Imitative and Direct Learning as Interacting Factors in Life History Evolution

Bullinaria, John

DOI:
[10.1162/ARTL_a_00237](https://doi.org/10.1162/ARTL_a_00237)

License:
None: All rights reserved

Document Version
Peer reviewed version

Citation for published version (Harvard):
Bullinaria, J 2017, 'Imitative and Direct Learning as Interacting Factors in Life History Evolution', *Artificial Life*, vol. 23, no. 3, pp. 374-405. https://doi.org/10.1162/ARTL_a_00237

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is an accepted manuscript which will appear in final published form in *Artificial Life*, published by MIT
https://www.mitpressjournals.org/doi/pdf/10.1162/ARTL_a_00237

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Imitative and Direct Learning as Interacting Factors in Life History Evolution

John A. Bullinaria

School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK

j.a.bullinaria@cs.bham.ac.uk

Phone: +44 (0) 121 414 2590

Abstract: The idea that lifetime learning can have a significant effect on life history evolution has recently been explored using a series of artificial life simulations. These involved populations of competing individuals evolving by natural selection to learn to perform well on simplified abstract tasks, with the learning consisting of identifying regularities in their environment. In reality, there is more to learning than that type of direct individual experience, because it often includes a substantial degree of social learning that involves various forms of imitation of what other individuals have learned before them. This paper rectifies that omission by incorporating memes and imitative learning into revised versions of the previous approach. To do this reliably requires formulating and testing a general framework for meme-based simulations which will enable more complete investigations of learning as a factor in any life history evolution scenarios. It does that by simulating imitative information transfer in terms of memes being passed between individuals, and developing a process for merging that information with the (possibly inconsistent) information acquired by direct experience, leading to a consistent overall body of learning. The proposed framework is tested on a range of learning variations and a representative set of life history factors to confirm the robustness of the approach. The simulations presented illustrate the types of interactions and trade-offs that can emerge, and indicate the kinds of species specific models that could be developed with this approach in the future.

Keywords: Artificial life, Life history, Evolution, Learning, Imitation, Memes

1. Introduction

The key details of individual lifetimes clearly exhibit a wide variety across different species, with life history factors (such as growth stages, age at first reproduction, number of offspring, degree of parental investment, age of menopause, lifespan, etc.) interacting and varying in line with the relevant physical and environmental differences. This variation has emerged from evolution by natural selection, and the study of *life history evolution* is a well established field that aims to understand the driving forces underlying the diversity [29, 51, 72, 76, 77, 78]. One factor that is known to be relevant [7, 51, 53, 59], but has so far received relatively little attention, is how learning within lifetimes, and learned information passed across generations, influences life history evolution. The study presented here actually began with the question of why some species are precocial, requiring little parental care, while others are altricial, requiring long term care and protection, and why humans are particularly altricial, even compared to other primates [57]. It is clear why protection may be necessary while growing to become big and strong enough to survive independently, but is it not so clear whether *learning* to perform well enough to survive without protections is a significant factor. The reason why this has not been studied before is probably that it is impossible to investigate these aspects analytically, meaning one must resort to full-scale agent-based artificial life simulations, and setting up such simulations in a realistic yet computationally feasible manner has proved challenging [17]. Agent-based models of aspects of life history evolution do already exist (e.g., [76]), but the setting up of reliable simulations to answer questions about learning and protection is not so straightforward, and actually needs to be constrained by developing a robust framework that is able to model learning in life history evolution much more generally. This paper therefore presents and tests a general simulation approach that will enable informative large scale life history evolution simulations involving learning, and illustrates its power by presenting results from a series of models of the evolution of parental protection.

One of the principal difficulties with such an endeavor is in identifying appropriate simplifications that allow the artificial life simulations to explore particular issues reliably and repeatedly, without all the confounding factors inherent in real biological systems, yet remain realistic enough that useful results and predictions can emerge from them [82]. The standard approach is to concentrate on the aspects of interest, and adopt typical general purpose approximations for all the other details. For example, simulations of cooperative breeding would concentrate on the group structure and environment, and gloss over details such as individual and social learning [76]. This paper systematically develops a general framework for modeling the effects of learning, tests it and generates baseline results using the simplest specifications of the other details, while providing pointers to the numerous additional features and relaxed assumptions that may be necessary for more realistic models of specific species and environments in the future.

Artificial neural network models that learn from representative streams of input-output samples have already provided good accounts of many forms of behaviour in humans and other animals [1]. The literature is full of successful simulations of normal and abnormal learning in infants, various patterns of adult behaviour (including factors such as generalization performance, reaction times, priming, and speed-accuracy tradeoffs), and different types of brain damage and neuropsychological deficits. Such computational models have also formed the basis of artificial life simulations of whole populations of individuals that learn by experience within their world, compete with each other, and evolve by natural

selection. That in turn has led to improved accounts of particular aspects of human development and life history, such as ages at first reproduction and the periods of protection that parents often offer their offspring [17]. Those models have elucidated the trade-off between learning quickly and learning well, and demonstrated how evolution can successfully balance the trade-off by leading to the emergence of extended periods of parental protection during which learning can be completed slowly and effectively, without the interference of performance-based natural selection pressures.

Although that previous study [17] began with artificial neural network based individuals that learned from sets of training patterns, it then moved on to formulate more computationally efficient non-neural network abstractions of that kind of learning process which rendered larger scale evolutionary simulations possible. All those simulations were based on the assumption that the learning was achieved in each individual by direct experience, by identifying important patterns in the training data that represented the world they lived in. In reality, for humans and many other animal species, not all learning is by that kind of direct experience. A significant component arises by mimicry of what other individuals have learned before them, or by instruction from those other individuals [10, 19, 30, 42, 47, 48, 52, 56, 62, 65, 87, 88]. That kind of imitative or social learning is better formulated in terms of *memes* (i.e., elements of culture, behaviour or information) being passed between individuals [2, 6, 12]. Any complete account of the effect of learning on life history evolution will clearly need to involve both types of learning, and that means incorporating imitation and memes into the earlier simulations that only involved direct learning. There will inevitably be costs and benefits to individuals arising from both types of learning, and, as always, there will be trade-offs between them [69, 77, 78]. Even though many of the relevant trade-offs may appear clear from a theoretical point of view, the interactions will be complex, and their consequences are found to depend on the associated simulation or environment details and parameters. The only feasible way to proceed is by performing comprehensive series' of simulations that explore the effect of the various details, establishing the range of behaviors that can emerge, and then fixing the details to match the particular species and scenarios of interest. The need to do that reliably and efficiently is where the above-mentioned general simulation framework comes in. The difficulty is that there are many crucial details that can easily cause problems if they are not dealt with carefully, and there are many potential variations to consider, some of which can lead to big differences in the results, and others that make very little difference. It is hoped that this paper, by specifying a complete simulation framework with these issues highlighted and tested in detail, will enable other researchers to set up their own reproducible simulations much more quickly and reliably.

A groundbreaking study by Higgs [43] has already attempted to simulate the evolution of learning by imitation, but that did not consider the crucial issue of how imitative learning interacts with more traditional neural learning by direct experience. Moreover, it is not even clear what is the best way to bring those two rather different forms of learning together into a single framework. An important aspect of the earlier study of learning in life history evolution [17] was the demonstration of how it is possible to abstract out almost all the details of the direct neural learning from training samples, yet still be left with an evolutionary system that results in the emergence of the same life history traits. Although that abstraction process was primarily aimed at reducing the computation times, and hence allowing more realistic larger-scale simulations, it also provides a basis for representing known information in a way that

facilitates the incorporation of both direct and imitative learning into the same framework. That is what allows this paper to introduce a parameterized account of memes and imitation into the earlier study [17], leading to a general framework for investigating the effect of both forms of learning on the various meme transmission and life history factors. Simplifications are required to provide baseline simulations that can form a robust starting point for more detailed models later, but care is taken to highlight all the crucial assumptions, approximations and design choices as the framework is developed.

The proposed approach is designed to be general enough that it can be validated by existing simple mathematical models, but also go beyond those models by relaxing the simplifications they rely on. For example, Rogers [71] considered how biology constrains culture by looking at two simplified models: one with no genetic evolution, and one with no cultural evolution. The simulation framework of this paper can be tested against those extreme cases, and also explore the continuum in between. It can also incorporate factors crucial for life history evolution simulations, such as evolving age distributions and competition between individuals of different ages, that are virtually impossible to deal with adequately in tractable mathematical models, which usually have to resort to using discrete generations [26].

The remainder of this paper is organized as follows: The next section specifies the underlying life history evolution simulation framework, and the following two sections explain how the direct individual learning and imitative social learning processes can be modeled efficiently within it. Then the details and results of a series of simulations are presented that are designed to test and explore many of the key issues and variations associated with the two forms of learning. As an example of the potential of the general framework, a selection of representative life history evolution simulations are then presented. The paper ends with some discussion, conclusions and suggestions for future work in this area.

