

Imitation and Social Learning

Karl H. Schlag

Department of Economics

Vienna

Austria

karl.schlag@univie.ac.at

May 7, 2011

Next version forthcoming in Encyclopaedia of the Sciences of Learning, Springer Verlag

Synonyms

Copying; acquiring an understanding as a group

Definition

Imitation is the act of copying the behavior of someone observed. It is the most common learning rule, and as behavior, can be observed among animals^[GL] as well as among humans^{[AHO],[HW]}. Not to imitate means to experiment or innovate, to try something new instead of choosing one of the behaviors that has been observed.

To learn means to acquire knowledge or understanding. Social learning is about learning within a social group as opposed to individual learning where only the learning of a single person is considered. Concern is for how individuals within the group learn, and hence how the whole group learns. Social learning does not require that individuals within the group actively change their understanding about the environment. To exhibit learning as a group it is already sufficient that more within the group adapt a superior action.

Theoretical Background

Imitation is a simple behavior that has two basic ingredients. One needs to be able to observe what others have done and one needs to be capable of doing what they have done. Often one is also able to observe the consequences of their choices. However, one does not need to understand why someone observed has made this specific choice nor understand how this choice has generated the observed performance.

To imitate means to economize on explicit or implicit costs of evaluating each possible choice. Its virtues are clearest when others have superior information. Not much theory is needed to explain why imitation is valuable under such circumstances. If someone else knows what is best and faces the same situation, then it is best to (unconditionally) imitate this person.

In this essay we present a theory that demonstrates why imitation can be valuable even if no individual has superior information. We show how imitation can enable the group to aggregate the different experiences of its members. Again, there are not many new insights if individuals face an environment where there is no variability in the performance of an action, when the same action always yields the same payoff. In such an artificial setting, imitation of someone who has performed better improves one's own performance. But what if there is variability in the performance of an action? This is the more natural setting. In the following we allow the performance of an action to depend on the current state of the environment, so that an action that sometimes performs well may perform badly on other occasions? What is the value of imitation for the person and for those in the group that are observing each other? It is the value of imitation in this context of

social learning that we wish to identify and to qualify. The original arguments can be found in [SWI], for a non-technical treatment see [AFS].

We will be considering a group of individuals who repeatedly make choices in the same environment. Between choices each meets someone else within the group and observes the action and performance of the other. We wish to answer two questions: Why imitate, and if so, then when? The “why” is simple. It is in order to aggregate individual experiences, based on an understanding that the social group together has more information than each separate member. In the case of effective social learning behavior, an action that is used by more individuals will reflect that this action is on average more successful. Hence, information on the success of each action will implicitly be stored by the frequencies with which they are used in the group. The question *when* to imitate seems obvious if the performance of others is observable. It seems most plausible to imitate those that have performed better, a rule we call Imitate if Better. It is theory that helps us to understand why this is generally not the right thing to do. It is theory that helps us identify what aspects of imitative behavior we should look for, and in what types of environments imitation is more likely to be valuable. Not until we have an understanding of which rules are best can we make predictions about social learning with an understanding that results are not driven by an ad-hoc choice of behavior but instead by behavior that is in some sense optimal.

Important Scientific Research and Open Questions

Imitate if Better is the classic behavior (or rule) that is associated with imitation if only one other is observed. We start by explaining why this behavior need not have good properties. To imitate those that perform better can induce, on average, more to choose an inferior action. In fact, if there are many in the group then it can be very likely that all choose an inferior action in the long run. The reason is that those that perform better need not be choosing a better action when performance itself is variable (or noisy). One needs to see an example to completely understand this counterintuitive result. Assume that all individuals are using Imitate if Better. Assume that there are two actions D and N. Action D is deterministic in the sense that it always yields a payoff equal to 0.1. Action N is noisy and yields 40% of the time the value 1 and 60% of the time the value 0. In terms of average performance, N is the superior action. The average payoff of N is $0.4*1+0.6*0=0.4$ which is larger than the (average) payoff of D which is equal to 0.1. Now consider two individuals seeing each other, one using D and the other using N. 60% of the time, so in 60% of such matches, N yields 0, and hence the individual using N switches to D. 40% of the time the individual using N gets payoff 1 and hence the individual using D switches to N. Thus, the likelihood of a switch from N to D is 60% while that of a switch from D to N is 40%. It is more likely that there will be a switch from N to D than vice versa. In the next round one expects more to play D. The number of those choosing the superior action N is expected to decrease. If there are many individuals, with initially both D and N being chosen, then it is easy to prove that with a large probability no one will choose the superior action N in the long run.

