

An evolutionary model of competition and imitation

How much should one imitate others as one progresses in learning and competition?

Hakim Khalafi^{a,1},

Supervised by Prof. Jerker Denrell and Prof. Chengwei Liu^b

^a*Centre for Complexity Science, University of Warwick*

^b*Warwick Business School, University of Warwick*

Abstract

The role of imitation in human learning is thought to be one of the key factors of our success as a species. Prior research has focused on answering central questions such as what circumstances makes imitation beneficial, and who one should imitate. We aim to examine the question of what degree of imitation is appropriate as one advances from beginner to expert in competition and learning. To study this we have designed an evolutionary agent based model with individuals that imitate and compete with each other. The model suggests that beginners should imitate and learn from a mixed group of above average individuals. Experts however should aim to imitate others as little as possible, and when necessary only imitate other experts.

Keywords:

agent based model, evolutionary dynamics, imitation models, social learning

Introduction

The remarkable success of the human species is often attributed to our capacity for learning and accumulating knowledge [1]. An important aspect of learning is social learning - learning through observing, imitating and interacting with others [2]. Much progress has been made in understanding these concepts but a complete understanding of social learning and imitation is yet to be achieved [3].

Prior work focuses on questions such as why copying is beneficial [3, 4] and who to copy [5, 6]. These issues span many academic fields such as the social sciences [7], economics [8] and biology [9].

However, few models address how individuals change their imitative behaviours as they *progress* in competition. The few papers that do are empirical studies of imitation in children and infants [10, 11]. While valuable, these studies do not aim to examine the question of how imitation relates to mastery of a skill and competition in general. This is our aim.

Email address: H.Khalafi@warwick.ac.uk (Hakim Khalafi)

¹The author to whom correspondence should be addressed.

To address this we have developed an agent based model that simulates competition in a changing environment. Strategies of agents are accompanied by a notion of progress that tracks how long their strategies have been successful.

We present this new model with properly defined imitative and competitive dynamics. This is done in order to study our questions of what level of imitation is appropriate at different stages in learning and competing. Another question we're looking to answer is how selective individuals should be of who they imitate, as they progress.

Method

Chosen method

The model is an evolutionary agent based model with several interacting components. The techniques obtained from the literature are standard agent based modeling techniques and the following analysis of results is fairly standard as well.

The reason we chose this approach is because it fits our purposes well.

A closely related mathematical framework that is sometimes used to model imitation and learning phenomena is evolutionary game theory (EGT). For the readers familiar with these methods it is worth to note that the underlying mechanisms in our model differ from those of EGT.

We decided to not use this framework for two reasons. The first reason is that we do not wish to use pairwise contests to determine payoffs as this is not realistic for what we are trying to model. The second reason is that it is not immediately obvious how to define a notion of progress within this framework.

Prior to constructing our model, research was done to make sure assumptions we have made are reasonable. These assumptions are explained throughout the next section.

Our model is completely based in MATLAB.

Model

In brief, there is a population of individuals that evolve and adapt to their environment. After going through a payoff process each individual performs mutation and imitation behaviours in order to adapt. Each process is explained in detail below and a summary can be found in table 2 at the end of this section.

We will from here use the term 'individual' instead of agent in the model.

Environment

We begin by assuming that there is an inherent environment optimum value that individuals want to achieve.

This environment optimum value is denoted $O(t)$ and varies with time t . For reasons of realism and simplicity we model this with a one dimensional Gaussian random walk. A realisation of this random walk can be seen in fig. 1. The random walk has step sizes that are zero-mean normally distributed $\mathcal{N}(0, \sigma_I^2)$.