2. The Life History Evolution Framework

The agent-based artificial life approach to life history evolution generally involves simulating simplified abstract versions of real biological populations, with the simplifications driven by the questions of interest. In the current context, that means maintaining populations of individuals of various ages, each specified by a relevant set of innate parameters, that can learn from their environment during their lifetimes to improve their performance, and compete with each other for survival and to produce children. At each stage, the fitness of each individual will depend on how well it has managed to learn to perform so far. The precise environmental details, abilities required for good performance, and learned behaviors associated with them, need not be specified, as long as realistic patterns of fitness acquisition can be established. The requirement that individuals experience performance-based competition to survive and reproduce results in populations of increasing learning ability emerging due to evolution by natural selection. Moreover, since individuals of all ages are competing with each other, there is also a pressure to learn to achieve good performance *quickly*, to minimize the period during which they are outperformed by older individuals. The associated trade-off between speed and accuracy of learning can then lead to the evolution of risky learning strategies that do not arise in over-simplified “generational” simulation approaches which have weaker selection pressures and do not match real environments so well [15]. One of the life history factors of particular relevance to this is the period of care or protection from competitors that parents of many species offer their young [20, 55, 57]. There is an inevitable trade-off between the

costs associated with that protection and the costs of adopting quicker-but-poorer learning strategies, which evolution will balance in different ways depending on the details of the population and its environment.

The simulation framework proposed here is inherently modular, so it is straightforward to represent some factors more (or less) accurately if they are especially important (or irrelevant) for models of particular species or scenarios. Though, for reliable testing of the learning details, which are the core contribution of this paper, it helps to keep the non-learning aspects as simple as possible. Consequently, it will initially be assumed that the simulated environment contains a single species and a fixed total food resource that limits the population to a fixed size, and that population size is maintained by replacing the individuals that die by children of the best performing adults. The death rates will then need to be specified for the particular species and environment being modelled to give realistic birth rates and age distributions, and the population size determined as a compromise between computational efficiency and improved realism. Having multiple species with population sizes that vary freely according to their ability to survive and procreate is also possible, and will be required for modeling many realistic scenarios, but that would over-complicate the testing of the crucial learning features. Having a changing environment will also affect the learning strategies that emerge [26, 71], but, to facilitate the development and testing of the learning framework, the environment and the appropriate learned behaviors associated with it will be kept stable over the course of each simulation, leaving all the potential types and rates of variation to be explored in future studies.

In practical terms, the simulations simply need to maintain a set of parameter values for each individual in the population (age, performance, brain size, learning rates, parent identifiers, number of children, protection period, etc.), updating them at each time-step, and over-writing the parameters of any individuals that die with those of the children that replace them. The time-step size needs to be set so the life history of the modeled species is simulated at a sufficient level of granularity; for example, one update per simulated year might be appropriate for humans. Then each simulation run will start from the relevant random initial population, and needs to be executed for more than enough simulated years that the population averages of the various parameters and performances stabilize. Typically, a few very long test runs are performed to establish a suitable simulation length, and then many runs of a sufficient set number of simulated years are carried out so that the variance across the final populations can be established. Unless stated otherwise, all the results presented in this paper are evolving or final population averages over 20 independent runs (using different random number seeds), and the graphs have error bars representing the standard deviations over the 20 population averages.

The selection and reproduction processes will be kept the same across all the simulations in this paper, to aid fair comparisons, though there remains scope for varying these details in the future. Survival selection corresponds to the weakest individuals being more likely to die at any point in their lifetimes. Such competition-based deaths may result in real populations from a direct fight with another individual, or arise indirectly due to being an easier target for predators or being less likely to find enough food to survive. A standard *tournament selection* approach [27] is sufficient to model all such deaths, with an appropriate number of random pairs of individuals competing each simulated year by having their performance compared, and only the winners surviving. For the test simulations in this paper, these

performance comparisons are based purely on what the individuals have learned so far, but they could also involve performance adjustments that depend on factors such as age or how many children they are looking after. Older individuals will also be prone to dying of old age. In the preliminary simulations here, old-age deaths will take place with a constant probability each year after individuals have exceeded their species' natural lifespan, but the framework is general enough to allow more realistic old-age death rates, potentially involving other factors including performance. The required number of children to maintain the population size are each produced from their two parents using the standard evolutionary operators of crossover and mutation, with the parents chosen each simulated year from the eligible individuals, again selected by taking the winners of performance comparisons of randomly chosen pairs of individuals [27]. Crossover corresponds to having each child's innate parameter values (e.g., learning rates, protection period, etc.) chosen randomly from the corresponding ranges spanned by their two parents, and the mutations are random constants added to each inherited parameter to allow a significant chance of its value falling outside the parental range. Particular distributions of litter sizes, delays between offspring and mutations will need to be set in accordance with the species being modeled. In this way, diversity is maintained, but there is a tendency for the best performing individuals to survive and produce children that inherit their good features, and natural selection gradually optimizes the population for their environment. These details capture many of the key aspects of real animal populations, and although they clearly involve a great deal of simplification, they constitute a manageable level of approximation that has proved to be effective in numerous previous studies [14, 15, 16, 17, 18].

The performance of real individuals clearly depends on a whole range of interacting learned skills appropriate for their environment, but an important feature of the simulation approach here is that the details of what exactly is being learned make no difference to the life history evolution results – all that is required is a reasonably realistic mapping between the learning time and the level of performance that has been achieved. The earlier study of direct learning [17] began with a fairly typical neural network approach that approximated the full learning process using standard fully connected multi-layer perceptron neural networks trained on simple classification/categorization tasks using gradient descent. It was primarily concerned with the particular life history issue mentioned above, namely how the neural learning can affect the periods of protection that many species offer their young. Bestowing such protection clearly involves costs for the parents (or alloparents) that will often increase their own chances of dying. The specific costs will depend on the species, but in real populations will typically include the consequences of providing food, shelter, teaching and protection from predators [20, 55]. The precise details again need not be specified, as the whole protection process can be approximated sufficiently well by simply not allowing the children to die in any way that can be prevented by the protectors, such as due to lack of food or shelter, or being killed by predators or other individuals. Implementing that in the simulations is just a matter of not allowing protected children to die as a result of the performance-based competition for survival until they have reached a certain age. In real animals, the protection will rarely be that effective, but assuming this extreme case enables a clear exploration of the relevant issues, leaving the effect of less-than-perfect protection to be studied in the future. In that way, a series of simulations showed that significant learning advantages and better adult performances are achieved if the children have longer periods of protection, but only if the children are not allowed to procreate while protected

[17]. Moreover, by treating the protection period as an innate parameter that evolves alongside the innate neural learning parameters, it was found that extended protection periods did emerge. The crucial factor is that the lack of competition, either to survive or reproduce, means that learning can proceed more slowly and reliably. The associated cost to the child is the reduced period remaining for reproduction, so a compromise protection period evolves which balances the reduced procreation period with the improved learning and associated increased chances of surviving and being chosen as a parent. That trade-off was also shown to depend on other life history factors, such as the average life-span. In that case, instead of the protection period remaining at a fixed duration determined by the learning task complexity, an increased life-span meant a longer potential procreation period, and more of that could be given up in return for improved learned performance throughout adulthood.

The original aim of this study was to establish whether those general results persist for more complete accounts of learning, that include imitative social learning as well as direct individual neural learning. To do that requires a modeling framework which can accommodate both direct and imitative learning, and formulating and testing that has ended up being another important outcome of the study.

3. Modeling the Direct Individual Neural Learning Process

For large-scale life history evolution simulations involving learning to be computationally feasible, it is necessary to model the learning process with something far less time consuming than standard gradient descent-based neural network training. A key achievement of the earlier study [17] was the demonstration that it is possible to approximate complete neural learning processes by a single performance level that varies with age in a simple parameterized fashion, without changing the qualitative patterns of life history evolution that emerge. This idea is taken even further here, by combining into a single measure of performance *all* the numerous abilities an individual needs to learn to perform well in their natural environment, but the approach and key simulation results, outlined in this section, are essentially the same as in that earlier study [17]. Later sections will show how this can be transformed into the new models that include imitative learning.

There are several considerations that narrow down the choice of function of age to represent the direct individual learning process. Although expertise acquisition will obviously be rather species specific, there are known commonalities across species [37], so a simple “general purpose” function would be an appropriate starting point. Also, the performance improvements brought about by learning should be separated from the eventual decline that is usually associated with old age [37]. In fact, much learning in later life is about compensating for physical deterioration, and that may need to appear explicitly in some future models, but it would be an unwanted source of confounding factors that is best avoided for the initial models of this paper. A suitable compromise is to require a simple monotonic function that eventually levels off, much like a smoothed version of the learning curves typically found in neural network models. Thus, it will be assumed that knowledge is acquired at an approximately constant rate, so the performance rises approximately linearly with age, till an appropriate maximum level is reached. The associated time granularity of the simulations is something that needs to be set in line with the lifetimes being modeled, so the performance increases are relatively smooth. In this paper, all population details will be updated once per simulated year, which is reasonable for human-like life-spans. Then, to

avoid potential complications of performance scales that vary during evolution, it will be assumed that the learned knowledge simply ranges from 0 to 100% at each stage of evolution. Finally, to reflect the nature of real environments, a certain degree of stochasticity is required.

