessary). Overall, those promoted end up being the ρ-fraction who performed best at the lower level, only "minor" deviations required in order to achieve the desired internal compatibility of skills within each firm.

Finally, we turn to the process of socialization i.e. the mechanism by which the junior managers choose their skills. In this respect, of course, the two alternative formulations con-

---

[12] For example, one could suppose that the internal operations conducted within each firm require the use of identical "protocols" if internal communication is to be useful, or that they are undertaken through technologies displaying acute "O-ring fragilities" in case of overall incoherence (cf. Kremer (1993)). Of course, such a schematic view of a "firm" ignores, among other things, that there are considerations concerning across-level compatibility that will typically be very important as well.

templated above, (S.1) or (S.2), lead to quite different interpretations. The first one, (S.1), can be understood as a context where junior managers enter a firm with no prior skills and acquire them by imitating the top managers of their own firm. In contrast, the socialization mechanism embodied by (S.2) does not reflect "in-house training". Instead, it is a much less rigid (although unbiased) mechanism, top managers only having an equal ex-ante weight with possibly quite asymmetric ex-post influence. (For example, this influence could depend on the extent to which they can affect the curriculum of the "business schools" that train junior managers.) In both cases, (S.1) or (S.2), the outcome of socialization is assumed subject to exogenous random perturbations, whose main role is the introduction of new skills and could thus be interpreted as skill "innovations".

3 Analysis

Given any $\varepsilon > 0$; the (perturbed) process is clearly aperiodic and positively irreducible. The latter simply follows from the fact that there are some states (for example, the state $\omega^A \equiv (1, 1, \ldots, 1)$ where strategy A is played at each level) that can be reached from any other state after a suitable chain of perturbations. Let $\Delta(\Omega)$ denote the set of Borel probability measures on Ω: By standard results in the theory of Markov chains, there is a unique invariant distribution in $\Delta(\Omega)$ that summarizes the long-run behavior of the process, independently of initial conditions. This distribution will be denoted by μ_ε, in order to account explicitly for its dependence on ε:

Heuristically, we want to conceive of ε as small and study the behavior of the process under these circumstances. Formally, this is done by focusing on the limit invariant distribution:

$$\mu^* \equiv \lim_{\varepsilon \to 0} \mu_\varepsilon ;$$

that will be shown to be a well defined element of $\Delta(\Omega)$. The states in its support are labelled stochastically stable states.

Our aim is to characterize μ^* and explore its dependence on the key parameters of the model, i.e. the promotion rate ρ and the height of the hierarchy n: Depending on the payoff configuration of the underlying game (cf. Table 1), the analysis may be decomposed into two (generic) cases: one where $d > c$; and another where $d < c$: In the former case, the criteria of efficiency and risk dominance coincide, i.e. the efficient strategy A is also risk-dominant in the sense of Harsanyi and Selten (1988). A conflict between these two alternative criteria can only occur (although not necessarily) if $d < c$. Since this latter context is where the more interesting conclusions arise, our detailed analysis will be mostly restricted to it, the results applying when $d > c$ only briefly outlined at the end of our discussion.

First, we deal with the case where socialization is rigid, as formalized by (S.1). The following result, proven in the Appendix, characterizes the long-run behavior of the process in this case.

Theorem 1 Let $d < c$ and consider the process defined by (S.1), (M), and (P). Given any $\rho \in (0,1)$, there exists some $\bar{r} \in \mathbb{N}$ such that, provided $r \geq \bar{r}$,
(a) if $\rho < 1/2$, $\mu^*(\omega^A) = 1$, where $\omega^A = (1,1,\ldots,1)$;
(b) if $\rho > 1/2$, $\mu^*(\omega^B) = 1$, where $\omega^B = (0,0,\ldots,0)$.[13]

The previous result establishes that, provided the number of groups is sufficiently large, the long-run behavior of the system crucially depends on the institutional parameter ρ: If $\rho < 1/2$ (i.e. promotion is quite selective), the unique long-run outcome involves agents playing the efficient strategy at every level. Instead, when $\rho > 1/2$ (promotion is not very selective), the unique long-run outcome has everyone playing the inefficient strategy at each level.

To understand the intuition underlying these conclusions, it is useful to focus on the simplest context where $n = 1$, i.e. the hierarchy consists of just two levels. In this case, it is easy to see that, in terms of the unperturbed dynamics, there are only two states which display basins of attraction of positive measure. They are the two homogeneous states (ω^A and ω^B) where all individuals at level 1 (and therefore at level zero as well) play a common strategy. Then, as customary in recent evolutionary theory, the key subsequent task involves an assessment of what are the relative sizes of their respective basins of attraction. In fact, since the state space may be heuristically conceived as one-dimensional, the issue can be essentially reformulated as that of identifying where lies the "point" which separates both basins of attraction – or, somewhat differently, in which of the two basins lies the "intermediate" state $\hat{\omega}$ $(= \hat{\omega}_1) = \frac{1}{2}$ (for simplicity, r is taken to be even).