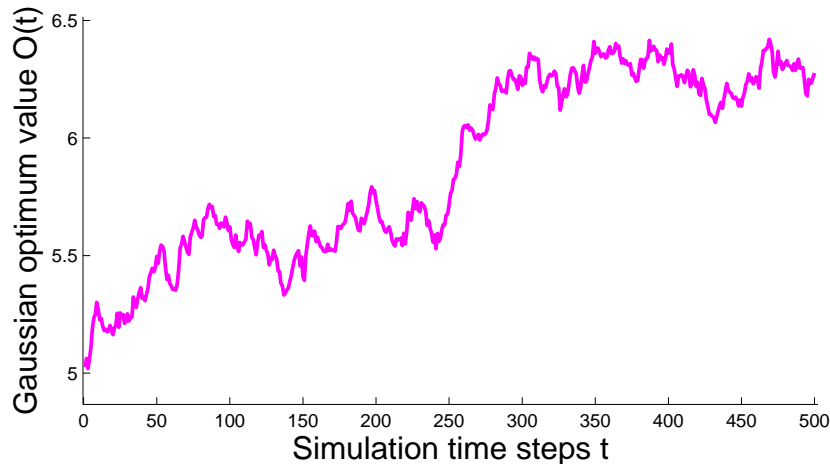


Figure 1: The background environment optimum is assumed to follow a Gaussian random walk. Plotted is a typical Gaussian random walk. The random walk has step sizes that are distributed as $\mathcal{N}(0, \sigma_I^2)$.

The standard deviation σ_I that is defined by us for each simulation can be used as a measure of environment instability. A random walk that has step sizes that vary with a larger σ_I is said to be more unstable.

There are n individuals in the simulation that adapt to this optimum. They are entirely defined by their strategies. Each individual has four strategies that determine their inherent behaviour. We can therefore construct a n -by-4 matrix that keeps track of these values. A typical realisation of this matrix at a certain time in the simulation can be seen in table 1.

Individual	Estimation \mathcal{E}	Mutation \mathcal{M}	Imitation \mathcal{I}	Selectivity \mathcal{S}
Individual 1	4.76	0.1	0.1	0.1
Individual 2	4.8	0.12	0.2	0.2
...
Individual n	4.6	0.07	0.2	0.1

Table 1: A typical strategy matrix of individuals in the simulation here shown as a table. This is for a certain time t . The matrix values will vary as the individuals evolve and environment changes. As an example, for this point in time the optimum could be 4.7 with most estimations being around that value.

The first of these strategies, estimation, is each individual’s estimation of the optimum value in the next time-step. Estimation is used in the evolutionary process to determine the fitness

of each individual. The closer an individual's current estimation $\mathcal{E}_{i,t}$ is to the next optimum value $O(t + 1)$ of the system, the more fit he is.

Given that the optimum value varies stochastically, it is not possible for individuals to estimate what the next optimum value will be exactly. However it is possible for individuals to evolve so that their other strategies give them the *highest probability* of estimating the right value.

The underlying assumption is that an individual in a competitive situation wants to be as good as he possibly can, but that he cannot predict exactly what his surrounding environment expects of him in the future. The strategy estimation is a simplification of these dynamics.

The rest of the strategy types will be explained in depth after the evolutionary mechanisms have been introduced.

Payoffs

We start the simulation at time $t = 0$ by initializing the strategy matrix as a random matrix with each value standard normally distributed.

After this, each individual receives a payoff proportional to how close their estimation is to the environment optimum in the next time step.

More specifically, the payoff of individual i at time $t + 1$ is denoted $P_{i,t+1}$. It is calculated by inserting the individual's current estimation $\mathcal{E}_{i,t}$ in to the probability density function (PDF) of a normal distribution. The mean of this normal distribution is the optimum of the next time step (eq. 1).

$$P_{i,t+1} = f(\mathcal{E}_{i,t}; O(t + 1), \sigma_{EP}^2) \tag{1}$$

This way, estimations that are close to the optimum will be close to the peak of the PDF of the normal distribution. Individuals with these near-optimum estimations will therefore receive a relatively large payoff. While estimations that are far from the optimum are further away from this peak where corresponding PDF value is low. These individuals will therefore receive a smaller payoff.