These considerations lead to having individual learned performances that increase from 0 to 100% in steps drawn randomly each simulated year from the range $[0, 2\delta]$. That means the learning will be completed in an Expected Learning Time (ELT) of $100/\delta$, and the population mean performances in life history evolution simulations for particular fixed values of the learning rate δ will fall almost linearly with that ELT. If protection periods are allowed to evolve in the simulations, they rise approximately linearly with the ELT, because there is no need for parental protection after perfect performance has been learned, though the linearity fails if the ELT reaches the point at which individuals start dying of old age [17]. Not surprisingly, the best mean performances are achieved when the learning is fastest, with very high learning rates δ and very low ELT, which allows all the individuals to achieve perfect performance before their first round of competition to survive or procreate. Consequently, if the learning rate δ is allowed to evolve, it quickly takes on very high values leading to very low ELT, so that if the protection period is also allowed to evolve, it quickly drops to zero [17]. That is clearly unrealistic, because it is not possible with real neural processing to just keep on increasing the learning rate parameter (e.g., gradient descent step size) and expect the learning time to decrease with it. Eventually, at some point that depends on the given learning task, the increasing step size causes the learning process (e.g., the approximation to true gradient descent) to break down, with a consequent deterioration in learning performance. In that case, an evolutionary process is able to identify optimal values for the learning parameters [15], and having slower learning with longer protection periods consistently emerges to provide a clear advantage [17].

To represent the neural learning processes more realistically, in a way that has faster learning correspond to riskier learning strategies which eventually result in persistent poor performance, the function of age needs to include an additional component that restricts the learning when the learning rate δ becomes too large. That can be achieved most simply by having each individual's learning process terminate at some random age, and hence at some random point in the learned performance range $[0, 100]$, with a probability $\rho\delta$ of termination each simulated year that increases linearly with both their learning rate δ and a parameter ρ that represents the “learning task difficulty”. The result of just evolving the protection period in this case is shown in **Figure 1** for a range of fixed values of δ , with a population size of 200, human-like natural lifespan of 60 years, and four representative values of ρ . The left graph confirms that this simple parameterized formulation of the learning process is sufficient to result in the required clear peaks in mean performance levels when they are plotted against the $ELT = 100/\delta$. As the task difficulty ρ increases, the peak in mean performance is lower, and the expected learning time to reach it is higher, again as required. That is the expected consequence of a higher ρ leading to a lower learning rate δ at which significant deviations from the earlier $\rho = 0$ case arise. The right graph shows that the value of ρ has very little effect on how the evolved protection period varies as a function of the ELT, confirming that it is the learning time rather than the task difficulty that is driving the evolution of the protection period.

If the learning rate δ is also allowed to evolve, along with the protection period, the clear peak in mean performance for each value of ρ will drive any successful evolutionary process to the associated optimal

value of the $ELT = 100/\delta$, and the corresponding non-zero protection period. The ELT and protection period that actually emerge for each of a wide range of values of ρ are plotted in the left graph of Figure 2. As the task difficulty represented by ρ increases, the resulting evolved ELT also increases as required, and the best protection period is always a few years longer than that. That difference between the evolved protection periods and the corresponding $ELTs$ is due to the stochastic nature of the learning process and the fact that the evolutionary operators lead to distributions of learning rates and protection periods, and the obvious evolutionary advantage of having protection periods that are long enough to accommodate a significant number of individuals that are slower than average to learn.

The parameter ρ is an abstract measure of learning difficulty designed to represent the difficulty that real individuals have in acquiring the information required for good performance in their environment. Relatively straightforward performance corresponds to low ρ , can be learned quickly, and has a short associated protection period. More complex skills and performance correspond to higher ρ , take longer to learn, and benefit from extended protection periods. In the earlier study [17], the individual performance levels emerging from this abstracted learning model were compared directly with those arising from full evolutionary neural network simulations involving the learning of non-linear decision boundaries in real-valued input spaces, and a good qualitative correspondence was found when $\rho = 0.04$. The mean evolved ELT in this case is around 10 years and the mean evolved protection period is around 14 years, which tie in reasonably well with human-like time-scales. The right graph of Figure 2 shows how the median performance levels vary with age for individuals in the final populations of these $\rho = 0.04$ simulations, when the protection period and learning rate δ are both evolved (Ev.), and, for comparison, when the learning rate δ is evolved but the protection period is set at a selection of fixed numbers of years (1, 10, 20). This pattern of results is qualitatively the same as for the full neural simulations, despite the stochastic linear learning and uniform distribution of residual errors being rather rough approximations of the real neural learning processes [17]. The emergent protection periods allow slower learning, which enables better adult performance, but excessively long protection periods are avoided because of the associated diminished procreation period. Allowing procreation while protected is an obvious variation, and simulations of that confirm it reintroduces competition and the need to learn quickly, and that destroys the evolutionary advantage of extended protection periods. The consequences of allowing procreation while protected, and of varying the natural life-span, in the abstracted learning models are also found to be qualitatively the same as for the full evolving neural network models [17].

There clearly remains considerable scope for more accurate parameterizations to represent specific direct neural learning processes for different species in real environments, as discussed previously [17], but the general approach described above will need little modification to sit alongside the implementation of imitative social learning that will be introduced in the next section.

4. Modeling the Imitative Social Learning Process

As noted above, the abstract direct learning approach of the previous section was originally designed for computational efficiency, so that more detailed life history evolution simulations would be feasible. However, it also facilitates the development of a modeling framework that allows the amalgamation of information from both direct and imitative learning into the same overall performance measure.

The underlying idea here is that it will often be more efficient to imitate the successful behaviour of other individuals than attempting to learn it from direct personal experience of the environment [52, 62]. This paper will follow the AI tradition of using the terms imitation and social learning in their most general sense, to cover all forms of acquired behavior copying between individuals, though other disciplines tend to use them with more specific meanings [61]. Whiten provides a useful taxonomy of all such processes [87]. There is certainly no dispute that humans have evolved to learn by imitation as well as by direct experience across a wide variety of tasks [63, 67, 70, 87], and there is evidence that imitation also takes place, though less commonly, in other species [6, 19, 24, 30, 47, 84, 88]. It has become standard practice to frame this kind of imitative social learning in terms of the transmission of units of cultural information and behavioral practices between individuals, with those *memes* being allowed to replicate and respond to natural selection pressures in a manner analogous to genes [2, 6, 12, 25]. Recently, there has been considerable interest in these ideas across a range of disciplines [48, 66], and the interesting preliminary work already carried out suggests that artificial life simulations may be the most fruitful avenue for exploring these issues in the context of other life history traits. Belew [4] and Best [5] have already introduced imitation based cultural factors into the much-studied Hinton and Nowlan [44] model of learning guiding evolution, but their approaches are far removed from the neural-inspired learning relevant to the life history factors of interest here. The limitations of those early studies have been addressed by Borenstein and Ruppin [8] in a manner that does incorporate neural learning mechanisms, but they actually prevented cultural evolution by not allowing meme transmission between generations and only allowing innate behaviors to be imitated.

It is the previously mentioned study by Higgs [43] that comes closest to exploring the life history issues of interest here, by investigating the evolution of populations of individuals that may invent and imitate memes. A variety of crucial factors were considered that affect how the meme numbers, imitation rates, and performance levels evolve. The main finding was that significant imitation rates *do* consistently emerge under a range of conditions, even when some memes make a negative contribution to the performance of the individuals possessing them, and there is an inherent cost associated with the ability to imitate. Intuitively, it is not difficult to understand why. If there exists a set of memes with a range of positive and negative contributions to the overall performance, then not imitating them will leave the performance at some baseline, while imitating them will result in a range of performance levels above and below that baseline. Any selection on the basis of performance will then favour those individuals that have imitated the good memes, and hence favour higher imitation rates. Moreover, since it favours individuals that have acquired and can pass on those good memes, the good memes will tend to propagate in the population at the expense of the bad memes. This simple picture will be complicated by memes that act together (i.e., *memeplexes*), and by the interplay of genetic and social contributions to fitness, and also by the interaction of genetic and mimetic replicators [5, 7, 12], but these factors are all things that can be investigated in future simulations within the general framework being proposed here.

For the purposes of the current study, the next important undertaking is to identify an effective way to incorporate the general idea of memes and their imitation into the earlier investigation of learning as a factor in life history evolution that only involved direct individual learning from the environment [17]. That is the problem to be addressed in the next section.

5. Formulating an Imitative Learning Framework

Since extracting reliable general conclusions from artificial life simulations relies on avoiding potential confounding factors, that usually means keeping the models much simpler than when aiming to model particular real life scenarios. Also, for the simulations to remain relevant to a diverse range of species and environments, and to facilitate comparisons between them, the models need to be parameterized (e.g., in the way the task difficulty parameter ρ was introduced above). The aim now is to develop such a parameterized framework for learning and evolution that is flexible enough to cover social learning from other individuals in the most general sense (that includes all processes leading to some form of behavior imitation), as well as direct individual learning from the environment.