In state $\hat{\omega}$, half of the individuals at level zero choose each of the two strategies. Thus, after they are randomly matched, the resulting payoff profile involves half of the A strategists obtaining each of the two possible payoffs (a and d) and half of the B strategists obtaining each of the two possible payoffs (c and b). Thus, if $\rho < 1/2$, more than half of the population promoted to level 1 are A strategists (all those who obtain a), which implies that the fraction

---

[13] The analysis would remain unaffected if ρ could vary across levels but the same strict inequality relative to $1/2$ would apply at all of them. On the other hand, in the knife-edge case where $\rho = 1/2$, the set of stochastically stable states consists of all those where, at every level, any of the two actions is homogeneously played (possibly a different one at each level).

of this type must increase next period. Instead, when $\rho > 1/2$; the opposite happens. Now, those promoted include all B strategists who have obtained a payoff of b and, since $c > d$; also some of those who have obtained c: Thus; it follows that those A strategists who have obtained d may be promoted only if all B strategists (either with a payoff of b or c) are promoted as well. Clearly, this implies that more than half of the population promoted is a B strategist, which raises the frequency of this type next period.

The above considerations suggest that, provided $c > d$; state $\bar{\omega}$ must belong to the basin of attraction of ω^A or ω^B if, respectively, $\rho < 1/2$ or $\rho > 1/2$: In fact, it may be computed (cf. Lemmas 1 and 4 in the Appendix) that, if the number of locations is sufficiently large, the basins of attraction of ω^A and ω^B are (roughly) constituted by those states were the fraction of A strategists is respectively larger or smaller than ρ. Thus, it follows that whether the basin of attraction of ω^A is larger or smaller than that of ω^B essentially depends on ρ being smaller or larger than $1/2$:

The previous discussion can be readily extended to any finite $n > 1$: For, when socialization is governed by (S.1), the generation of newcomers entering at any given period may be conceived as belonging to one of n separate cohorts, promotion and socialization operating on each of them in a fully independent manner. Thus, concerning promotion, the selection mechanism can be seen as applying in "parallel" for every such cohort, operating on the single generation that represents it at any t. On the other hand, pertaining socialization, note that each cohort is fully renewed every n periods, its newcomers being socialized by those who currently occupy the uppermost level and also belong to it. With this interpretation of our setup in mind, the former heuristic argument can be applied to any n-level hierarchy. This indicates that, again, ω^A or ω^B should display the largest basin of attraction depending on whether ρ is smaller or larger than $1/2$: Therefore, in each of these cases, the respective state will be observed "almost always" in the long run, as the magnitude of the perturbation becomes small and escaping any basin of attraction becomes progressively more difficult.

In contrast with the previous conclusion, the analysis is now shown to be significantly different when the socialization process is flexible (i.e. is given by (S.2)). For, in this case, the number of hierarchic levels turns out to play a very crucial role in the long-run dynamics of the process, as indicated by the following result.

Theorem 2 Let $d < c$ and consider the process defined by (S.2), (M), and (P). Given any $\rho \in (0;1)$; there exists some $\bar{r}; \bar{n} \in \mathbb{N}$ such that, if $r \geq \bar{r}$ and $n \leq \bar{n}$; $\mu^*(\omega^A) = 1$:

The previous result establishes that, no matter how mildly selective promotion is (i.e. how large is ρ), a sufficiently short hierarchy ensures long-run efficiency if the socialization process is given by (S.2). Intuitively, this result reflects the combination of the following two ideas.

On the one hand, if the social hierarchy is short, a small fraction of "lucky" A strategists (i.e. those who meet their own type) can rise through it, even if their initial frequency is low. This cannot happen if there are many repeated operations of the promotion mechanism (i.e. many levels) since, in that case, the frequency of consistently lucky individuals is bound to be too small to sustain even a single location. In contrast, this argument cannot be applied to a small initial frequency of B strategists. The fact that B is the inefficient strategy implies that, even with a single operation of the promotion mechanism, no B strategists will be promoted when most of the population is composed of A strategists obtaining the maximum payoff $a.$

The previous considerations then need to be combined with the fact that (S.2) allows positive socialization probability to any location at the highest level. Thus, while a small perturbation of an originally homogeneous B-profile allows for the possibility that future newcomers adopt A through socialization, the converse is impossible since no B strategists will reach the highest level if they start in small frequency. In a sense, we may interpret this as an indication of the benefits that may be afforded from the persistence of some amount of noise (here, on socialization) when combined with a selection process that is not too stringent (here, a promotion mechanism that does not involve too many repeated rounds).

The central role played by the height of the hierarchy under flexible socialization is further underscored by the following result.

Theorem 3 Let $d < c$ and consider the process defined by (S.2), (M), and (P). Given any $\rho \in (0,1)$, there exist $\bar{r}$ and $\bar{n}(r)$ such that, if $r \geq \bar{r}$ and $n \geq \bar{n}(r)$,
(a) $\rho < 1/2 \Rightarrow \mu^*(\omega^A) = 1$;
(b) $\rho > 1/2 \Rightarrow \mu^*(\omega^B) = 1.$

In contrast with Theorem 2, the previous result indicates that, if the number of hierarchical levels is sufficiently large, (S.2) leads to the same conclusions as (S.1). In particular, the selection of the efficient or inefficient strategy in the long run hinges upon whether the promotion mechanism is sufficiently selective ($\rho < 1/2$) or not ($\rho > 1/2$). The essential intuition underlying this result is a combination of ideas already explained for Theorems 1 and 2. On the one hand, if the social hierarchy is high enough, no strategy whose initial frequency for a certain cohort lies outside the relevant basin of attraction may hope to sustain any viable representation at the highest level. Thus, under these circumstances, the stochastic socialization mechanism contemplated by (S.2) does not introduce any considerations different from those derived from a direct comparison of basins of attraction. These are precisely the considerations reflected by Theorem 1, which explains the coincidence between the conclusions established by this latter result and Theorem 3.