The underlying normal distribution $\mathcal{N}(O(t + 1), \sigma_{EP}^2)$ will be referred to as the evolutionary distribution. A visualization of this payoff process can be seen in fig. 2.

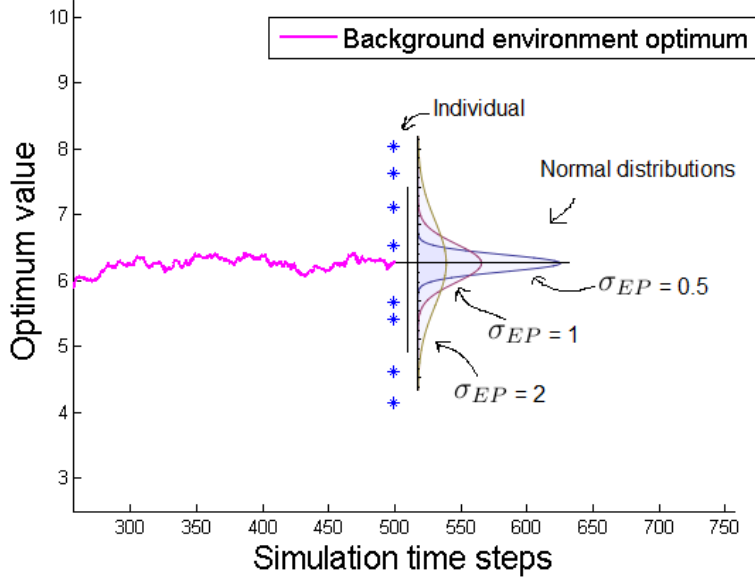


Figure 2: The payoff mechanism. The blue stars are individual estimations and the purple line is the background environment optimum. Depicted also are a couple normal distributions with different σ_{EP} . The individuals receiving highest payoff are the ones closest to the optimum of the next time step. Note however that choosing a larger σ_{EP} means individuals further away from the optimum receive more payoff.

To this distribution belongs a standard deviation σ_{EP} . Increasing the σ_{EP} of the evolutionary distribution (and thereby changing the shape of it) corresponds to giving relatively more payoff to outliers and relatively less to individuals close to the optimum.

This payoff mechanism was chosen because of its favorable properties. Namely, varying σ_{EP} can be thought of as varying the *evolutionary pressure* in the system. A small σ_{EP} means high pressure as few optimum centered individuals then receive a high payoff, whereas a high σ_{EP} means more accepting dynamics in which most individuals in the system receive a payoff, albeit a small one.

Another property of this mechanism is that it acts on the whole population at once. That is, each individual is at any given time competing against the whole population of individuals.

Next we describe how payoffs are translated to survival and replication of individuals in the system.

Replication

First the payoffs are rescaled so that they sum to n , the amount of initial individuals. We then round each payoff number to nearest integer.

The resulting payoff numbers will be the amount of times an individual replicates.

There are therefore n 'slots' in the next time step which individuals in the current time step compete for. The individuals that receive the highest payoff get the most slots, while individuals that receive the least payoffs get no slots at all.

Thus the amount of individuals in the simulation that are allowed to exist at any one time is n . But the amount of *unique* individuals that are fit enough and chosen for survival, varies.

As an example consider a particular time t where 30 individuals were chosen out of $n = 200$ to be fit enough according to this mechanism. Then each of these 30 individuals will reproduce proportional to their payoff to fill up the 200 slots at time $t + 1$. The individuals with highest payoff get the biggest proportions.

These dynamics lead to each individual being completely replaced by its replicas in the next time step. Therefore in order to measure how successful strategies belonging to a lineage of individuals have been in their time evolution it is necessary to maintain a notion of progress.

Progress and beginners

Progress tracks how many time steps a lineage of individuals has been able to survive in the simulation. At the beginning of the simulation the progress of every individual is set to zero. As the individuals undergo a payoff round and then replicate, the progress of each replica is incremented by one. An individual with progress of 600 time steps for example has had a lineage of ancestors that have survived for 600 time steps in the simulation.