The most relevant previous approach is the one developed by Higgs [43], but many of its details are not compatible with the current aims. The main problem is that it does not consider any of the processes taking place during individual lifetimes, which are clearly of utmost importance when modeling life history evolution. Moreover, it is based on non-overlapping generations, which means a total absence of the competition between individuals of different ages that underlies many of the issues of interest here. There are also several other details that introduce unnecessary complications into the analysis, such as Gaussian distributions for the meme contributions to performance, a non-linear relation between the learning ability and probability of imitation, and unbounded numbers of memes that can be invented. So, rather than following the approach of Higgs [43], the abstract direct learning approach described above will be extended in a minimal computationally efficient manner to include the key concepts of memes and their spread by unspecified forms of imitative social learning.

Consider, for example, a practical context for early humans, such as spear throwing or animal trapping. Each aspect of it will have an associated set of “memes”, each corresponding to a particular pattern of behavior, and each individual will have at most one meme from each such set, namely the behavior they use for that aspect of that context. Some of these memes will tend to be better than others, in the sense of providing higher contributions to performance. Harmful or bad memes might correspond to actions that tend to kill fellow humans rather than the hunted animals. The memes might be learned by direct individual learning or practice, or they may be learned by copying what other individuals have learned in the past. Copying and natural selection will tend to result in the more useful memes propagating at the expense of the others. However, simulating such a level of detail would be difficult, because keeping track of large sets of memes for each aspect of each context, with associated copying fidelities and such like, will introduce a considerable amount of complexity into the models, and there are a large number of potential ways that could be done that would have to be explored and compared. Realistically, incorporating all such details will have to be left for the future, and a simplified set of good and bad meme pairs maintained with no specific associated contexts.

The simplest starting point is to assume that there exists an overall set of M memes $\{m_j : j = 1, \dots, M\}$ and that each individual at each stage of their life will have acquired some subset of them and stored them in their brain. Although most animals continue to learn throughout their lives, some kind of leveling off of performance is required, or the natural selection will unrealistically end up being more about age than about what is learned. Moreover, if there is no limit on the number of memes that can be learned by each

individual, all positively contributing memes will tend to be learned by all individuals, and there will be no selection pressure among those memes or between individuals. In practice, there has to be some limit on learning, or an effective “brain capacity”, so each individual i will have a brain with a maximum capacity of B_i memes. At this point, there is no need to specify exactly how these memes were created, or what exactly they represent, nor worry about the details of the meme transfer processes, such as imitation, emulation and teaching, which are still the subject of some debate [24, 30, 45, 46, 54, 62, 65, 83, 84, 87]. It will also be sufficient to begin with the simplifying assumption that all the memes are of equal complexity and imitability, though to have reasonably realistic simulations it will be necessary for different memes to contribute unequally to the overall performance of the individuals that possess them. Of course, mental representations do not really exist as small discrete sets of memes that can be copied directly, and they will not contribute to performance or fitness in a simple additive manner [2, 46, 63], but this simplification is a useful starting point for baseline simulations that simply need a convenient way to “keep score” of how socially learned information or behaviors are affecting individual performances [39]. These approximations can always be relaxed later if they are believed to be having a significant effect on more specific models [40]. This approach is also an approximation in that memes themselves can evolve, and some are known to have evolved to the extent that they can no longer be acquired directly by individuals. Boyd and Richerson [11] have already developed mathematical models showing how this primarily happens when social learning is common, such as in humans, and this is another complication that should be straightforward to incorporate later into the general evolutionary framework proposed here if it is deemed to be relevant for modeling particular scenarios.

The models of Higgs [43] were based on memes of this general type with performance contributions drawn from a Gaussian distribution, and although that makes sense, it has the serious disadvantage that some memes have extremely large contributions that can end up masking the effect of all the others. The preliminary version of the current study [18] went to the other extreme and simply had performance contributions of +1 for good memes and -1 for bad memes, which led to some interesting results, but having no selection pressures to drive the propagation of one good meme over another proved to be too much of an over-simplification. The remainder of this study will therefore have uniformly distributed meme performance contributions in the range [-1, +1], which deals with both the problems of the earlier approaches. In practice, the actual contributions of particular memes will depend on the environment, but that will only need accommodating in models of scenarios with significant changes of environment over time or location, which will not be the case for the test simulations of this paper.

Each individual is born with a baseline performance level of 0 (corresponding to their innate abilities), and throughout their lifetime will acquire memes that either improve or worsen their performance depending on whether they contribute positively (good memes) or negatively (bad memes). Good memes represent correct information concerning a useful behaviour that has a particular level of importance to the individuals’ performance, and bad memes represent information that is incorrect or contradictory in some way to a good meme, and thus reduces their performance by an equal amount. The idea of opposing pairs of memes existing as *cognitive attractors* has previously been studied in the models of Henrich and Boyd [39, 40]. The process of reconciling incompatible ideas and correcting harmful information can then be accommodated by simply having corresponding pairs of memes of equal and opposite contribution (i.e.,

positive and negative) cancel each other out if they are both acquired by the same individual. The brain capacity B_i of each individual i limits the total number of memes it can acquire, and hence the maximum level of performance it can achieve in its lifetime will be limited to that associated with having the set of B_i memes of highest contribution. Obviously, real memes are much more complicated, but this provides a manageable starting point for the simulations, with the intention that more realistic variations can be explored within the same framework later. The simulations are run in exactly the same way as described in Section 2, except that each individual now also needs to have its set of acquired memes and associated performance contribution updated at each time step, along with all its other parameters.

The remainder of this paper will be framed in terms of simple meme imitation, but the same modeling framework will apply for more complex social processes of emulation, teaching and learning. Although there is no need to specify the precise mechanisms by which the memes are created and transferred between individuals, a few more details are required to allow the simulations to be run and replicated. Obviously, successful imitation requires that memes are somehow introduced into the population in the first place. That will usually be by direct individual learning, but to test the imitation process, simulations sometimes need to be run with that disabled. In those cases, with probability r each simulated year, one randomly chosen individual acquires (in some way that need not be specified) one randomly chosen meme if their brain has not already reached full capacity. To avoid any bias towards imitation in the simulations, all the initial evolutionary populations begin with each individual i having an innate imitation rate α_i that is zero. Then repeated crossovers and mutations of the form described in Section 2 will enable the α_i to evolve from there up to a maximum of 1 if that proves beneficial. What the imitation rate means in practice is that during each simulated year, each individual can acquire up to $\alpha_i \phi B_i$ memes from other individuals, where ϕ is a parameter that specifies the maximum rate at which memes can be copied, and the brain capacity B_i is included to establish a scale independence to the imitation rate. One could also easily allow the forgetting of memes, though that will not be done in the simulations of this paper. It is expected that different patterns of results will emerge depending on which species is being modeled, and it is the setting of parameters like ϕ , along with the competition-based and old-age death rates, the natural lifespan, and so on, that tie each simulation to a particular species and environment.

6. Testing the Imitative Learning Framework

To test what emerges from evolution in the above framework with only imitative learning, and to establish suitable values for the various parameters, an initial series of simulations were run with just the imitation rates α_i allowed to evolve. The populations were updated once per simulated year, which automatically provides a realistic gap between births. For each case, average results and standard deviations were computed from 20 independent runs, each of which continued for a predetermined number of simulated years that was considerably longer than needed for the populations to stabilize. **Figure 3** shows the evolving imitation rates and associated individual performance levels obtained for a baseline setup that was found to produce reasonably realistic results in a reasonable amount of computation time, namely a population of 200 individuals with fixed equal brain capacities $B_i = 100$, and fixed meme parameter values $M = 500$, $\phi = 0.1$ and $r = 0.01$, with memes copied from random other individuals, 10% of the population dying each year due to unsuccessful competition, and 20% of individuals aged over 60 years dying each

year of old age. These human-inspired baseline parameter values will be used as the default set in all the subsequent simulations presented in this paper, unless explicitly varied. The performance advantage associated with acquiring memes causes the average imitation rate α to rise quickly, leveling off close to 1.0 after about 50,000 simulated years. The average remains slightly below the ceiling level of 1.0 because the mutations mean some individuals will always have values falling below that. The corresponding performance levels rise more slowly, taking about 1,000,000 years to level off, because the distribution of memes takes longer to optimize. The tournament based selection of parents and survival give the good memes an advantage over bad memes, so the number of bad memes initially rises more slowly than the good memes, then begins to fall, and eventually becomes negligible after about 500,000 years, as shown in the left graph of Figure 4. Then, over the next 500,000 years, the surviving population of good memes is refined to further improve the performance levels of those possessing them. The behaviour during the lifetime of a typical evolved individual begins with an approximately linear acquisition of memes over the first $1/\alpha, \phi = 10$ years, at which point the brain stops acquiring memes and maximum performance is achieved, as shown in the right graph of Figure 4. Most deaths arise due to competition during the meme acquisition period, with relatively few individuals surviving long enough to die of old age, though these numbers can easily be adjusted to model particular species and scenarios by varying the proportion that die each year due to competition and/or introducing protection of offspring.

