

BONN ECON DISCUSSION PAPERS

Discussion Paper 15/2003

Traits, Imitation, and Evolutionary Dynamics

by

Wendelin Schnedler

July 2003



Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post  World Net
MAIL EXPRESS LOGISTICS FINANCE

Traits, Imitation, and Evolutionary Dynamics

Wendelin Schnedler*

July 31, 2003

Abstract

In this article, a modelling framework for the information transmission between agents in an evolutionary game setting is proposed. Agents observe traits which reflect past and present behaviour and success of other agents. If agents imitate more successful agents based on these traits, the resulting dynamics are a multivariate stochastic process. An example for such a process is simulated. The results resemble the replicator dynamics to a remarkable degree. If traits moderately depend on the past, this accelerates convergence of the dynamics towards a stable state. If the dependence is strong, the stable state is not reached.

Keywords: replicator dynamics, imitation, evolution of cooperation, information transmission, simulation

JEL-Codes: C73, D83

*W.Schnedler@bristol.ac.uk, Leverhulme Centre for Market and Public Organisations (CMPO), University of Bristol, Bristol BS8 1TN, England

The author likes to thank Wolfgang Leininger, Erwin Amann, Ingolf Dittmann, and René Fahr for helpful comments. All remaining errors are my own. Financial support by DFG and the Leverhulme Trust are gratefully acknowledged.

1 Introduction

In the context of evolutionary game theory, it is often assumed that agents in a population interact in randomly drawn pairs and afterwards revise their strategy using the gained information (see for example Fudenberg and Maskin 1990). A particular form of revision is imitation. Imitation is usually driven by the intention to increase one's success. Hoping that present and future performance are positively related, an agent might be interested in learning from another agent who is currently more successful. Supposing that the two agents meet, the learning process can be structured in the following way:

- The agent compares her success with that of the other agent to determine to what degree she should change her behaviour.
- The agent then attempts to assess the behaviour of her role model in order to copy it.

Both elements of the learning process require knowledge. First, the agent must know her and the other player's "success". Second, she must know her and the other player's "behaviour". Typically, success is taken to be the utility obtained in the immediately preceding interaction while behaviour is taken to be the strategy used in that interaction (see Weibull 1995, pages 155ff or Schlag 1998). Schlag (1998) proves that adopting the strategy with a probability proportional to the difference in utility is the optimal learning rule when there is no additional information available. Schlag (1998) also points out that a population following this rule evolves according to a pay-off monotone dynamic. Weibull (1995, p. 158) notes that in the neighbourhood of a stationary state, this dynamic resembles the replicator dynamics, which are usually employed to describe the purely biological evolution of genetically determined behaviour. This surprising similarity of simple social learning (imitation) and biological evolution has first been described by Björnerstedt (1993), who analysed an imitation rule where agents randomly adopt other agent's behaviour and where unsatisfied agents do so more often.

In all these models, imitation is based solely on recent information: the current action and payoff of the role model. Why don't agents also use past information? The underlying rationale seems to be that agents meet anonymously. But if two strangers meet, they will have difficulty identifying the success and behaviour of the other perfectly. Often, the assessment then has to rely on signals which partly depend on the past and which imperfectly reflect the present. Is this difference important and does it influence the resulting evolutionary dynamics?

Success or behaviour of an unknown person are usually judged by clothing, manners, body language etc. Frank (1987) argues that some of these signals are beyond the control of the individual (“hard-wired”) to explain the evolution of cooperation in a population. This suggests that behaviour is assessed on the basis of character traits which are to a certain degree robust to changes or difficult to change. Hence, a particular source of influence (and possibly disturbance) on the observable success and behaviour are past utilities and strategies: A person with a scar from a fight might be regarded as dangerous although she turned her back on crime, while a driver of a prestigious car is assumed to be successful, although she might not be anymore.