We also add agents of progress zero ('beginners'), that have standard normally distributed strategy values, in each time step. We do this for two reasons. The first reason is that without these beginners a possible simulation outcome is that no individual is fit enough for reproduction. This leads to a wipe out of the population. Adding beginners restarts the simulation dynamics. The second reason is that we are interested to see if these beginners can start competing with old established strategies or even outcompete them.

Finally with these evolutionary dynamics in mind, we describe the remaining behaviours that occur for each individual.

Selectivity, Imitation and Mutation

After replication has occurred, each new individual imitates a selected portion of the population by replacing part of his values with the average of this selected portion. After this the individual mutates each of his strategies. An in depth explanation of these behaviours follows.

Selectivity $\mathcal{S}_{i,t}$ determines how large part of the population the individual i will imitate at time t . Via a mapping this value will range between zero and one.

A selectivity of zero means that the individual with highest payoff will be imitated, 0.3 means that the average of the top 30% of the population with highest payoffs will be imitated and one means that the average of the complete population (100%) will be imitated. We denote this average for the different strategies estimation, grade of mutation, grade of imitation and selectivity with $A_{\mathcal{E}}$, $A_{\mathcal{M}}$, $A_{\mathcal{I}}$ and $A_{\mathcal{S}}$ respectively.

The rationale behind this is that each individual will have the possibility to be selective about who he chooses to copy as is usual in imitation models [12]. Selectivity also goes by the name success-bias in the literature and means being biased towards only successful individuals in the selection of who to copy.

The extent of which the individual imitates is decided by his grade of imitation.

The **Grade of Imitation** $\mathcal{I}_{i,t}$ determines how much an individual includes strategies of others with his own. The strategies of others is the average of a selected portion of the population according to the selectivity mechanism above. Grade of imitation is also mapped to a value between zero and one. The imitation step updates the individuals strategy values as follows

$$\begin{aligned}\mathcal{E}_{i,t} &= (1 - \mathcal{I}_{i,t})\mathcal{E}_{i,t} + \mathcal{I}_{i,t}A_{\mathcal{E}} \\ \mathcal{M}_{i,t} &= (1 - \mathcal{I}_{i,t})\mathcal{M}_{i,t} + \mathcal{I}_{i,t}A_{\mathcal{M}} \\ \mathcal{I}_{i,t} &= (1 - \mathcal{I}_{i,t})\mathcal{I}_{i,t} + \mathcal{I}_{i,t}A_{\mathcal{I}} \\ \mathcal{S}_{i,t} &= (1 - \mathcal{I}_{i,t})\mathcal{S}_{i,t} + \mathcal{I}_{i,t}A_{\mathcal{S}}\end{aligned}$$

As an illustrative example, consider a particular individual that has a selectivity of 0.3 and grade of imitation of 0.4. Each of his strategy values will be updated to a weighted average between his own values and the average of the top 30% of the population. The weights are according to the grade of imitation, so 60% of his strategies will be his own values and 40% the values of others.

If the individual has estimated the optimum environment value to be 5, and the average of the top 30% of the population estimates it to 5.3, the individuals new estimation will be $(1 - 0.4) * 5 + 0.4 * 5.3 = 3 + 2.12 = 5.12$.

The **Grade of Mutation** $\mathcal{M}_{i,t}$ determines how much individual i will vary his strategy values at time t . This is done by adding a random number r to each strategy value of agent i . The number r is distributed as $\mathcal{N}(0, \mathcal{M}_{i,t})$. This means that the following updates of each strategy value occurs

$$\begin{aligned}\mathcal{E}_{i,t+1} &= \mathcal{E}_{i,t} + r_1 \\ \mathcal{M}_{i,t+1} &= \mathcal{M}_{i,t} + r_2 \\ \mathcal{I}_{i,t+1} &= \mathcal{I}_{i,t} + r_3 \\ \mathcal{S}_{i,t+1} &= \mathcal{S}_{i,t} + r_4\end{aligned}$$

This is done in order to introduce an element of randomness in the model that allows the individuals to evolve.