Is this environment too complex for simple learning rules to perform well? No, Imitate if Better is simply not the appropriate rule to use. In fact, on closer look, it is not that intuitive. While imitation of those that perform better sounds appealing when one's own payoff is 0.1 and observed performance is 1, it is much less appealing when one's own payoff is 0 and observed payoff is 0.1. The difference in performance should play a role, and it does, both behaviorally^[AHO] and theoretically. Consider the following alternative imitation rule called the Proportional Imitation Rule^[SWI]. Never imitate anyone who performed worse. Do not always imitate those that do better, but instead choose the probability of imitating the observed proportional to how much better the observed performed. We illustrate with the above example. This means to choose a constant $\sigma > 0$ that is common for all in the group and then to switch when using D and observing N with payoff 1 with probability $(1-0.1)*\sigma=0.9*\sigma$ and when using N, receiving payoff 0, and observing D to switch with

probability 0.1σ . The increase in choice of N among those pairs where D observed N is then equal to $0.4\sigma - 0.6\sigma = -0.2\sigma$. Hence more are expected to choose the superior action N in the next round.

So why not choose the superior action N at the start? Typically, information about the underlying process is limited, i.e. we do not know that N is superior. The decision theoretic solution to learning under limited information, Bayesian learning, is to first (subjectively) assess the likelihood of each possible environment and then to choose what is best, given the information one gathers over time. This is not very practical as it does not tell us how to determine these likelihoods. Moreover, updating choice based on information is extremely difficult even in the simplest of situations. Instead of searching for a rule that is best according to such a specific subjective assessment we search for a rule that is good in all environments and hence does not require such a subjective assessment. We search for one rule that on average causes performance in the population to increase over time regardless of how payoffs are generated. This captures learning in its most basic notion as it only relies on minimal a priori knowledge, in contrast to decision theoretic (or Bayesian) learning which is more about updating initial assessments (or beliefs).

So, what kind of rules are good in all environments when all use the same rule? They have to be imitative: either choose previous action again or switch to the action of the observed. Otherwise, if some individuals go off trying new actions not observed then these may be inferior to those currently used in the population and hence can reduce overall payoffs in the population. Imitate to avoid making mistakes. Use a rule that is responsive to the magnitudes of performance. For instance, Imitate if Better did not perform well. In fact, it turns out that only a small class of imitation rules does the job. Among these one can argue that the Proportional Imitation Rule (PIR) is best. Of course, when all use the PIR, then some may do worse in the next round, however others will do better in a way that average performance in the population is expected to increase over time. PIR incorporates differences in performance in behavior to wash out good luck and bad luck and to focus on average performance. From a broader perspective we find that society can function as a storage or memory of the success of strategies, where the success of a strategy is mirrored by the proportion of individuals choosing it. Strategies can only remain popular over time if they are superior. For individuals within this population we find a particular form of random behavior to be valuable as switching to those that did better is probabilistic. To act probabilistic means to incorporate the magnitude of performance into one's own behavior. When all use PIR, then the frequencies of choices made within the population evolve according to the replicator dynamic (known from evolutionary game theory). This connects social learning dynamics with asexual reproduction dynamics.^[W]

Once we have derived the simple but insightful result above, we can discuss the role of assumptions that our logic is based on.

In what types of environments is the above true? The environment has to be stationary in the sense that each action generates either a constant or a noisy payoff where the underlying process generating payoffs does not change over time. Individuals must be identical in the sense that the payoff generated by an action may not depend on which individual chooses this action. In particular, payoffs should not depend on what others choose in the group. For instance, when payoffs are generated by some interaction then one observes others within the group and interacts with others outside the group where each individual in the group potentially interacts with the same outside the group. Payoff differences must be bounded as switching occurs with probability equal to σ times the difference in payoffs, where σ is the same for all environments, and this probability has to be bounded above by 1.

When is Imitate if Better a good rule? It is best among imitation rules when the performance of actions is not subject to noise. It is also a good rule^[SIB] when noise is idiosyncratic, which means that the payoffs generated by any two actions only differ according to their mean and not according to other properties of their

distribution such as variance (as assumed in the location shift model of statistics and when referring to homoscedastic errors in econometrics). Here payoff differences do not have to be bounded.

Above, we assumed that performance of others is observable. What should one do when one does not observe their performance and one only observes their behavior? Use the proportional reviewing rule that is defined as follows. Draw a threshold, each threshold being equally likely, and switch if own performance lies below this threshold. It is as if one is using a satisficing rule (cf. [S]) where the aspiration level above which one is satisfied is drawn according to a specific distribution.

Up to now we have assumed that each individual only observes one other individual between making choices. What should one do if the individual observes the behavior and the performance of two or more others? It is hard to compare rules. However, one rule that performs well is the sequential proportional observation rule.^[HS] This rule prescribes to place the observed individuals in a random order, to then consider each individual one by one, to imitate the action of the next in line with probability proportional to the observed payoff and then to continue to the next. Population dynamic approaches the adjusted replicator dynamic when many are observed, a further connection between social learning and evolutionary game theory.^[W] Note that this behavior is based on the proportional observation rule, to imitate the observed with probability proportional to the performance of the observed, empirically found to be relevant for understanding the behavior of fish.^[P]