Contrary to this line of thought, the past normally leaves no trace in evolutionary models of cooperation (Frank 1987, Harrington 1989, Amann and Yang 1998): Once adapting a new strategy, agents immediately signal this strategy. In a genetical context, where agents are programmed to play certain strategies, this assumption seems reasonable. In a social context, where strategies are propagated by imitation, it is less clear why focussing on the present is legitimate. If the evolution of cooperation is seen as a social rather than a genetic phenomenon, it is thus sensible to examine what happens to evolutionary dynamics, if agents are judged and imitated on the basis of traits which do incorporate past events.

This article serves two purposes. On a conceptual level, it suggests a model of information transmission, describes how this information is used for imitation and interaction, and how the resulting dynamics can be simulated. On an applied level, it examines the influence of the past on the evolution of cooperation. It turns out, that some influence of the past on signals has a beneficial effect in the sense that the population reaches an asymptotically stable state with some cooperation more smoothly. If there is too much influence of the past, however, the information used when imitating is out of touch with reality and the population never reaches this state.

The following section suggests a way to model the dependence of signals on the past using the idea of *traits*. Section 3 applies this information transmission mechanism to a particular game which explains the evolution of cooperation. Section 4 ties in the information transmission mechanism into the imitation procedure while section 5 explains how the mechanism affects the interaction between agents. Section 6 considers the dynamics which are induced by imitation and interaction based on past-dependent signals. As these dynamics are hard to describe in a closed form, the problem is adapted for simulation in section 7. Finally, simulations are carried out for the pre-

viously introduced game in section 8 Section 9 concludes.

2 Dependence on the past

We want to assume, that the experiences of agent i , can be described by events e_t^i , where t denotes the time, when the event occurred. Each event consists of the strategy employed s_t^i and the utility gained by this strategy u_t^i : $e_t^i = (u_t^i, s_t^i)$. These events carve themselves into the appearance of the agent and form traits θ_t^i . Lacking additional information, an agent who wants to learn from agent i has to rely on signals which are based on these traits.

In comparison with the complete history of events, traits are less informative. First, it is less clear what the agent actually did. There is a *loss of modal information*: upon meeting a person with a scar, we might not know, whether it was caused by a fight or by an accident. Second, the time dimension might be blurred. In other words, there is a *loss of temporal information*: for example, it might be impossible to say, when exactly the scar was inflicted.

To represent the loss of modal information, we assume a function, which maps events e_t^i into a possibly lower dimensional trait space. This function is called *trait function*. Supposing that strategies can be expressed as k -dimensional vectors, that utility can be measured on the real line, and that the trait space is l -dimensional the trait function g formally becomes:

$$\begin{aligned} g : \quad \mathbb{R}^{k+1} &\rightarrow \mathbb{R}^l, & k+1 \leq l \\ (u_t^i, s_t^i) &\mapsto g(u_t^i, s_t^i). \end{aligned} \quad (1)$$

The loss of temporal information will be taken into consideration by weighting the transformed events with respect to time. Formally, we employ an *intensity measure* $\mu_{t_0}^t(\cdot)$, which has the same properties as a probability measure and assigns a weight to any time point between the starting time t_0 and the present time t . Using this measure, we define the *traits* of an individual i at t as:

$$\begin{aligned} \theta^i : \quad [t_0, t] \times \mathbb{R}^l &\rightarrow \mathbb{R}^l \\ \theta_t^i &= \int g((u_{\tilde{t}}^i, s_{\tilde{t}}^i)) d\mu_{t_0}^t(\tilde{t}) \end{aligned} \quad (2)$$

Finally, the information must be decoded to be useful for the receiver. This step is necessary because the receiver is ultimately interested in utility and strategy of the agent and not in the traits. To represent the decoding, we

introduce a *signal function*. The signal function g^- maps back from the l -dimensional trait space to the $k + 1$ -dimensional space of strategies and utility:

$$\begin{aligned} g^- : \mathbb{R}^l &\rightarrow \mathbb{R}^{k+1} \\ \theta_t^i &\mapsto g^-(\theta_t^i) \end{aligned} \quad (3)$$

Figure 1 depicts, how traits are composed and decomposed, when time is discrete.