The four strategies and other key concepts are summarized with the following table for quick reference.

Concept	Symbol	Description
Estimation	\mathcal{E}_i	The estimation of the next optimum value by individual i .
Grade of Mutation	\mathcal{M}_i	The amount individual i varies his strategy values.
Grade of Imitation	\mathcal{I}_i	The extent individual i copies the strategies of others.
Selectivity	\mathcal{S}_i	How selective individual i is when choosing group to imitate.
Progress	-	Number of time steps a lineage of individuals has survived.

Table 2: A summary of concepts relevant for the results section

Simulation set-up

We were interested in

1. The prevalence of imitation as an individual progressed in competition.

We measured this by calculating the average grade of imitation for every individual of every progress level for all time steps in the simulation. When this had been established we examined

2. How selective individuals were of whom they imitate as they progressed in competition.

This was measured by calculating the average selectivity for every individual of every progress level for all time steps. The result of this sparked an interest to measure

3. What the prevalence of imitation would be, if every individual always was forced to imitate the absolute best.

The absolute best here is the individual with highest payoff.

Our main result is the first item in this list. The second and third result are answers to questions raised by the first result.

The following parameters were kept constant throughout all simulations and results below. They were chosen to reflect reasonable and realistic dynamics.

Parameter	Symbol	Value	Reason/Explanation
Evolutionary pressure	σ_{EP}	0.05	Simulates high evolutionary pressure environment.
Instability of optimum	σ_I	0.03	Simulates an environment that slowly changes.
Amount of individuals	n	200	Amount of individuals in the simulation.
Amount of beginners	–	50	Amount of beginners added per time step.
Total time steps	–	10 000	Every simulation ran for this many time steps.
Amount of simulations	-	100	Every plot is the average of 100 simulations

Each mean value for the progress figures below represent the cumulative grade of imitation for all individuals at a specific progress level, divided by the amount of those individuals.

The evolutionary dynamics of the system ensure that the measured statistics only come from successful individuals. This is because individuals with disadvantageous strategies were weeded out by evolutionary mechanisms.

Results

Figure 3 shows that it becomes less beneficial for an individual to imitate as he progresses and becomes more experienced in his competitive environment.

More precisely the close up plot to the right shows a spike in the amount of imitation for absolute beginners.

For intermediates there is a linear decline in imitation. For the top few with the most progress the most beneficial level of imitation is zero. That is no imitation at all.

As a clarification of what is going on consider that the maximum progress possible is 10 000, and happens if any individuals were good enough from time $t = 0$ to survive until the end of the simulation. As this rarely happens however the top individuals have only progressed to about 7000 time steps.