Our analysis has focused on the payoff scenario with $d < c$ where the more interesting considerations arise (recall the discussion preceding Theorem 1). However, for the sake of completeness, we end this section with the statement of the result that applies to the alternative payoff context.

Theorem 4 Let $d > c$ and assume that the socialization process satisfies (M), (P), and either (S.1) or (S.2). Given any $\rho \in (0,1)$, there exists some $\bar{r} \in \mathbb{N}$ such that, provided $r \geq \bar{r}$, $\mu^*(\omega^A) = 1$.

If $d > c$, the A strategists dominate the B strategists, both in symmetric and asymmetric encounters. Under these conditions, it is quite intuitive that, as asserted by Theorem 4, strategy A should be selected in the long-run within either of the two socialization scenarios considered ((S.1) and (S.2)) and independently of any further details of the evolutionary process (in particular, for any number of levels displayed by the hierarchy). We dispense with a formal proof of this result, since it follows from a direct adaptation of the arguments spelled out in the Appendix for the alternative payoff scenario.

4 Summary

As in much of recent evolutionary literature, our objective here has been to provide a clear-cut criterion of long-run selection in population games. However, in contrast with received evolutionary theory, the main considerations underlying this criterion are of an "institutional" nature. Essentially, two are the key parameters involved: the promotion rate ρ and the height of the social hierarchy n.

Focusing on the most interesting case where $d < c$, the effect of ρ and n on the long-run outcome has been shown to depend on the particular details of the socialization scenario. Specifically, the main conclusions can be summarized as follows.

(1) Under (S.1), where the socialization mechanism essentially reproduces in a deterministic fashion the highest-level profile, the only relevant consideration is whether ρ is large ($\rho > 1/2$) or not ($\rho < 1/2$):

(a) If $\rho > 1/2$ (i.e. if promotion is rather lax), the unique long-run state has everyone playing the inefficient strategy B.

(b) Instead, if $\rho < 1/2$ (i.e. promotion is stringent), the unique long-run state has the whole population playing the efficient strategy A:

(2) In contrast, under (S.2), where socialization is independently conducted by each lowest-level location, the analysis also depends on n, the height of the social hierarchy. Specifically:

(a) If n is small, then the efficient outcome is attained in the long run, independently of ρ.

(b) Instead, if n is sufficiently large, ρ regains a relevant role and the same conclusions of (1a) and (1b) above obtain.

These conclusions reflect institutional considerations that pertain to how high and steep is the "social pyramid" and how these two features interact with the socialization mechanism. Even though these considerations are mostly absent in recent evolutionary literature, related ideas have been explored by two recent papers in this field.

Concerning the role played by the gradient of the social pyramid (inversely related to the stringency of promotion requirements), the relevant factors in our model are somewhat reminiscent of those operating in a recent paper by Binmore and Samuelson (1997). These authors study an evolutionary model where agents are mere "satisficers", thus deciding whether or not to revise their previous strategy depending on how their current payoff compares with a given aspiration level. If this aspiration level is low enough, their model selects the inefficient equilibrium in the long run, for reasons that are similar to those reflected by Theorems 1 or 3 for small ρ (cf. their Proposition 10). In fact, in our context, the promotion rate ρ may be heuristically conceived as determining the (relatively defined) "aspiration level" required by society for promoting agents to higher levels. And, as in the model studied by Binmore and Samuelson, we reach the conclusion that a rise in "social aspirations" favors long-run efficiency because it increases the fragility of the inefficient configuration, i.e. it shrinks the latter's basin of attraction.

On the other hand, concerning the role played by random socialization (i.e. (S.2)) in leading to long-run efficiency under a short hierarchy (low n), the nature of our argument may be related to considerations studied by Robson and Vega-Redondo (1996). These authors focus on an evolutionary context of the kind proposed by Kandori et al. (1993), but with the crucial difference that genuinely random matching introduces an additional source of randomness into the process. In a certain sense, such increased randomness due to matching plays a role analogous to that induced by stochastic socialization in the present context. In Robson and Vega-Redondo (1996), few mutations and “lucky” matching may lead to the spread of efficient behavior through imitation. Here, somewhat analogous considerations arise when the hierarchy is short and socialization is stochastic. For, under these conditions, only a few lower-level locations need to be perturbed in order for further “lucky” socialization to lead to the spread and consolidation of efficient behavior through promotion and (perturbation-free) socialization. But, of course, in both contexts such relatively “easy” mutation-driven transition to the efficient state is inherently asymmetric, i.e. when only few mutations arise at the efficient state they are selected against (by either imitation or the promotion mechanism, respectively), thus proving insufficient to perform the transition away from that state. Indeed, since mutations are assumed rare, the long-run selection of efficiency in both Robson and Vega-Redondo (1996) and the present case is essentially a consequence of that asymmetry.

Appendix

Proof of Theorem 1. Our first task involves a characterization of the unperturbed dynamics of the process. To this end, I focus first on the promotion dynamics, as described by (P) above. The promotion mechanism is exactly the same at each level $k = 0, 1, 2, \ldots, n - 1$. Thus, it may formalized by a common function $g : \Gamma_r \to \Gamma_r$ that, given any fraction of A-strategists $x \in \Gamma_r$ prevailing at any such level k and some t, determines the corresponding fraction $g(x) \in \Gamma_r$ prevailing at $k + 1$ and $t + 1$.