Figure 1: The Composition of Traits

Traits carry the information about utility and strategy but this information is blurred by the past.

To simplify the model, one could assume, that the present trait is composed from the last trait and the present event as depicted in figure 2. We want to call such traits *updatable*. If the new trait is a convex combination of the old trait and the present event, the resulting intensity measure for continuous time becomes the exponential density.

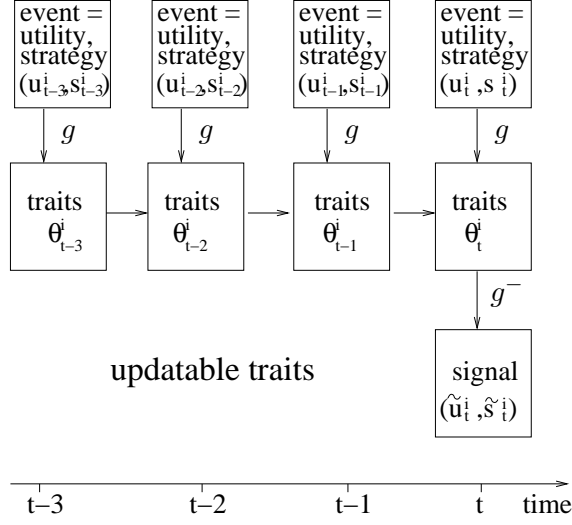


Figure 2: Updatable Traits

If traits are updatable, they are composed of the last trait and the current event.

3 An Example: The Amann-Yang-Game

Information transmission between agents plays an important role for a series of models trying to explain the evolution of cooperation (Frank 1984, Harrington 1987, Robson 1990, Amann and Yang 1998). While most of these models are more concerned about stability and instability of cooperation in populations, Amann and Yang (1998) describe an actual dynamic process leading to some cooperation in a variation of the prisoners' dilemma. This process is fueled by the (rather mechanical) replicator dynamics, where strategies grow proportional to their success. If the evolution of cooperation is seen as a social phenomenon, it makes sense to replace the replicator dynamics by imitation. The result can then be compared to the result of the replicator dynamics.

There are two players and four pure strategies in the game of Amann and Yang:

- cooperate always (C),
- defect always (D),
- cooperate upon meeting a partner exhibiting a "cooperative trait", oth-

erwise avoid the interaction (CA), and

- defect upon meeting a partner exhibiting a "cooperative trait", otherwise avoid the interaction (DA).

Obtaining the information about the trait, which is necessary for strategies CA and DA , leads to costs κ for the agent. If the interaction is avoided, both participants get a side-payment which is larger than the payoff when both agents defect but smaller than the payoff when both agents cooperate. Besides this, the payoff matrix for an interaction is identical to that of the ordinary prisoners' dilemma.

To represent the strategies of agent i , we use the following vectors: $C = (1, 0, 0, 0)'$, $D = (0, 1, 0, 0)'$, $CA = (0, 0, 1, 0)'$, and $DA = (0, 0, 0, 1)'$. Together, with the utility obtained from the interaction, the events experienced by agent i are represented as $e_t^i = (u_t^i, (s_t^i)')'$, where s_t^i is one of the above vectors.

Amann and Yang assume that there is no influence of the past: Signals are identical to traits which in turn reflect the strategy and utility from the interaction. Using the framework of the previous section, these assumptions can be embedded by setting trait and signal function to be identity functions and assigning all weight to the present.¹

Bearing in mind, that real life signals often do not only rely on the present but also on the past and that a loss of modal information occurs, the following alternative might be more appropriate. First, we want to assume, that traits can only reflect, whether an individual intended cooperation or not and whether it was careful or not. We will call the former cooperative trait (\hat{c}_t^i) and the latter risk trait (\hat{r}_t^i). Given a third trait measuring the utility (\hat{u}_t^i), the trait function maps from the five dimensions describing the event to the three dimensions of traits. Formally, the trait function can be represented by the matrix

$$g = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \end{pmatrix}, \quad (4)$$