6.1 Choice of copied individuals

In reality, successful social learning is generally more sophisticated than simply imitating random other individuals at a constant rate, particularly if the environment is rapidly changing [69], though the current simulations provide a useful baseline for future work. Even for stable environments, further simulations reveal interesting dependencies on which individuals are chosen to be imitated as a source of acquired memes, so this is something else that needs to be set appropriately for modeling particular scenarios. Good strategies for social learning are known to be complex, and have been discussed in some detail by Henrich and Gil-White [41] and Laland [56]. For current purposes, it is sufficient to simply have preferential imitation of better performing individuals, which is an old idea [65] that has been confirmed many times [33, 67, 73]. If, rather than copying memes from random other individuals, a minimalist tournament-based selection process is used, copying from the best performing of two randomly chosen individuals, there is more selection pressure for good memes, and the bad memes are eliminated about twice as quickly, after only about 250,000 years. The bad memes can be eliminated even more quickly, in only about 150,000 years, if individuals first learn memes from their own parents, if they are still alive, before randomly choosing other individuals to imitate. This is because their parents have already gone through selection to become parents, and are generally older and more experienced than randomly selected individuals, so they are a better source of good memes. In fact, if individuals *only* imitate their own parents, the number of bad memes building up at any stage of evolution is virtually zero, though that also slightly reduces the number of good memes acquired on average, as shown in the left graph of Figure 5. The advantages arising from individuals only imitating their own parents, and how those benefits depend on the environment being stable, have previously been explored in a rather different type of model by McElreath and Strimling [64].

A related variation that affects what evolves has the choice of who to imitate based on some form of *cultural measure*, that might correspond to just one small aspect of behavior, or, in the extreme case, have no correlation at all with biological performance [9, 43]. To simulate the extreme version of this idea, each meme can be assigned a totally random cultural value, and individuals are chosen for imitation according to the total cultural value they have acquired. This does not prevent the imitation rates from still rising quickly towards 1, and the real performance-based pressures of selection for survival and procreation still manage to reduce the number of bad memes to low levels. However, as the right graph of Figure 5 shows, the elimination of bad memes that happen to have high cultural value is slowed down considerably. In reality, any cultural value of memes will not be totally random, and there will inevitably be interactions between memes that can affect their evolution. Such complications, including memes associating into memplexes [6, 42] and cognitive dissonance reduction [21, 75], are likely to increase or decrease the persistence of bad memes to varying degrees, and these are further issues that may need incorporation into future simulations of more specific scenarios. For example, models of the propagation of certain religious practices might be created in this way.

6.2 Brain capacity

Another important detail that needs to be tested is the setting of a maximum total number of memes that can be acquired by each individual. Since the memes come in good/bad pairs, and only one from each pair can be accommodated in a given brain at a given time, a maximum of $M/2$ of the full set of M memes can be stored, even if the relevant “brain capacity” B_i is bigger than that. In the above simulations, the brain capacity B_i of all individuals was set to be the same fixed value B that was much smaller than the maximum number of storable memes, with $B = 100$ compared to $M/2 = 250$. In such cases, the total number of different memes carried by the population will generally be greater than B , and each individual will tend to acquire a different subset of them before reaching their capacity. The selection pressures can then act on those differences leading to the eventual disappearance of the bad memes as seen in Figure 4. If those simulations are repeated, with everything the same apart from allowing the fixed brain capacities to be more than the maximum number of storable memes (i.e., $B > M/2$), imitation is still beneficial to performance, so the imitation rates α_i still quickly rise to 1, but it is more difficult for the selection pressures to act. In this case, the whole population tends to settle down into acquiring the same set of $M/2$ memes, with only slightly more good memes than bad. Any newly injected meme, whether introduced artificially or via direct individual learning, will either already be among the existing memes and not affect anything, or it will cancel out the opposing meme in one individual, leaving capacity that will quickly be filled again with the original meme being copied from another individual, so the poor set of memes and associated poor performance persists. This persistent poor performance is seen clearly in Figure 6, for the case of $B = M/2 = 250$ in the top-left graph, and for $B = M/2 = 100$ in the top-right graph. Other details can also affect the numbers of good and bad memes, but the lack of improvement over time is a recurring feature. For example, if high performance individuals are preferentially imitated, there is greater selection pressure before a full set of $M/2$ memes is established, and the circulating set of memes tends to settle down into containing a higher proportion of good memes, but still a stable set containing many bad memes persists, as seen for $B = M/2 = 250$ in the bottom-left graph of Figure 6. Another factor that affects

what emerges is the injection rate of new memes. Decreasing that from $r = 0.01$ to 0.002 leads to less selection pressure and allows more bad memes to become established, as seen for $B = M/2 = 100$ in the bottom-right graph of Figure 6. Interestingly, the extreme strategy of each individual only imitating their own parents is able to prevent any significant build-up of bad memes for any brain capacity and any injection rate.

It is clear that the effective brain capacity relative to the total number of potential memes, in conjunction with the meme injection rate and choice of imitation strategy, is a crucial factor that can affect the results that emerge. The choice of who to imitate will normally be relatively easy to set in line with whatever scenario is being modeled, and the meme injection rate should emerge automatically from the direct individual learning, but setting the relative brain size may be more difficult. In practice, there are numerous factors that drive and constrain brain evolution [49, 80], and the number of memes M that can be learned directly and copied will usually co-evolve with the brain size B . Ultimately, that co-evolution will need to be modeled, but for simplified models with fixed B and M it seems reasonable that the number of memes that can be acquired by individuals will always be much smaller than the maximum number of possible memes (i.e., $B \ll M/2$). This paper will therefore continue with the earlier assumption that this is how the simulations should be run, but fixing B and M may need more careful consideration when attempting to model some particular scenarios. Introducing different meme representations, or more complex learning strategies, such as remembering deleted bad memes so they are not relearned later, may affect such brain size consideration, though determining the best way to implement such things in a realistic fashion is another big task that will have to be left for future investigations.

6.3 Mimetic transitions

One of the central recurring features of the earlier Higgs study [43] was the emergence of “mimetic transitions” at which dramatic increases occur in the imitation rates and meme numbers, and the various factors affecting the timing of those transition were investigated. However, those delayed transitions depend on the evolutionary approach adopted, and in all the simulations carried out so far within the framework proposed here, such “transitions” invariably take place immediately at the beginning of the evolutionary process. The reason is that many of the key parameters involved (namely ϕ , r , death rates, lifespan, brain size, etc.) have been set at the same fixed values for all individuals, that are appropriate for specific imaginary imitating species designed to demonstrate the key ideas with evolutionary simulations that take a feasible time to run. In practice, for fully realistic simulations, these parameter values would themselves vary between individuals and evolve, starting from values appropriate for simple non-imitating species, and then “mimetic transitions” can be expected to arise as they change.

One example that illustrates this, and conveniently relates to the issue of brain capacity just discussed, has individual brain sizes B_i that are not fixed, but allowed to evolve from very small values alongside the imitation rates α_i . Figure 7 shows how a typical population evolves in the proposed framework for the case of $M = 200$ memes and random copying selection. Nothing much happens for nearly 400,000 years, and then suddenly the brain sizes B_i , imitation frequencies $\alpha_i \phi B_i$, meme counts and performance levels rise dramatically, like in the mimetic transitions observed by Higgs [43]. The timing of the transitions depend on random mutations and coincidences, and is rather variable, but eventually the imitation rates α_i always

rise to 1 and the brain sizes B_i always increase to $M/2$. The meme counts then level off in a similar way to that seen in Figure 6 for the fixed $B_i = M/2$ cases, though with performance-based imitation selection they can sometimes end up varying more like those in Figure 4. The final proportion of good and bad memes depends on various factors (such as meme injection rate, brain size mutation rate, and imitation strategy), but as long as there are more good memes than bad, a larger brain will be beneficial to each individual and therefore the brain sizes evolve all the way up to $M/2$, even though that sometimes makes it impossible for the population as a whole to eliminate the bad memes, as was seen in the simulations of Figure 6. In real populations, the number of potential memes M will tend to increase in line with the brain sizes B_i rather than remain constant, and various biological costs will constrain the brain size increases [49, 80]. Both of these factors are likely to keep the brain capacities well below $M/2$, and affect the nature of the mimetic transitions, making these further details that are worthy of investigation in the future. These complications will be avoided for the remaining simulations of this paper by reverting to having fixed equal brain sizes B_i that are many times smaller than the fixed number of memes M .