The detailed specification of $g(\cdot)$ becomes somewhat complicated by the “indivisibilities” induced by the location construct on promotion (i.e. the need to fill homogenous locations at the upper level). However, if $\nu \equiv 1/r$ is small enough, it is shown below (cf. Lemmas 2-4) that both of these consideration play no important role. In particular, the analysis of the process can be conducted through a function $f : [0, 1] \to [0, 1]$ that abstracts from the location discreteness, both in matching and promotion. Formally, such a function may be defined as follows:[14]

$$f(x) = \begin{cases} \frac{1}{½}x^2 & 0 \quad x \quad \frac{1 ¡ \sqrt{4½ ¡ 3}}{2} & \text{(a)} \\ \frac{1}{½}(x + ½ ¡ 1) & \frac{1 ¡ \sqrt{4½ ¡ 3}}{2} < x \quad \frac{1 + \sqrt{4½ ¡ 3}}{2} & \text{(b)} \\ \frac{1}{½}x^2 & \frac{1 + \sqrt{4½ ¡ 3}}{2} < x \quad \sqrt{½} & \text{(c)} \\ 1 & \sqrt{½} < x \quad 1: & \text{(d)} \end{cases} \qquad (1)$$

To understand the expressions specified in (1), consider first Range (a). It applies when the fraction of A-strategists x satisfies $x < 1=2$, $x(1 ¡ x) \quad 1 ¡ ½$; and $x^2 \quad ½$: Within this range for x; the A players who meet B players (who receive the lowest possible payoff) display a frequency no larger than $1 ¡ ½$. Therefore, if the requirement of location homogeneity is dispensed with at the upper level, only those A players who meet players of the same type (x^2 of them if locations play no role in matching) will be promoted. Scaling this magnitude by the factor ½ at which the relevant population size decreases across consecutive levels, one is led to the expression contemplated in (a).

Next, Range (b) corresponds to those x such that $x(1 ¡ x) > 1 ¡ ½$: In this case, some of the A players who meet B players will be promoted – specifically, the fraction is $x(1 ¡ x) ¡ (1 ¡ ½)$, again abstracting from locations. When this fraction is added to x^2 (i.e. those A players meeting their own type), (b) is obtained after normalizing the result by the factor ½.

Range (c) is simply the counterpart to (a) if $x > 1=2$ and $x^2 \quad ½$: Finally, Range (d) considers those x such that $x^2 > ½$: In this case, no B players are promoted – only A players who meet their own type may be so. This implies that the frequency of A players in the upper level next period must be 1, which is the content of (d).

For our purposes, the relevant features of f(¢) are established by the following first Lemma.

Lemma 1. The set of fixed points of the function f(¢) is given by f0; 1; ½g: Moreover, for all x 2 (0; 1), we have $f(x) > x$, $x > ½$:

To verify this Lemma, first observe that f(¢) must have at least one interior fixed point since the function f(¢) specified in (1) is continuous and, from (a) and (d), it follows that there are some $x^0 < x^{00}$ such that $f(x^0) < x^0$ and $f(x^{00}) > x^{00}$: It is immediate to see that such fixed points can only exist in ranges (a) and (c), in both of which we have $f(x) = \frac{1}{½}x^2$: Since $x^¤ = ½$ is the unique fixed point in this case; the desired conclusion follows.

---

[14] Some of the Ranges (a)-(d) below may be empty or not well-defined, in which case there is a corresponding reduction in the relevant cases that need to be considered. For example, if $½ < \frac{3}{4}$; Range (b) disappears and $f(x) = \frac{1}{½}x^2$ for all $x \quad \sqrt{½}$:

Even though $f(\cdot)$ does not reflect an accurate description of the model (because it abstracts from the requirement of location homogeneity), its usefulness derives from the fact that, as advanced, it represents a sufficiently good approximation of the promotion mechanism. Specifically, we may establish the following Lemmata.

Lemma 2. $\exists \nu > 0$ such that $\forall r \in N; \forall x \in \Gamma_r; |g(x) - f(x)| \leq \nu/r:$

Lemma 3. Given ρ; there exist $\hat{r} \in N$, $\hat{x} < 1; x > 0;$ such that if $r \geq \hat{r}$;

(i) $[x \in \Gamma_r; x \geq \hat{x}] \Rightarrow g(x) \geq f(x);$

(ii) $[x \in \Gamma_r; x \leq x] \Rightarrow g(x) \leq f(x):$

Lemma 4. Given ρ and any $\delta > 0; \exists \tilde{r} \in N$ such that if $r \geq \tilde{r}$;

$[x \in \Gamma_r; |x - \rho| > \delta] \Rightarrow \operatorname{sgn}(g(x) - x) = \operatorname{sgn}(f(x) - x):$

To prove Lemma 2 note that, in view of the promotion rule (P), there can be at most a fraction of $\varepsilon \equiv 1/r$ individuals who are not promoted when others with lower payoffs are so. Thus, in terms of the induced frequencies at the upper level, one obtains

$$\max_{x \in \Gamma_r} |f(x) - g(x)| \leq 1/r: \tag{2}$$

which leads to the desired conclusion.