If this matrix is multiplied from the left to the event vector, we obtain the three-dimensional contribution of the event at time t to the traits, where the

¹Amann and Yang also analyse a variation of their game which includes random noise in the signalling process. However, this noise is limited to the transmission of signals in interactions and does not concern imitation.

first component relates to the utility, the second to carefulness, and the third to cooperation. The actual trait $(\hat{u}_t^i, \hat{r}_t^i, \hat{c}_t^i)'$, is then obtained by weighing these contributions according to the time of occurrence. In case of updatable traits in discrete time, the present contribution is weighed by a factor μ while the last trait is assigned the weight $(1 - \mu)$.

How does an agent infer the utility and strategy from traits? Because the trait space is of lower dimensionality, an assumption needs to be made to recover strategies from traits. For example, assuming independence between carefulness and cooperation allows to retrieve the carefulness from the risk trait and the cooperation aspect from the cooperative trait, while using the utility trait as an indication for utility. The respective signal function transforming traits at time t into signals is:

$$\begin{pmatrix} \hat{u}^i \\ \hat{s}^i \end{pmatrix} = g^-((\hat{u}_t^i, \hat{r}_t^i, \hat{c}_t^i)') = \begin{pmatrix} \hat{u}^i(t) \\ \hat{r}^i(t) \cdot \hat{c}^i(t) \\ \hat{r}^i(t) \cdot (1 - \hat{c}^i(t)) \\ (1 - \hat{r}^i(t)) \cdot \hat{c}^i(t) \\ (1 - \hat{r}^i(t)) \cdot (1 - \hat{c}^i(t)) \end{pmatrix}. \quad (5)$$

This signal function concludes the description of a more sophisticated but complicated alternative to the transmission of information proposed by Amann and Yang.

4 Imitation and traits

Having defined a model of information transmission alone, does not lead to population dynamics. It is also necessary to specify how agents use the information. We want to employ an imitation procedure that is structured in the way we described in the introduction. The success – now measured by the utility signal – is used by an agent i to assess the value of her current strategy. We suppose that the value v^{ij} , which agent i assigns to her strategy upon meeting agent j , is monotonous increasing in the difference between j 's utility signal and i 's utility $u^i - \tilde{u}^j$. Further, we assume the value to lie in the interval $[0; 1]$, where the maximum is reached for all $u^i \geq \tilde{u}^j$ and the minimum will be attained for $u^i - \tilde{u}^j$ having the smallest possible value.

If mixed strategies are allowed, the new strategy of i can be gained by blending i 's current strategy s_t^i and j 's strategy signal \tilde{s}_t^j using the value v^{ij} :

$$s_{t+\epsilon}^i = s_t^i \cdot v_t^{ij} + \tilde{s}_t^j \cdot (1 - v_t^{ij}) \quad (6)$$

If only pure strategies are allowed, the value v_t^{ij} can be taken to indicate the probability of i keeping her strategy. Once an agent has decided to adapt to the strategy of the other, this strategy is identified using the signal vector. Each component of this signal vector can be regarded as the probability of the role model using the respective strategy.

For pure strategies the imitation procedure is identical to that of Weibull (1995, p. 158). However, the information framework here is richer: It encompasses the case of Weibull as a special case where the signal is identical to the trait and all weight is put on the present.

5 Interaction and traits

Up to now, we have developed a model to describe the usage of information within the imitation process. It seems plausible, that the information coded in traits is not only available to learning agents but also to interacting agents.