6.4 Transmission fidelity

The next simplification to explore is the unrealistic assumption that imitation is always carried out with perfect accuracy [32, 47, 63], so the effect of varying the copying fidelity f needs to be investigated. That is easily incorporated into the above framework by having a fraction $1-f$ of copied good memes acquired incorrectly, and thereby degraded into worse memes. However, there are numerous ways that idea could be implemented, and varying the details can lead to the evolution of somewhat different patterns of results. The meme performance contributions have been assumed to be distributed uniformly in the range $[-1, +1]$, with good memes defined as those with positive contribution x , and the corresponding bad memes having negative contribution $-x$. Therefore, the obvious way to implement a copying error of a good meme with contribution x would be to replace it by the associated bad meme with contribution $-x$. It seems improbable that an error in copying a bad meme would create a good meme rather than another bad meme (e.g., because there are usually many more bad ways of doing something than there are good ways), so it is natural that the good and bad memes should be treated asymmetrically in that respect. Figure 8 shows how the average performance and meme counts then evolve for a relatively modest 1% error rate (i.e., $f=0.99$), with $M = 500$, fixed relatively small brain sizes $B_i = 100$, and evolved imitation rates α_i . There is a clear reduction in average performance compared to the corresponding perfect copying (i.e., $f = 1.0$) case shown in Figures 3 and 4. The selection pressures for survival and procreation are now insufficient to restrict the proliferation of bad memes, but still the imitation rates α_i are close to 1, and the average performance levels are well above 0. If the choice of which individuals to imitate is performance-based rather than random, the degradation in performance is significantly reduced, as shown in Figure 9, with 4% errors required before the poor performance levels of Figure 8 are reached. The effect of decreasing the copying fidelity is apparent in Figure 9, and shown more clearly in Figure 10. As the number of copying errors increase beyond about 2%, the average performance levels begin to fall drastically, and this is reflected in decreasing numbers of good memes and increasing numbers of bad memes. Eventually, for more than about 10% errors, there are so many bad memes being introduced into the population that imitation becomes a bad strategy and the imitation rates begin to fall to 0, reducing the

numbers of good and bad memes and the average performances to near 0. These simulations thus confirm that this very simple model of transmission fidelity is sufficient to give reasonably realistic results. For more realistic models, costs can be associated with higher fidelity copying, and the performance-based selection pressures will then act to balance them against the advantages of improved imitation.

There are several other ways that copying errors might reasonably be implemented, and **Figure 11** shows typical results that arise for three of them with performance-based imitation and 1% error rate, and all the other details and parameter values the same as for Figures 8, 9 and 10. First, a copied good meme with contribution x could end up as a random meme with worse performance contribution, i.e. a meme with contribution drawn uniformly from the range $[-1, x]$. This has the effect of introducing many new memes into the population, both good and bad, allowing the good memes to dominate even more quickly than before, though bad memes persist in the emergent steady state. A slightly different approach would be to replace an incorrectly copied good meme by a random bad meme, i.e. one drawn uniformly from the contribution range $[-1, 0]$. That means new bad memes are introduced considerably more frequently than good memes, so they initially dominate. Eventually the selection pressures allow the good memes to take over again, but the numbers of bad memes remain high in the stable state. A less drastic meme degradation process would be to simply subtract 1 from the contribution x of an incorrectly copied good meme, which again leads to bad memes in the range $[-1, 0]$, but means the highest contribution memes become only slightly bad, rather than potentially very bad. In this case, the stable state again has the good memes dominating, but many bad memes persist. In all three cases, there is still an advantage to imitation and the imitation rates still quickly evolve towards 1 as in the perfect fidelity case, but as the copying fidelity is reduced, the average performances fall and eventually imitation becomes a poor strategy in a similar way to that seen in Figure 10.

If the simulations were to be based on more realistic meme sets that include a whole range of memes with different performance contributions for each behavioral context, as discussed in Section 5, it would then be natural to have a poorly copied good meme become a random poorer quality meme associated with the same context. Mapping that idea onto the simplified good/bad meme pairs adopted so far leads to the copying-error implementation used for Figures 8, 9 and 10, so that will be employed for the remaining simulations presented in this paper. Ultimately, the copying fidelities and error types will need to be set appropriately for each model, and, if required for particular realistic scenarios, they can be made to vary across different memes or meme types or “imitation” types, or across different individual species, ages or performances. For many species, direct individual learning will be more useful than social learning that is limited by low fidelity, and modeling the trade-offs between the two forms of learning will be required to understand which behaviors are best acquired by which form of learning, and how that varies between species [52, 67, 69]. This issue will be explored in the full dual-learning simulations of Section 8.2.

6.5 Stochasticity

Finally, the imitation of a deterministic number of memes $\alpha_i \phi B_i$ every year in the above simulations is another over-simplification that can easily be avoided. It is a straightforward extension to implement a stochastic version that has the same average meme acquisition rates, but imitation of a random number of memes each year drawn uniformly from the range $[0, 2\alpha_i \phi B_i]$, like the earlier use of a stochastic direct

learning rate drawn from $[0, 2\delta]$. The results emerging from evolutionary simulations of this type are found to follow the same pattern as for the deterministic versions used to generate Figures 3 to 11, though they are inevitably slightly noisier.

7. Formulating the Combined Learning Framework

There is certainly a lot more to memes and imitation than has been introduced above [2, 6, 12, 48], but the proposed simulation framework already includes enough of the crucial details to make progress on the issue of learning in life history evolution. With the key meme-based imitative social learning processes formulated and tested, the direct individual lifetime learning can now be incorporated into the same framework. Mathematical models involving both individual and social learning already exist, such as in the study of the evolution of between-group differences by Henrich and Boyd [38], and the investigation of cumulative cultural learning by Ehm and Laland [26], so the vital issues are understood. However, such models omit many of the crucial details relevant to the life history evolution factors of interest here, such as the speed versus accuracy of learning, and competing individuals of different ages.

The existing abstract stochastic direct learning process outlined in Section 3 has two components: an approximately linear knowledge acquisition with the associated performance increasing each year by an amount drawn uniformly from the range $[0, 2\delta_i]$, and a probability $\rho\delta_i$ of that learning stalling at some random point, where δ_i is an evolvable direct learning rate and ρ is a measure of direct learning difficulty. Ideally, the form of direct learning in the new meme-based framework should be consistent with that, and also match as closely as possible the stochastic imitative learning of memes at a rate drawn uniformly from the range $[0, 2\alpha_i\phi B_i]$ each year, where α_i is the evolvable imitation rate, ϕ is a measure of ease of copying, and B_i is the brain capacity. The natural way of doing that has a random number of random good memes learned directly each year drawn uniformly from the range $[0, 2\delta_i\psi B_i]$, where δ_i is an evolvable direct learning rate and ψ is a measure of ease of direct learning. The learning task difficulty component, that prevents the evolution of unrealistically high learning rates, also has a straightforward meme-based implementation in the form of a probability $\rho\delta_i$ of learning a bad meme rather than a good meme. Then evolution will balance the trade-off between learning quickly with high learning rates δ_i and having a high probability $\rho\delta_i$ of learning performance-reducing bad memes, with results equivalent to the full neural network simulations of the earlier study [17].

The direct learning parameters ψ and ρ will, like the corresponding imitative learning parameters ϕ and f , need to be set to match the particular scenario being modeled. In the test simulations presented in this paper, these parameters are each assigned a single value, but one can easily have the overall performance based on multiple tasks or skills, each represented by different parameter values. The evolved learning rates must lead to appropriate expected learning times $1/\delta_i\psi$, and the meme numbers and acquisition rates should align with the relevant form of imitative learning. Simulations of just the direct learning process show that, whatever the chosen parameter values, the evolutionary process quickly settles to a stable state, but particular parameter values need to be identified that coordinate the direct learning with the associated imitative learning timescales. **Figure 12** shows how the learning rates δ_i and meme numbers evolve for a baseline direct-learning-only population of 200 with $M = 500$, $B_i = 100$, $\psi = 0.1$ and $\rho = 0.04$, which lead to human-like learning timescales similar to those previously simulated for imitative learning on its own.

Consequently, the values $\psi = 0.1$ and $\rho = 0.04$ will be added to the set of default parameter values used for the remaining simulations presented in this paper.

In the simpler preliminary version of this study [18], all the good memes had an equal performance contribution, so the overall performance was not affected by which particular memes were learned. In the improved framework here, that has memes with a range of contributions, the direct learning process will not necessarily acquire the highest contribution memes, so imitation might provide a better source of high quality memes. This ties in with real human learning in which much of our most valuable information and behaviors tends to be acquired from other individuals rather than by direct learning. A complication that arises from having two distinct learning mechanisms is that it becomes possible for an individual to acquire both a good meme and the corresponding bad/incorrect “version” via different routes. In reality, resolving meme inconsistencies is known to be a complex issue [22], but the minimalist solution adopted above of having corresponding good and bad memes simply cancel each other out worked well, so that approach will be followed here too. Any newly copied meme will cancel out the opposing version if it was previously acquired by either direct learning or by copying, and any new meme acquired by direct learning will cancel out the opposing version if previously acquired by copying. However, the direct learning probabilities only make sense if only the first of each opposing pair of memes acquired by direct learning is kept, and the rest ignored. Finally, now that memes arise from direct learning, the artificial injection of random memes to start the imitation process is no longer required. More realistic models may require the inclusion of some memes that can only be acquired by social learning, and others that can only be learned directly, but that complication will not be included in this initial study.