Next, we confirm the validity of Lemma 3. Given $\rho \in (0; 1)$; choose $\hat{r}$ large enough so that there exists $\hat{x} < 1$ and $x > 0$ such that

$$[x \in \Gamma_r; x \geq \hat{x}] \quad \Rightarrow \quad x(x - \frac{1}{\hat{r}}) > \rho$$
$$[x \in \Gamma_r; x \leq x] \quad \Rightarrow \quad x(1 - x) + \frac{3}{\hat{r}} < 1 - \rho$$

Then, if $r \geq \hat{r}$, every individual promoted when $x \geq \hat{x}$ is an A-strategist, which implies that $g(x) = f(x) = 1$ for any such $x \in \Gamma_r$: This proves Part (i) of Lemma 3. On the other hand, if $r \geq \hat{r}$, it is immediate to check that no A-strategist matched with a B-strategist is promoted if $x \leq \hat{x}$, which guarantees that no more than x^2 A-strategists are promoted. This implies that $g(x) \leq f(x) = x^2/\rho$ for such an $x \in \Gamma_r$; thus proving Part (ii).

Now, I turn to proving Lemma 4. Given ρ; fix any $\delta > 0$ and choose $\hat{r}$, $\hat{x}$; and x as in Lemma 3 so that the set $U \equiv \{|f(x) - x| : |x - \rho| \geq \delta; x \leq \hat{x}; x \geq x\}$ is non-empty. (If δ is so large that the set U is necessarily empty, the result applies voidly.) Define $\hat{u} \equiv \min\{u \in U\}$: By Lemma 1 and the continuity of $f(\cdot)$, $\hat{u} > 0$: Thus, if $\tilde{r}$ is chosen such that $\tilde{r} > \max\{\hat{r}; \nu/\hat{u}\}$; where ν is as in Lemma 2, this latter Lemma and Lemma 3 lead to the desired conclusion.

To simplify our ensuing argument, it is useful to partition the population present throughout the whole process into n infinitely-lived “cohorts”. Specifically, let

(a) the individuals placed at $t = 1$ within each one of the n levels $k = 1, 2, \ldots, n$ belong to a separate population cohort;

(b) the individuals entering the system at any t belong to the same cohort as those currently placed at level n.

Relying on the population partition induced by (a)-(b), it should be clear that the study of the evolutionary process may be performed as a juxtaposition of n fully independent dynamics, one for each of the n different cohorts. The state space for each of these dynamics may be identified with Γ_r, a typical state $x \in \Gamma_r$ indicating the frequency of individuals of the cohort who (every n periods) reach level n and adopt strategy A.

If $\varepsilon > 0$, such cohort-based stochastic process is clearly ergodic, for the same reasons as the original (multi-level, multi-cohort) process is. Thus, it has a unique invariant distribution, that will be denoted by μ_ε. Making $\varepsilon = 0$, one obtains its unperturbed counterpart, whose transition function $\phi : \Gamma_r \to \Gamma_r$ may be defined as follows:

$$\phi(x) = g \circ g \circ \overset{n \text{ times}}{\cdots} \circ g \equiv g^n : \Gamma_r \to \Gamma_r \tag{3}$$

where we implicitly introduce socialization in (3) by postulating (cf. (S.1)) that the newcomer population exactly reproduces the profile prevailing at the current highest level.

Clearly, the states $x = 0$ and $x = 1$ both define (singleton) limit sets of the unperturbed dynamics induced by ϕ. Given some $\delta > 0$, let $N_\delta(\rho) \equiv \{x \in [0, 1] : |x - \rho| \leq \delta\}$. By Lemmas 1 and 4, it follows that, if r is chosen large enough, any additional limit states of the unperturbed dynamics (if there are any) must lie in $N_\delta(\rho)$ where, of course, δ can be chosen arbitrarily small. To discriminate among such multiplicity of limit states, we now turn to an explicit analysis of the perturbed dynamics.

This analysis will rely on the well-known graph-theoretic methods developed by Freidlin and Wentzel (1984), often used in recent Evolutionary Theory. Particularized to the cohort-based dynamics considered here, they involve the following concepts.

Definition: Let $x \in \Gamma_r$. An x-tree Y is a directed graph on Γ_r (i.e. a subset of $\Gamma_r \times \Gamma_r$) such that every state $x' \in \Gamma_r \backslash \{x\}$ is the initial point of exactly one “arrow” $(x', x'') \in Y$ and from any such state x' there is a path $\{(x', x''), (x'', x'''), \ldots, (x^{(r-1)}, x^{(r)})\} \subset Y$ whose end point $x^{(r)} = x$.

For every x, $x' \in \Gamma_r$, denote by $p_\varepsilon(x, x')$ the probability of a transition from x to x' when the perturbation probability is ε. Provided $p_\varepsilon(x, x') > 0$, we define the cost of the transition from x to x' as follows:

$$c(x; x') = \min\left\{\frac{1}{\nu}\left|x - x''\right| : x'' \in \Omega_r,\ x' = \psi(x'')\right\}. \tag{4}$$

Heuristically, we may think of $c(x; x')$ as the minimum number of locations where the socialization outcome needs to be perturbed in order for the ensuing operation of the promotion dynamics $\psi(\cdot)$ to produce a transition from x to x'. If such a transition cannot occur, even resorting to suitable perturbations of the process, it will be formally convenient to make $c(x; x') = \infty$.