In principle, all information that is available while imitating should also be accessible while interacting. In other words, actions of the agent may condition on this information. If the agent conditions on additional information, the strategy space is enlarged. This has implications for imitation: The new strategies must be identifiable, otherwise they cannot be imitated. That means, there have to exist signals which allow the agent to distinguish between this and other strategies. These signals, however, can again be used to condition actions upon them. So, the strategy space is once more enlarged. Such recursions can only be avoided when the information for imitation is “richer” than that for interacting. Formally, the information partition must be finer while imitating. This is true for the Amann-Yang-Game considered earlier: While agents condition on cooperation and risk signal when imitating, they only condition on the cooperation signal when interacting. Had they also conditioned on the risk signal, this would have enlarged the strategy space by four strategies, which again would have needed new signals, leading to additional strategies.

Allowing actions to condition on signals, has another consequence. It implies that the utility of an agent does not only depend on the choice of the strategy but also on her traits. This seems rather realistic: Past behaviour does not only influence what other people learn from us but also how they behave towards us.

6 Induced Dynamics

The framework introduced to model the influence of the past on imitation is rather general. This generality comes at a price: The model is hardly analytically tractable. Consider the simple setting of updatable traits. In order to describe the evolution of strategies, it is not sufficient to keep track of the strategies themselves; the prevailing traits also have an effect as they determine the imitation behaviour. So, a state of the population is characterised by the joint distribution of traits and strategies, while the transition from one time point to the next time point is stochastic.

Stochastic transitions are sometimes approximated by deterministic dynamics (see e.g. Benaim and Weibull forthcoming). In fact, many deterministic dynamics such as the replicator dynamics can be understood as a deterministic approximation of stochastic phenomena. This generally works because populations are assumed to be large, so that the law of the large number applies and the distribution of strategies is adequately summarised by the shares of strategies within the population. It then suffices to analyse the deterministic process governing these shares.

As soon as traits play some role, a representation of the common distribution by the shares falls short from recognising important information. In particular, the correlation between strategy and utility traits is neglected but crucial for the next state of the population. It is thus impossible, to get rid of the stochastic element by just focusing on shares. Because the process cannot be simply represented by the dynamics of shares, finding an analytical solution is difficult and will not be pursued, here. Alternatively, we obtain an idea how the population evolves by carrying out a simulation for a specific game.

7 Adapting the problem for simulation

Before we can simulate dynamics, we have to deal with the problems arising from the "correct" representation of the model in a computer programme.

First, we have to overcome the limitation of computers to create continuously timed events. We cannot simulate agents who are able to change their strategy at any time but only at discrete time points. However, we can imagine the agents to change their behaviour in between two of those time points and think of the time points as discrete measurements to which the changes are assigned. If we suppose poisson distributed changes and a constant expected

rate for such changes to occur, reducing the time between two discrete time points is equivalent to decreasing the number of changes for each time unit. Consequently, we can approximate continuously timed revisions to any accuracy by using a sufficiently small share of agents revising their strategy between two measurements.

Second, we cannot evaluate the integral in formula (2) because time events are discrete. We solve this problem by assigning weights to the intervals between two measurements instead of assigning them to single time points, where the weights are taken to be the area below the graph of the density function. The construction of the weights ensures, that their sum equals one. Each transformed event is then weighted with the weight of the preceding interval. The result is an approximation of the integral by a sum, which again can be made arbitrarily precise by increasing the measurements for each time unit.

Third, there is always a low probability that each strategy gets totally extinct by chance. This is a consequence of the fact, that digital computers can only store approximations for real numbers since they have to represent numbers by finite states. If a strategy share falls below the lowest representable real number above zero, it simply vanishes. The strategy will never return, even if conditions for this strategy are optimal. To keep all strategies present, we introduce spontaneous changes of behaviour which are not caused by imitation but by mistakes or experiments of the agents. For pure strategies, we assume that each agent adopts one of the strategies with a certain probability. For mixed strategies, we will not allow any component of the strategy vector to drop below this probability. The probability should be sufficiently low not to exert too much influence on the induced dynamics but sufficiently high to assure the presence of the strategies. This last requirement is similar to the idea of Levine and Pesendorfer (2000) that the imitation mechanism (and not experimentation) is the driving propagation mechanism.