Having formulated and tested both the direct and imitative meme-based learning processes separately, they can now be simulated simultaneously, and their combined interaction with evolvable life history traits like protection periods can be investigated. It is clear already that there are numerous implementational details that can have a significant effect on what emerges from the simulations, and those differences will reflect the diversity of life history patterns that have emerged from the biological evolution of different animal species. It is this ability to vary the details and parameters, and thereby perform a systematic exploration of the different trade-offs and interactions underlying specific traits, that makes the proposed simulation framework approach so powerful. However, given the number of free parameters that were deliberately introduced, so that they can be set to appropriate values for modeling a wide range of particular scenarios, it will be impractical to provide results for all possible combinations. The following section presents an illustrative selection of results for a few concrete scenarios, primarily concerned with the evolution of learning-driven protection periods, which serve to demonstrate the range of factors that can be investigated within this framework.

8. Testing the Simulation of Life History Evolution

As always, the details of the simulations need to match the particular species and scenario being modeled, and the introduction of additional life history factors, such as protection periods, means that some of the baseline parameter values used above are no longer appropriate. In particular, different death rates per year emerge for real species depending on numerous other factors, such as the given environment, the population size in relation to the available food resources, the number of predators, the age distribution of

the population, the distribution of life-spans, the number of children born each year, the level of protection those children receive, and so on. In fully realistic simulations, all these interacting factors co-evolve and settle down into an ecosystem of many species, with occasional extinctions and formations of new species, with birth rates, death rates, etc. that emerge automatically depending on the environment. For smaller practical simulations with fixed population sizes, key parameters such as the death rates need to be set by hand to appropriate values so that all the other details settle down to appropriate values close to those observed for the species under consideration, e.g. humans.

8.1 Individual learning only

Before simulating the two learning processes together, it is important to check that the earlier life history evolution results for direct learning only [17], as outlined in Section 3, are replicated in the new meme-based formulation. Previously, a 10% death rate per year was reasonable when the whole populations was eligible for death by competition, but that leads to very few individuals surviving for long after extended protection periods. That in turn results in unrealistic age distributions, unrealistic numbers of children per year per eligible individual to maintain the population size, and very little performance-based selection pressure. Introducing protection therefore means the competition-based death rate needs to be adjusted to bring everything back into line with real populations. Testing a range of parameter values shows that a reduced death rate of 2% per year and reduced task difficulty parameter $\rho = 0.01$ are appropriate for the emergence of realistic age distributions, with the learning rates and protection periods evolving as shown in Figure 13. Initially, before significant protection periods have evolved, the learning rates quickly grow to values similar to those in Figure 12, but they then fall as the protection period increases, because slower learning leads to better performance after the protection period ends. The protection period eventually levels off, because procreation is not allowed while protected, and there is a trade-off between reduced procreation period and improved performance-based selection during that period, as was also found in the earlier non-meme-based simulations [17].

A potentially important detail here concerns whether at least one of an individual's own two parents must be alive for them to receive protection. In some real populations, only parents offer protection, while in others, protection may be available from extended families or alloparents [13, 36, 68]. For the models underlying Figure 13, living parents were not required for a child to receive protection, but very similar results arise when they *are* required, because it is rare for both parents to die before the protection period ends. However, it is to be expected that alloparenting will still emerge, if appropriate, in more realistic simulations that allow its evolution and take full account of its benefits [68]. Another important variation allows offspring to procreate while being protected. In that case, the simulated individuals cease to have a cost or disadvantage associated with being protected, so enormous protection periods emerge, and deaths due to competition drop to zero. Also, selection pressure from an early age to procreate causes higher learning rates to evolve, resulting in faster learning but inferior final performance. This again replicates the earlier study [17] outlined in Section 3. In this case, both parents will usually die before the end of the innate protection period, so requiring living parents for protection leads to deaths through competition again. However, that has relatively little effect on the learning rates and performance, because selection pressure to procreate from an early age already results in high learning rates.

8.2 Social versus individual learning

Although the details depend on the chosen parameter values, the broad pattern of results across all the direct learning variations is consistent with the earlier non-meme-based simulations [17]. This means the imitation-based learning can now be incorporated with some confidence as discussed in Section 7. First, the simplest case of no protection period was tested, leaving only the direct learning rates δ_i and imitation rates α_i to evolve in the simulations. If the imitation is perfect ($f = 1.0$), the direct learning is the only source of bad memes, so very low learning rates ($\delta_i \sim 0.02$) evolve to keep the number of bad memes low, and the imitation rates rise quickly to near 1.0. For slightly lower copying fidelities (e.g., $f \sim 0.98$), both the evolved imitation and direct learning rates end up close to 1.0. When the copying fidelity is much lower (e.g., $f < 0.9$), poor imitation introduces so many bad memes that the evolved imitation rates become very low ($\alpha_i \sim 0.05$) and memes are mainly acquired by direct learning. The full dependence on copying fidelity in this case is shown in the top-left graph of Figure 14, which is consistent with and extends the results from various mathematical models [9, 26, 28, 71].

In reality, the trade-offs between social and individual learning will be complicated by the range of meme types and complexities associated with different tasks and contexts, different copying fidelities for different memes, and the interaction of particular meme copying fidelities with other acquired memes and performance levels. Complex behaviors may appear more difficult to imitate with high fidelity, but they are also likely to be more difficult to acquire directly by trial and error. Indeed, some memes have evolved to the extent that they can no longer be learned directly [11]. Other behaviors might look simpler, and appear easier to imitate, but they may rely on precise motor control that is most easily acquired by individual practice. Some behaviors may even require active teaching for them to be passed on with sufficiently high fidelity, and that may limit the possibilities for some species [45, 54, 83]. The simulation framework proposed here, however, is general enough that it can easily be extended to accommodate such complexities, for example, by implementing separate copying fidelities for each meme, and by allowing those fidelities to depend on the presence of other memes. Nevertheless, it will clearly not be easy to model all these things accurately, and in many cases there is still a lack of empirical data, but experimental methodologies do already exist to explore many of the relevant issues, and data does exist in the literature, for some species at least [24, 30, 84].

8.3 Evolvable protection periods

When protection periods are also allowed to evolve in the simulations, the broad pattern of evolved direct learning and imitation rates is the same, but the details change to those shown in the top-right graph of Figure 14. The lack of competition during the protection period allows the direct learning to be slower and more accurate in this case, with the number of bad memes kept to lower levels. For high copying fidelities (e.g., $f > 0.99$), the learning is mostly by imitation, and for low fidelities (e.g., $f < 0.95$), it is mostly by direct learning. For all fidelities, the evolved protection period is slightly longer than the typical learning time, which is a feature found quite generally, including in the neural network-based direct learning simulations of Figures 1 and 2. For the simulations here, based on the same parameter values chosen earlier to result in human-like age distributions and learning timescales, protection periods

emerge around 16 years for low fidelities, and around 14 years for high fidelities, with a significant dip in between to around 8 years, when both types of learning are being used together, as shown in the bottom-right graph of Figure 14.

The fact that evolution again chooses between imitation and direct learning depending on the copying fidelities, to minimize the number of bad memes and maximize performance, is easily understandable. However, in the simpler preliminary study [18], it was found that errors arising from faster direct learning could be corrected by high-fidelity imitation, allowing faster direct learning to emerge without a negative impact on the final adult performance. In that case, extended protection periods allowed more time to correct the too-fast direct learning errors, rather than reducing the errors by allowing slower and more accurate direct learning. The key difference between the two sets of simulations is the distribution of meme performance contributions and how the bad memes arise and are corrected. As discussed earlier, the more sophisticated approach of this paper is more reliable and realistic, but this does highlight the importance of testing the robustness of the results with respect to such technical details.

8.4 Procreation while protected

Some of the most interesting results from the earlier study with only direct learning [17] concerned the advantages and disadvantages of allowing procreation while individuals are being protected. As discussed above, it was found that if procreation while protected (PWP) *was* allowed, the evolved protection periods rose to the extent that there were no competition-based deaths, and individuals only died of old age. Then the consequence of the selection pressure to procreate from an early age was faster and poorer quality learning, and that inevitably led to poorer adult performance. Equally large protection periods emerge, for the same reasons, in the full meme-based framework here. The bottom-left graph of Figure 14 shows the evolved learning and imitation rates in this case. The learning rates are again higher than without PWP, for all copying fidelities f . Obviously, it is not really possible for everyone to be protected till old age, because a reasonable fraction of the population has to do the protecting, and the costs of protecting will also limit the emergent protection periods [20, 68], but the extreme case simulated here does point to the direction of the effects that will remain to some extent in more realistic simulations.