Given any state x, denote by $\mathcal{Y}_x$ the set of x-trees. For any such tree $Y \in \mathcal{Y}_x$, its cost $c(Y)$ is defined equal to $c(Y) \equiv \sum_{(x', x'') \in Y} c(x; x')$, i.e. the sum of transition costs across its constituent arrows.

Define $\mu^* = \lim_{\varepsilon \to 0} \mu_\varepsilon$ as the limit invariant distribution of the cohort-based evolutionary process. By direct adaptation of standard arguments (see Kandori, Mailath and Rob (1993) or Young (1993)), μ^* may be shown to be a well defined element of $\Delta(\Omega_r)$. Moreover, its support can be characterized as follows:

$$\mathrm{supp}(\mu^*) = \{x \in \Omega_r : \exists Y \in \mathcal{Y}_x \text{ s.t. } c(Y) \leq c(Y') \text{ for all } Y' \in \mathcal{Y}_{x'},\ x' \in \Omega_r\}. \tag{5}$$

That is, the stochastically stable states $x \in \mathrm{supp}(\mu^*)$ are those that display some x-tree of minimum cost across all possible such trees, $\bigcup_{x' \in \Omega_r} \mathcal{Y}_{x'}$. Based on this characterization, the main additional step in the argument is embodied by the following lemma.

Lemma 5. There exists some $\bar{r}$ such that, provided $r \geq \bar{r}$,

(a) if $\rho < 1/2$, $\forall x \in \Omega_r$, $[x \neq 1, Y \in \mathcal{Y}_x] \Rightarrow \left[\exists \hat{Y} \in \mathcal{Y}_1,\ c(\hat{Y}) < c(Y)\right]$;

(b) if $\rho > 1/2$, $\forall x \in \Omega_r$, $[x \neq 0, Y \in \mathcal{Y}_x] \Rightarrow \left[\exists \check{Y} \in \mathcal{Y}_0,\ c(\check{Y}) < c(Y)\right]$.

Consider first the case where $\rho < 1/2$. Choose δ and $\bar{r}$ in Lemma 4 such that $\rho + \delta + 1/\bar{r} < 1/2$. Then, it follows from this Lemma that, if $r \geq \bar{r}$, there is some 1-tree $\hat{Y}$ with $c(\hat{Y}) \leq (\rho + \delta)r + 1$. Specifically, consider a tree $\hat{Y}$ that has:

(T.1) every state $x \leq (\rho + \delta)$ connected to the state $x' \equiv x + \nu$ $(= x + 1/r)$;

(T.2) for all $x > (\rho + \delta)$, a path that connects x to state 1 at zero cost.

On the one hand, by construction, the arrows described in (T.1) induce a total cost no larger than $(\rho + \delta)r + 1$. On the other hand, the complementary (costless) paths described in (T.2) are possible since (by Lemma 4) any state $x > (\rho + \delta)$ belongs to the set $B_1 \equiv \{x \in \Omega_r : \psi^q(x) = 1,\ q \in \mathbb{N}\}$, the basin of attraction of state 1 in terms of the unperturbed dynamics. Overall, therefore, (T.1) and (T.2) define a 1-tree $\hat{Y}$ at the maximum indicated cost.

Now, we argue that for all $x \not= 1$; any corresponding x-tree involves a cost larger than $(½ + ±)r + 1$: First, note that, since only limit states of the unperturbed dynamics may be stochastically stable, no state in B_1 other than $x = 1$ needs to be considered: Focus, therefore, on any given state $x \not2 B_1$. For any such x and any $Y \; 2 \; Y_x$; $c(Y) \; ¸ \; (½ + ±)\, r + 2$ since, by Lemmas 1 and 4, at least $((½ + ±)\, r + 2)$ perturbations are needed to link state 1 to any such x. (Note that $((½ + ±)\, r + 1) < \frac{1}{2}r$:) Thus, $c(Y) > c(\hat{Y})$; where $\hat{Y}$ is as constructed above. This proves Part (a) of Lemma 5. To prove its Part (b), let $½ > 1{=}2$ and choose ± and $\bar{r}$ in Lemma 4 such that $½ \; ¡ \; ± \; ¡ \; 1{=}\bar{r} > 1{=}2$: Then, an argument reciprocal to the one put forward for Part (a) leads to the desired conclusion.

Lemma 5 implies that, provided r is large enough, $¸^¤(1) = 1$ when $½ < 1{=}2$ and $¸^¤(0) = 1$ when $½ < 1{=}2$: In turn, since the dynamics associated to each cohort is independent of those of others, this implies that, for large r; $^1{}^¤(!^A) = 1$ when $½ < 1{=}2$ and $^1{}^¤(!^B) = 1$ when $½ < 1{=}2$. ■