8 Dynamics in the Amann-Yang-Game

To examine whether a population driven by imitation develops similar to the replicator dynamics and to analyse the influence of the past on the evolution of cooperation, a simulation is carried out for the Amann-Yang-Game.

First, we summarise some of Amann and Yang's (1998) results: When the price for information κ is not too high the Amann-Yang-Game has two Nash-

equilibria. One consists of a mixture between the strategies D and DA and should be referred to as the "defective equilibrium". This equilibrium is not locally stable. The other "cooperative" equilibrium is a mixture of the strategies C , CA , and D and it is an asymptotically stable fix point. For specific starting values, the shares of the strategies in the population approach this equilibrium fluctuatingly (see figure 3), while they exhibit a particular pattern: Careful cooperators (CA) successfully invade a population of defectors. After a certain level of cooperation is reached, being careful is not necessary anymore and agents become careless (C). This in turn gives defectors the opportunity to spread.




Figure 3: Replicator Dynamics

The replicator dynamics exhibit a cyclical pattern, where careful cooperation is followed by cooperation which entails defection.

Is the defective equilibrium still locally unstable when agents imitate more successful agents on the basis of traits? To answer this question, the initial state of the simulation will be the defective equilibrium. Assuming that the defective equilibrium has been stable for a long time, we take all strategy traits to represent the true strategies and all utility traits to equal the equilibrium payoff. The value function from the imitation procedure will always be taken to increase linearly in the difference between utility and utility signal (see section 4). Trait and signal function will always be chosen as described

in formulas (4) and (5). The intensity measure μ will vary from simulation to simulation.

The simplest choice for the intensity function is to put all weight on the present. Doing so our model differs from the original model by Amann and Yang just with respect to the learning dynamic: Where Amann and Yang supposed the replicator dynamic, we will use the proposed imitation dynamic. Our choice yields the dynamics depicted in figure 4. This dynamics resemble the deterministic replicator dynamics to a remarkable degree – compare with figure 3. Amann and Yang prove the instability of the defective equilibrium with respect to small changes in strategy shares but they do not use this equilibrium as a starting point for the dynamics. Accordingly, they do not describe the complete transition from a defective to a partly cooperative population. That such a transition is possible can be seen from figure 4. The finding also suggests that starting condition do not matter for cooperation to arise.




Figure 4: Imitation Dynamics without Past

The social dynamics of imitation exhibit a pattern similar to the biologically motivated replicator dynamics. For legend see figure 5.

The similarity between the findings of Amann and Yang and the simulation

indicates that the social behaviour of imitation and the evolutionary concept of selection are closely related. While Weibull (1995) has proven this for the neighbourhood of asymptotically stable states, it seems to hold more generally: the pattern induced by imitation resembles the replicator dynamics already when shares are still far from the equilibrium levels. This adds justification for using the replicator dynamics to model social phenomena in general and in the case of Amann and Yang in particular.




Figure 5: Imitation Dynamics with Past

Moderate influence of past events on traits increases the similarity between imitation dynamics and replicator dynamics.

Up to now, the only innovation was the use of imitation instead of the replicator dynamics. Next, we introduce the influence of the past. If we select the exponential distribution as our intensity function and allow for a moderate influence of past strategies and utilities on traits, the resulting process resembles the replicator dynamics even more (see figure 5). The convergence towards the asymptotic equilibrium is accelerated: It takes fewer cycles until the shares are close to the cooperative equilibrium. This acceleration is due to better informed agents: When the past has no influence, the decision to imitate is based on two interactions, the interaction of the imitator and that of the role model. However, when traits incorporate the past, they capture the experience of a multitude of interactions –for example previous interactions of imitator and role model with other agents as well as the interactions

which entered previous learning. The multitude of interactions allows a more stable prediction of the success of strategies in the current population. Statistically speaking, the sample on which agents base their decision is larger and they are better in estimating the state of the population when traits are moderately influenced by the past rather than without such influence.