Above, we observed that Imitate if Better is a good rule when noise is idiosyncratic and one other is observed. But what if noise is idiosyncratic and many others can be observed? One might think that one should imitate the action that performed best on average. However, this rule does not induce good social learning properties. Averaging washes out noise which is good for understanding the value of a choice used by the majority. However, this is unfavorable when compared to a choice only used by few when noise is not symmetric. An inferior choice that sometimes generates very bad outcomes can spread when used by few if it is likely to generate outcomes above that of the majority choice. Concern for average performance is only a good rule if one is comparing groups where each choice is used by many. In order to allow for comparisons of many with few (and many with many) one should imitate the action of the individual that performed best.^[SIB]

What is the role of innovation? If all use an imitative rule then there will never be any innovation, no new action will ever be introduced. Innovators are needed for progress, they have to be willing to experience failures. In contrast, society based on imitative behavior has the ability to avoid failures in the aggregate and yet to adapt actions introduced by innovators and to test them against the existing strategies.

There are numerous open questions in the literature.

What if there are more prominent members in society, members that are more likely to be observed than others? We provide some intuition for the case where each individual only observes one other. The key to the above result is that there should be no expected change when both actions have the same average performance. So those observing the prominent member should anticipate that many are switching when adapting this person's action while only one switches, namely the prominent member, when the non-prominent perform better. To illustrate, assume that a prominent member is observed twice as likely as a normal member. Then normal members have to downscale their switching propensity from the rate σ used by the prominent member to $\sigma/2$.

What if the probability of observing someone depends on that person's performance? This sampling bias has to be taken into account which in special cases is easy to do. For instance, assume that each individual is matched with another one but only observes that person's choice with probability that is proportional to her performance. Then a good rule is to imitate whenever possible. Here good performance of the rule rests on the assumed linear relationship between observability and performance.

Other important open topics for future research include giving up the assumption that all use the same rule (cf [BS]), investigating how to deal with memory, how to learn from others facing similar but different situations and how to deal with changing environments.

The key to this research is to accept and model limits on information and knowledge, to consider simple objectives and to then gain insights into optimality of behavior that sounds plausible, like imitation. Insights can be valuable across many disciplines including anthropology, population biology, decision theory and economics. Theory shows where to look in the data, eg. see [AHO], [K], [P]. Once one has identified which rule is best one can make predictions. For instance, an investigation of the replicator dynamics^[W] can now be understood as an analysis of the outcome of “optimal” social learning with limited information.

In summary, there is more to imitation than simply copying behavior of others. To be an effective social learning rule, thus to induce the group over time to choose superior actions, each individual has to look at differences in performance before deciding whether to imitate. A very simple rule that determines when to imitate and only relies on the ability to observe behavior of one other, yields effective learning even when there is variation in individual performance. Learning becomes possible as society functions as a storage of what has been successful, more successful choices being adapted by more individuals.

References

[AFS] Alós-Ferrer, C. and K.H. Schlag (2009), Imitation and Social Learning, chapter in *The Handbook of Rational and Social Choice*, P. Anand, P. Pattanaik and C. Puppe (eds.), Oxford University Press.

[AHO] Apesteguia, J., Huck, S. and J. Oechssler (2007), Imitation - theory and experimental evidence, *Journal of Economic Theory* 136, 217 – 235.

[BS] Björnerstedt, J. and K.H. Schlag (1996), On the Evolution of Imitative Behavior, Discussion Paper No. B–378, Sonderforschungsbereich 303, University of Bonn.

[GL] Galef, B.G. Jr. and K.N. Laland (2005), Social learning in animals: empirical studies and theoretical models, *BioScience* 55(6), 489-499.

[HS] Hofbauer, J. and K.H. Schlag (2000), Sophisticated Imitation in Cyclic Games, *Journal of Evolutionary Economics*, 10(5), 523-543.

[HW] Horner, V. and A. Whiten (2005), *Animal Cognition* 8, 164–181

[K] Kendal, J.R., Rendell, L., Pike, T.W., and K.N. Laland (2009), Nine-spined sticklebacks deploy a hill-climbing social learning strategy, *Behavioral Ecology* 20(2), 238-244.

[P] Pike, T.W., Kendal, J.R., Rendell, L.E. and K.N. Laland (2010), Learning by proportional observation in a species of fish, *Behavioral Ecology* 21(3), 570-575.

[SWI] Schlag, K.H. (1998), Why Imitate, and if so, How? A Boundedly Rational Approach to Multi-Armed Bandits, *Journal of Economic Theory* 78(1), 130-156.

[SIB] Schlag, K.H. (1996), Imitate Best ' vs 'Imitate Best Average', Mimeo, University of Bonn, <http://homepage.univie.ac.at/karl.schlag/research/imitation/IDIO9.pdf>.

[S] Simon, H.A. (1955), A Behavioral Model of Rational Choice, *The Quarterly Journal of Economics* 69(1), 99-118.

[W] Weibull, J. (1995), *Evolutionary Game Theory*, Cambridge, MA: The M.I.T. Press.