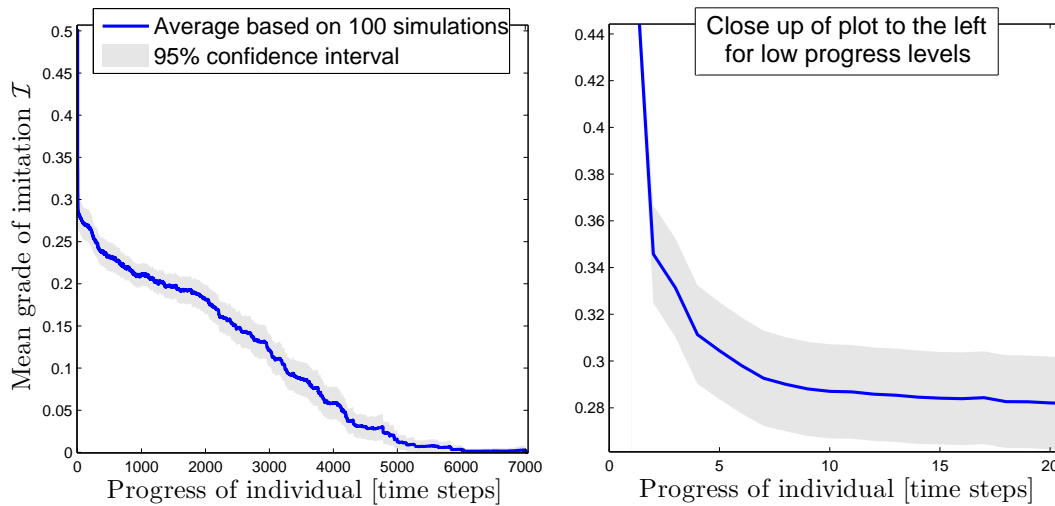


Figure 3: The mean grade of imitation as dependent on the progress structure of the population. This was calculated by measuring the grade of imitation for every individual of a certain progress for each simulation. The plot is an average of this for 100 simulations. The figure shows a steady decline in grade of imitation for individuals of increasing progress. A close up for absolute beginners can be seen to the right. Each simulation lasted 10 000 time steps.

A question that followed is how selective are individuals in their imitation behaviour?

Figure 4 shows that individuals become more and more selective of who they imitate as they progress in competition. Also here is a linear decline, with beginners imitating the top 30% of the population and becoming increasingly selective until the absolute top individuals only benefit from imitating other top individuals.

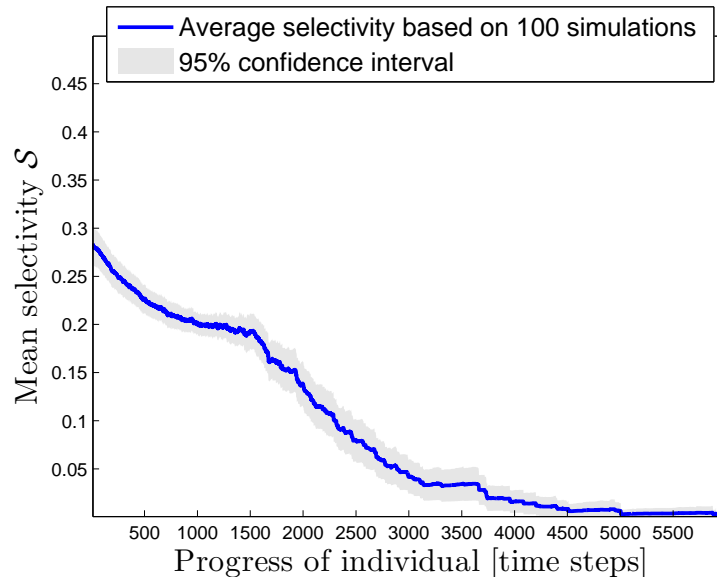


Figure 4: The mean selectivity as dependent on the progress structure of the population. This was calculated by measuring the selectivity for every individual of certain progress for each simulation. The plot is an average of this for 100 simulations. The figure shows that as individuals progress in competition, they become more selective of who they copy. Each simulation lasted 10 000 time steps.

Given these results we asked ourselves, why don't individuals *always* copy the top ranked individual? Results so far indicate that beginners tend to imitate a mix of the top 30%, what would be the harm in restricting this to the absolute top?

To find this out we fixed all imitation behaviour for all individuals to always copy the top ranked individual.

Figure 5 shows a comparison in imitation levels when restricting copying to the top individual vs. letting individuals be freely selective (fig. 3).

We can see a markedly lower imitation prevalence compared to when individuals were free to choose what proportion of the population they imitate. This suggests that as a beginner it might be harmful to strictly imitate the top individual, instead one should aim to imitate a mix of the top individuals.