Figure 6: Imitation Dynamics with Strong Influence of the Past
With a strong influence of past events on traits, agents adapt to an outdated situation and the strategy shares overshoot. For legend see figure 5.

However, the influence of the past needs to be "moderate" for this observation to be true. If we increase the influence of the past (by changing the parameter of the exponential distribution), the information about previous interactions which is captured by traits is outdated. As a consequence, agents adapt a population which is not existing anymore –see figure 6. The dynamics become sluggish because agents continue imitating outdated strategies while actually optimal strategies are driven to extinction. Accordingly, the dynamics oscillate between the extremes. Without experimentation, no oscillation would occur and the dynamics would come to a halt the first time, when one strategy takes over the population –in the depicted case this would be careful cooperation.

There is a clear trade-off with respect to the influence of the past: The greater the influence of the past on traits, the more information is used

when imitating; at the same time, the information gets older and thus less useful. So, imitation of successful agents leads to strategies closest to the best response for some but not too much influence of the past on traits.

9 Concluding remarks

Assuming that the attempt of an agent to imitate the present behaviour of a role model is systematically disturbed by the past of this role model, we defined a model for the informational flow between agents. Additionally, we formulated an imitation procedure based on this informational flow. This framework is rich but it is difficult to describe the resulting dynamics analytically. Hence, we used simulation techniques and a particular game to get an idea about how the influence of the past affects the distribution of strategies in a population. More specifically, we analysed how imitation using different types of available information affects the evolution of cooperation.

When the past has no influence, the dynamics resemble the replicator dynamics. This enforces the idea that the replicator dynamics, which originally referred to biological phenomena, are a suitable deterministic approximation to imitation, which describes a social behaviour. Moreover, a moderate influence of the past stabilises the dynamics and accelerates the convergence to a stable state, so that the resulting process is even more similar to the replicator dynamics. The intuition for this effect, is the following. The situation in the population is sufficiently stable so that past experiences are valuable to judge the current state of the population. If, however, the influence of the past on the signals becomes stronger, signals get less and less informative about the current population and agents start to adapt to the past. This might even lead to a situation where the distribution of strategies in the population never stabilises.

The theoretical concept of information transmission could also be applied to other games to see whether the results are robust. Eventually, this may lead to the identification of an analytical result which links the influence of the past, imitation, and evolutionary dynamics.

References

AMANN, E., AND C.-L. YANG (1998): “Sophistication and the Persistence of Cooperation,” *Journal of Economic Behavior and Organization*, pp.

91–105.

- BENAIM, M., AND J. WEIBULL (Forthcoming): “Deterministic Approximation of Dynamic Evolution in Stochastic Games,” *Econometrica*.
- BJÖRNERSTEDT, J. (1993): “Experimentation, Imitation and Evolutionary Dynamics,” Mimeo. Department of Economics. Stockholm University.
- FRANK, R. (1987): “If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?,” *American Economic Review*, 77, 593–604.
- FUDENBERG, D., AND E. MASKIN (1990): “Evolution and Cooperation in Noisy Repeated Games,” *The American Economic Review*, 80(2), 274–279.
- HARRINGTON, J. E. (1989): “If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?: Comment,” *American Economic Review*, 79(3), 588–593.
- LEVINE, D. K., AND W. PESENDORFER (2000): “Evolution Through Imitation in a Single Population,” Levine’s Working Paper Archive 2122, UCLA Department of Economics.
- ROBSON, A. J. (1990): “Efficiency in Evolutionary Games: Darwin, Nash and the Secret Handshake,” *Journal of Theoretical Biology*, 144, 379–396.
- SCHLAG, K. (1998): “Why Imitate, and If So, How?,” *Journal of Economic Theory*, 78, 130–156.
- WEIBULL, J. (1995): *Evolutionary Game Theory*. The MIT Press, Cambridge, Massachusetts.