Proof of Theorem 2. As in the proof of Theorem 1, the discussion will focus on a cohort-based representation of the promotion dynamics, as captured by the same function $Á(¢)$ de…ned in (3). However, in contrast with the previous formulation of the socialization dynamics, the present one is probabilistic – cf. (S.2). Let $^3{}_x(\bar{x})$ denote the probability that the pro…le $\bar{x} \; 2 \; ¡_r$ is attained at level zero when $x \; 2 \; X_r$ prevails at the nth level. Then, for any pair of states (x, x^0) which represent a feasible transition of the process, de…ne:

$$\hat{c}(x; x^0) = \min f \frac{1}{o} \bar{\;} x^{00} \; ¡ \; x^{000} \bar{\;} : x^{00}; x^{000} \; 2 \; ¡_r; \; ^3{}_x(x^{00}) > 0; \; x^0 = Á(x^{000}) g;$$

an adapted counterpart of the cost function postulated in (4), again making $\hat{c}(x; x^0) = 1$ when no transition from x to x^0 is possible. Moreover, for any $x \; 2 \; ¡_r$ and every x-tree $Y \; 2 \; Y_x$, we de…ne $\hat{c}(Y) \; ´ \; \mathbf{P}_{(x^0; x^{00}) \, 2 Y} \; \hat{c}(x; x^0)$:

As before, the cohort-based dynamics (induced by the composition of perturbed socialization and n-repeated promotion) displays a unique invariant distribution for any $^2 > 0$: Denoting by $\hat{¸}_2$ such invariant distribution, our aim is to identify the support of $\hat{¸}^¤ = \lim_{" \, ! \, 0} \hat{¸}_"$. In analogy with (5), we have the following characterization:

$$\mathrm{supp}\,(\hat{¸}^¤) = f x \; 2 \; ¡_r : 9Y \; 2 \; Y_x \text{ s.t. } \hat{c}(Y) \quad \hat{c}(Y^0) \text{ for all } Y^0 \; 2 \; Y_{x^0}; \; x^0 \; 2 \; ¡_r g: \qquad (6)$$

Let $\hat{B}_0 \; ´ \; f x \; 2 \; ¡_r : \; ^3{}_x(x^0) > 0; \; Á(x^0) = 0 g$ and $\hat{B}_1 \; ´ \; f x \; 2 \; ¡_r : \; ^3{}_x(x^0) > 0; \; Á(x^0) = 1 g$ stand for the respective basins of attraction of states $x = 0$ and $x = 1$ in terms of the unperturbed (cohort-based) dynamics.[15] By (S.2), $\hat{B}_0$ = ¡$_r$ n f1g and $\hat{B}_1$ = ¡$_r$ n f0g; which implies that states x = 0 and x = 1 are the only limit states of the unperturbed dynamics. In view of this fact, the essential step in the argument is embodied by the following Lemma.

Lemma 6. Suppose n = 1:

(a) There exists a 1-tree Y such that ¢(Y) 2° $^{i1=2}$ + 1:

(b) Every 0-tree Y has ¢(Y) ¸ (1 ¡ $\hat{x}$)° ^{i1}, where $\hat{x}$ is as in Lemma 3.

Since $\hat{B}_1$ = ¡$_r$ n f0g, in order to prove Part (a) of Lemma 6 it is enough to determine the minimum number of locations that need to be perturbed at state x = 0 so that the resulting state x^0 has Á(x^0) 6= 0: When n = 1; it is enough that such a state x^0 satisfy x^0(x^0 ¡ °) ¸ °, which is implied by x^0 ¸ 2°$^{1=2}$: Thus, relative to state x = 0, no more than $\frac{2°^{1=2}}{°}$ + 1 locations need to be perturbed, thus leading to the conclusion stated in Part (a).

To show Part (b), note that, as explained in the proof of Lemma 3, if x ¸ $\hat{x}$; Á(x) = 1: Therefore, relative to state x = 1; no fewer than ° ^{i1}(1 ¡ $\hat{x}$) locations must be perturbed in order to reach a state x^0 with Á(x^0) 6= 1: Since $\hat{B}_0$ = ¡$_r$ n f1g; this is the minimum number of perturbations which must be included in any path linking state x = 1 to state x = 0. Thus, as claimed, this is also the minimum cost that must be incurred by any 0-tree.

By Lemma 6, if n = 1 and r (= ° ^{i1}) is large enough, there exists some Y$^¤$ 2 Y$_1$ such that ¢(Y$^¤$) < ¢(Y) for all Y 2 Y$_0$: In view of (6), this implies that $\hat{¸}^¤$(1) = 1 and therefore ¹$^¤$(!^{A}) = 1: Since this conclusion applies for n = 1; it holds of course for all n n$_0$ and some su¢ciently small n$_0$. ■

Proof of Theorem 3. For concreteness, restrict consideration to the case ½ > 1=2 (if ½ < 1=2; the argument is analogous). Given some such ½ and any ±, choose $\hat{r}$ as in Lemma 4: Then, the following result obtains.

Lemma 7. Given any r ¸ $\hat{r}$; 9$\hat{n}$(r) such that if n ¸ $\hat{n}$(r); [x 2 ¡$_r$; x < ½ ¡ ±]) Á(x) = 0:

To prove Lemma 7, …rst note that, given r ¸ $\hat{r}$; there exists some q 2 N such that, from any pro…le at level zero x$_0$ < ½ ¡ ±; an application of q consecutive rounds of the promotion mechanism leads to a qth-level pro…le x$_q$ = g^q(x) < x; where x is as in Lemma 3. Thus, by Part (ii) of this Lemma, any further application of the promotion mechanism will promote, among the A strategists, only a subset of those who meet other A-strategists. (Note that, by Lemma 3, if there is a frequency x < x of A-strategists at any level k < n, a frequency of at