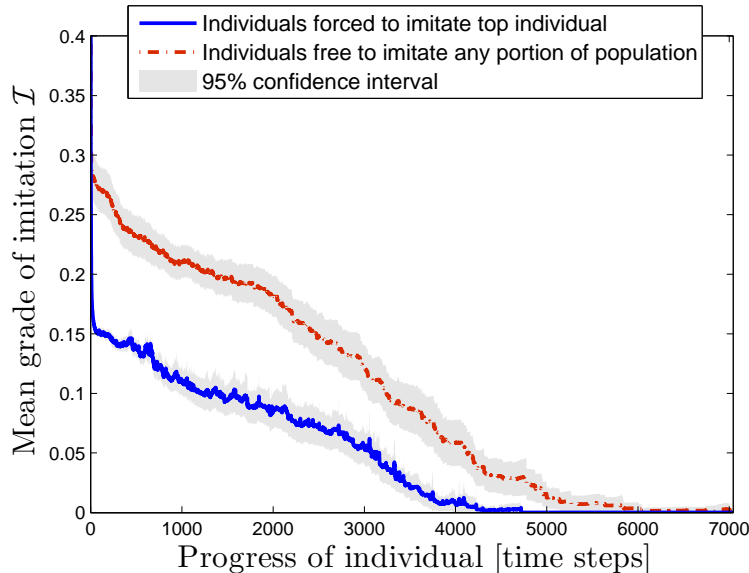


Figure 5: The blue line shows imitation levels for individuals only copying the top individual. This was done by forcing the selectivity parameter for every individual to zero. The plot is an average for 100 simulations. The figure shows less prevalence of imitation compared to when individuals were not forced to copy the top individual, shown in red. Each simulation lasted 10 000 time steps.

Discussions

Our main finding is that individuals imitate less as they progress in competition. As these results are for individuals that were evolutionarily successful, it implies that these are the actions of an individual that wants to maximize their probability of success. The results suggest that individuals at the top of their field should not aim to imitate others. The results also suggest that complete beginners should imitate the most at the start to get going and progressively lessen imitation as they become more established.

Our other findings are that beginners should imitate or learn from a group of individuals that perform above average. Moreover individuals at the top of their field should only imitate other top individuals. The fact that individuals imitate selectively at all is widely supported in the literature [13]. Our finding that a complete beginner learns the most from a group of above average individuals goes against the common idea that one should imitate only the absolute top performers. Recent studies on this phenomenon [6] indeed show that using the "copy only the top" heuristic can be risky.

In fact, to investigate this we forced the "copy only the top" heuristic in individuals and the result was system wide lower overall imitation prevalence. And surely, imitating the top performer is not the most beneficial action to take when extreme performance indicates system unreliability [6]; we believe that the top individual in our models became the top simply by being lucky enough at the start of simulations with his estimations. Subsequently self-reinforcing dynamics allowed this individual to remain at the top. The lower prevalence of imitation reflects the fact that imitation became less beneficial for individuals when implementing this restriction.

We have considered an alternative explanation of our results. Namely, it is possible that the reason the absolute top individuals have a lower prevalence of imitation is because they are the individuals that were around from the start where every other individual had almost stochastic strategies. Imitating these near stochastic individuals would be catastrophic if an individual was lucky enough to achieve a good estimation of the optimum. However the fact that "top level" is different for each simulation is enough for us to believe that this effect for any particular age level becomes washed out over time.

The main limitation of this model is the fact that it relies on many assumptions. Although great care has been taken to keep these assumptions sensible, they could combined lead to faulty and unrealistic dynamics. We remain optimistic however as many results are confirmed by established literature (see previous citations). Relating this model to real life data could result in a fine tuning and polishing of the model, an endeavor discussed below.

Conclusion and future outlook

The main conclusion from our model is that an individual should imitate the most when they're an absolute beginner at something, and that the individual should seek to learn from and imitate a mixed group of above average individuals. An individual at the top of their field however should imitate at all, and if needed only imitate other experts.

Future research that can be immediately applied is studying other aspects of our model. For instance we can vary our evolutionary pressure parameter σ_{EP} or the instability parameter σ_I to see what effect it has on the system and the population of individuals.

However, in the light of the limitations of our model, future research should be focused on relating the model to real life data for fine tuning and verification. We can think of two further ways future research can be carried out in this regard.

First, data may already exist in the form of research from previous studies. For instance data from behavioural economics could be used to infer how good people are at predicting stock market values that are inherently stochastic [14]. This could be related to the individuals in our model trying to estimate an optimum value. Other sciences that may be able to provide useful data on selectivity aspects are the biological and psychological sciences [15, 16].

However for a model as particular as ours we suggest that a new study should be carried out.

So secondly, we could design an online game that offers rewards for participation. The survey could be designed to resemble our model where people playing the game represent individuals. Users will be asked to make a prediction of the future value of say an imaginary company's stock, given only its current value. The closer their predictions are, the more they are rewarded. People taking the survey will also have the chance to see what the predictions of other people are, and how successful these people have been previously at estimating stock values. We can then collect data and analyse it in a similar way we did in this paper.

We let these three options for future research mark the end of this paper.

Acknowledgements

I would like to thank Jerker Denrell and Chengwei Liu for extensive discussions and comments on early drafts. This work was supported by the Erasmus Mundus programme.

References

- [1] L. L. Cavalli-Sforza, M. W. Feldman, Cultural transmission and evolution: a quantitative approach, Vol. 16, Princeton University Press, 1981.
- [2] C. M. Heyes, Social learning in animals: categories and mechanisms, *Biological Reviews* 69 (2) (1994) 207–231.
- [3] L. Rendell, et al., Why copy others? insights from the social learning strategies tournament, *Science* 328 (2010) 209–213.
- [4] R. Boyd, P. J. Richerson, Why does culture increase human adaptability?, *Ethology and sociobiology* 16 (2) (1995) 125–143.
- [5] K. N. Laland, Social learning strategies, *Animal Learning & Behavior* 32 (1) (2004) 4–14.
- [6] J. Denrell, C. Liu, Top performers are not the most impressive when extreme performance indicates unreliability, *Proceedings of the National Academy of Sciences* 109 (24) (2012) 9331–9336.
- [7] É. Danchin, L.-A. Giraldeau, T. J. Valone, R. H. Wagner, Public information: from nosy neighbors to cultural evolution, *Science* 305 (5683) (2004) 487–491.
- [8] J. Apesteguia, S. Huck, J. oechssler (2007). imitation-theory and experimental evidence, *Journal of Economic Theory* 136 217–235.
- [9] K. Dautenhahn, C. L. Nehaniv, *Imitation in animals and artifacts*, MIT Press, 2002.
- [10] A. N. Meltzoff, The human infant as imitative generalist: A zo-year progress report on infant imitation, *Social learning in animals: The roots of culture* (1996) 347.
- [11] F. Kaplan, P.-Y. Oudeyer, The progress-drive hypothesis: an interpretation of early imitation, *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions* (2007) 361–377.
- [12] B. Galef, K. N. Laland, Social learning in animals: Empirical studies and theoretical models, *BioScience* 55 (2005) 489–499.
- [13] L.-A. Giraldeau, T. J. Valone, J. J. Templeton, Potential disadvantages of using socially acquired information, *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 357 (1427) (2002) 1559–1566.
- [14] K. E. Wärneryd, *Stock-market psychology: How people value and trade stocks*, Edward Elgar Pub, 2001.

- [15] A. Meltzoff, R. Williamson, Imitation and modeling, in: E. in Chief: Marshall M. Haith, J. B. Benson (Eds.), *Encyclopedia of Infant and Early Childhood Development*, Academic Press, San Diego, 2008, pp. 127 – 137.
- [16] A. N. Meltzoff, M. K. Moore, Imitation, memory, and the representation of persons, *Infant Behavior and Development* 17 (1) (1994) 83 – 99.