---

[15] In principle, to de…ne exhaustively the basins of attraction one has to allow for any …nite number of iterations of the dynamics, as in the proof of Theorem 1. However, in the present case, this is immaterial..

most $x^2 = ½f(x)$ are promoted to level $k + 1$.) Suppose without loss of generality that $x < ½$: Then, for all $q^0 = q + 1; q + 2; :::; n;$

$$f(x_{q^0}) < \frac{x^2}{½} \qquad x$$

and, therefore,

$$Á(x_0) \;´\; g^n(x_0) \qquad \frac{\hat{x}^{2^{n ¡ q}}}{½^{(2^{n ¡ q}) ¡ 1}} = ½ \;µ\frac{\hat{x}}{½}¶^{2^{n ¡ q}} \qquad :$$

Hence, given r; there exists some large enough $\bar{n}(r)$ such that, if $n \;¸\; \bar{n}(r)$;

$$Á(x_0) < º = \frac{1}{r}$$

which implies $Á(x_0) = 0$; thus proving the Lemma.

Now, assume that ± and $\bar{r}$ have been chosen above so that $½ ¡ ± ¡ 1=\bar{r} < 1=2$: Then, along the lines of (T.1) and (T.2) above, a 0-tree $\hat{Y}$ can be constructed with $\hat{c}(\hat{Y}) \quad (1 ¡ (½ ¡ ±))r + 1$: On the other hand, from Lemmas 1 and 7, it follows that any 1-tree Y must involve a cost $\hat{c}(Y) \;¸\; (1 ¡ (½ ¡ ±))r + 2$: (Here, the argument is essentially the same as in the proof of Theorem 1.) Since states $x = 0$ and $x = 1$ are the only candidate states for stochastic stability in the present context, one may conclude that, under the conditions of Lemma 7; $\hat{¸}^¤(0) = 1$ and, therefore, $¹^¤(! ^B) = 1$: This completes the proof for the case $½ > 1=2$: As indicated, the case $½ < 1=2$ is analogous. ■

References

Alós-Ferrer, C. (1999): "Dynamical systems with a continuum of randomly matched agents", Journal of Economic Theory, forthcoming.

Bergin, J. and B. Lipman (1996): "Evolution with state-dependent mutations", Econometrica 64, 943-56.

Binmore, K. and L. Samuelson (1997): "Muddling through: noisy equilibrium selection", Journal of Economic Theory 74, 235-65.

Boyd, R and P. J. Richerson (1985): Culture and the Evolutionary Process, Chicago: University of Chicago Press.

Cole, H.L., G.J. Mailath and A. Postlewaite (1992): "Social norms, saving behavior, and growth", Journal of Political Economy 100, 1092-1125.

Cole, H.L., G.M. Mailath and A. Postlewaite (1998): "Class systems and the enforcement of social norms", Journal of Public Economics 70, 5-36.

Corneo, G. and O. Jeanne (1998): “Social organization, status, and economic growth”, Journal of Public Economics 70, 37-52.

Ellison, G. (1993): “Learning, local interaction, and coordination”, Econometrica 61, 1047–73.

Ely, J. (1995): “Local conventions”, mimeo, University of California at Berkeley.

Feldman, M. and C. Gilles (1985): “An expository note on individual risk without aggregate uncertainty”, Journal of Economic Theory 35, 26-32.

Fershtman, C., K. Murphy and Y. Weiss (1996): “Social status, education, and growth”, Journal of Political Economy 104, 108-32.

Foster, D. and P. Young (1990): “Stochastic evolutionary game dynamics,” Theoretical Population Biology 38, 219-32.

Fudenberg, D. and C. Harris (1992): “Evolutionary dynamics with aggregate shocks”, Journal of Economic Theory 57, 420-41.

Freidlin, M.I. and A.D. Wentzell (1984): Random Perturbations of Dynamical Systems, New York: Springer Verlag.

Harrington, J.E. (1995): “The social selection of flexible and rigid agents”, Department of Economics, The John Hopkins University, forthcoming in American Economic Review.

Harsanyi, J.C. and R. Selten (1988): A General Theory of Equilibrium Selection, Cambridge, Mass., MIT Press.

Judd, K. L. (1985): “The Law of Large Numbers with a continuum of IID random variables”, Journal of Economic Theory 35, 19-25.

Kandori, M., G.J. Mailath and R. Rob (1993): \Learning, mutations, and long-run equilibria in games”, Econometrica 61, 29-56.

Kremer, M. (1993): “The O-ring theory of economic development”, Quarterly Journal of Economics 108, 551-75.

Mailath, G., L. Samuelson and A. Shaked (1995): “Evolution and endogenous interactions”, Social Systems Research Institute, working paper no. 9426, University of Wisconsin at Madison.

Meyn, S.P. and R.L. Tweedie (1993): Markov Chains and Stochastic Stability, London: Springer Verlag.

Robson, A. and F. Vega-Redondo (1996): ”Efficient Equilibrium Selection in Evolutionary Games with Random Matching”, Journal of Economic Theory 70, 65-92.

Samuelson, L. (1994): “Stochastic Stability in Games with Alternative Best Replies”, Journal of Economic Theory 64, 35-65.

Vega-Redondo, F. (1997): “Unfolding social hierarchies in large population games”, A Discusión working paper, WP-AD 97-23, University of Alicante.

Young, P. (1993): \The evolution of conventions”, Econometrica 61, 57-84.