8.5 Performance analysis

It is reassuring that the evolved parameters make sense, but it is the performance levels that drive the evolutionary selection pressures and they also deserve scrutiny, particularly since understanding how the performance varies with the fixed parameters and strategies of simple models often provides useful clues about how those things are likely to evolve in more realistic models or in real life. Performance is clearly age dependent, so the average performance overall and average adult performance need to be considered separately, where *adult* means beyond the learning period. Here, adults are defined as individuals aged over 20 years, which is safely above the expected learning time for all cases. Figure 15 shows the average final population performances (left) and average final adult population performances (right), as a function of the copying fidelity f for the simulations of Figure 14. As expected, average adult performance is always greater than the corresponding average over all ages, and higher copying fidelities always result in similar or higher performance. The no-protection case can be considered as the baseline. Then, when

protection is allowed, it evolves along with lower direct learning rates, and the adult performances are significantly increased for all copying fidelities, except for the perfect fidelity case where the evolved direct learning rate is close to zero for both cases. For high non-perfect fidelities ($1.00 > f > 0.98$), the reduction in the direct learning rates are particularly dramatic, and that leads to large adult performance improvements. For lower fidelities, the adult improvement is smaller, and the whole population averages are actually reduced by the increased numbers of protected slower-learning non-adults with low performance. When procreation while protected (PWP) is allowed, the advantage of not needing to compete at young ages is lost, and the results come closer to the no-protection case (as in the earlier direct learning only simulations [17]). In this case, there is no removal of the worst performing individuals from the adult population, so the average adult performances are actually lower than the no-protection case, particularly when the copying fidelity is perfect.

9. Discussion and Conclusion

This paper has achieved its original aim of investigating the effect of learning on the evolution of protection periods, and in doing so has ended up formulating and testing a general modeling framework for exploring the effects of learning on life history evolution. This has involved the introduction of imitative social learning and memes into agent-based artificial life simulations in such a way that there is a consistent interaction with direct individual learning and other life history factors. Throughout, effort has been made to maintain flexibility, allowing a computationally efficient way of parameterizing and exploring any hypotheses in this field for any species. Inevitably, presentation and testing of the core framework has required numerous simplifications and approximations, but these have been highlighted when introduced, and the modular approach to the basic structures and ideas allow any such details to be represented more accurately as required in the future for more specific and realistic scenarios.

The precise patterns of results that emerge have been seen to depend on numerous details beyond those routinely optimized in evolutionary simulations [27], such as how exactly the memes are represented and distributed, how their copying fidelities vary, how the direct and imitative learning routes interact, how bad memes arise, and how meme inconsistencies are resolved. Some of the key possibilities have already been tested and discussed in this paper, though others remain for future work and may lead to further improved variations within the proposed framework. It will also be important to model more realistically the different learning strategies that emerge in environments which change significantly over individual lifetimes, or over generational timescales. Such variations have already been accommodated in some mathematical models [26, 38], and different time-weighted learning strategies are known to be superior in such situations [69]. One advantage of the proposed simulation approach is that such strategies can be encoded genetically and optimized by evolution along with everything else. Moreover, related to this are environmental conditions that vary with location, and incorporating those could allow the emergence of cultural differences between groups to be studied in more detail [38]. These factors illustrate the need for a well-defined, flexible and tested simulation framework for modeling the various interacting factors and trade-offs as accurately as possible in line with particular realistic scenarios, that facilitates reliable and reproducible comparisons with minimal confounding factors.

It is clear that the modeling framework presented in this paper can enable the systematic investigation

of an enormous range of interactions and trade-offs, allowing models to be tuned to simulate the known life history patterns of particular natural species (e.g., concerning offspring numbers [60], weaning [53], alloparenting [13, 36, 68], delayed maturity [79], diet [51], lifespan and age at first reproduction [34], and so on). The balance between the two forms of learning, parameterized here by ϕ , ψ , f and ρ , determines what emerges, and progress will be made by attempting to understand species specific differences in terms of evolved variations in such parameter values. For example, if the copying fidelity f is very low, a high imitation rate α_i may never evolve because it introduces too many bad memes into the population, and one ends up with direct learning only, as appears to be the case for many non-human animal species. The extreme cases, for which our intuitions are most reliable, tend to be too unrealistic. For example, perfect copying fidelity is unlikely to ever be possible in practice, and the reduction of direct learning to zero is problematic in that it will not allow any new information to be introduced into the meme pool and thus relies on good memes being consistently passed on through the generations. Length considerations have limited this paper to presenting results from only a small selection of the simulations carried out, primarily those confirming successful testing and optimization of the core modeling framework, and a few that illustrate the kind of issues that can be explored, along with some of the main variations that can be incorporated, and the key potential difficulties that need to be avoided. However, all the simulation results obtained so far are broadly consistent with existing intuitions, and simplified mathematical models, and that instills confidence in being able to take them further with some reliability to explore issues for which our intuitions are not so clear and controversy remains.

More detailed experiments studying further issues for particular species will be reported elsewhere, but it is worth noting at this stage a number of refinements which could be implemented in the presented framework without too much effort. First, relating to the key learning aspect, would be improvements to the way direct learning is parameterized, that are able to accommodate different types of animal learning in a more realistic fashion. For example, there are often sensitive or critical periods for learning, or even crucial sequences of such learning periods, or requirements for “memes” to be acquired in particular orders [14, 35, 50, 86]. Preliminary attempts involving more parameters and different distributions of good and bad memes have resulted in the various trade-offs being re-balanced slightly, but so far no fundamentally different behaviors have emerged. Several different imitation strategies have already been explored in this paper, but there is also scope for exploring conformist strategies that involve copying majority behaviors rather than particular individuals [9, 38, 42], and for more careful consideration of which strategies are appropriate for particular realistic scenarios [56, 69]. Incorporating the potential for associations between good and bad memes and the creation of memplexes [6, 42], and more realistic mechanisms for dealing with cognitive dissonance [22, 75], are other avenues for future improvements, though it is not clear what fundamentally new life history evolution results might emerge from them.

There are several things that more realistic models will require to refine the management of memes that so far arises by the simulation of natural selection based on the blind addition of meme performance contributions. The details of the processes available to reduce the numbers of bad memes can have a particularly dramatic effect on the simulation results. If insufficient mechanisms are available to remove or correct bad memes, cycles in the imitation rates can arise throughout evolution. For example, when the number of bad memes rises to very high levels, the best strategy may be to stop copying until all the

carriers have died, and then start again in conjunction with direct learning until the same problem arises again [18]. Such strategic switching between social and direct learning has been considered before and is known to be beneficial in other circumstances too [9, 11, 28, 52]. Conscious understanding that certain memes are bad (i.e., contributing negatively to performance), and their consequent deletion, may also be possible, for some species at least. These kinds of *critical social learning* have already appeared in simplified mathematical models [11, 26, 28], and been shown to have significant consequences. They could easily be studied in more detail within the framework proposed here, but the details will need careful consideration in light of the specific scenarios being simulated, particularly since the simulations reported above have already demonstrated the enormous effect even simple choices of which individuals to imitate can have on what emerges. A further complication, that is increasingly important for humans, is that some learning does not involve direct interaction with either the environment or other individuals, but comes indirectly via information stored in artifacts such as books or websites. Adding such additional sources of memes into the proposed framework would be straightforward, but it will inevitably render the simulations more complex to set up reliably.

Another important issue is that, while individuals are improving their performance by learning, they are often also improving it simply by growing, and that affects many other life history traits such as their physical ability to reproduce or learn particular tasks, and their need for nursing and protection. In principle, such factors would be straightforward to incorporate into the simulations for particular species, and it will be interesting to model how the learning and growing co-evolve with the other life history traits like parental protection. Also, in most species, performance begins a natural physiological-based decline beyond a certain age [37], and incorporating that will lead to more realistic accounts of deaths due to old age. Another factor closely related to life history evolution is the evolution of brain size and complexity [49, 80]. If the brain sizes are allowed to evolve in simulations of the form described above with a fixed number of potential memes, they increase until they can accommodate a full set of memes, as seen in Figure 7. In the evolution of real animals, the meme set is likely to co-evolve with everything else, and the brain sizes will not be able to increase indefinitely to accommodate them, because there are significant costs associated with having larger brains. It is the balancing of those costs against the improved performance resulting from the bigger brain that leads to particular brain sizes emerging [6, 49, 59, 80]. The general simulation framework proposed in this paper has been designed to make it easy to incorporate additional costs associated with any particular parameter, and that would limit the emergent brain sizes in a more realistic fashion. The difficulty is that models of the costs involved here are rather species specific, and care will be needed to implement them accurately enough for the simulations to provide reliable testable predictions.

The most relevant cost for the current study is that of learning itself, and that will clearly depend on the species and what they need to learn to perform well. If direct and imitative learning have significantly different costs, that will shift the balance between them in graphs such as those of Figure 14, but the qualitative pattern is likely to remain the same. If both forms of learning are sufficiently costly, a point may be reached where the learned performance is good enough that further learning ceases to be beneficial. This may be difficult to model, because the costs of learning are likely to depend on the performance levels already achieved. For example, once the initial costly learning to survive and observe

and copy reliably has been completed, further refinement of performance may be relatively cost-free. Moreover, if an infant is being protected, the cost of learning is likely to be less relevant than if they are not, but it may pass an additional cost on to their protectors. The simulations presented in this paper have deliberately not included any costs of learning, in order to provide a clear baseline case. Incorporating such costs in future models will simply require more complex functions to represent the learned performance levels, but identifying suitable functions for particular species and scenarios is likely to require some degree of experimentation, particularly when the models involve the interaction of several different species. It is also becoming clear that active teaching, obviously widespread in humans, is also more common in other animals than previously thought [45, 83], and the costs and benefits of that are further factors that could benefit from more accurate modeling.

Another learning cost issue relates to the interaction between learning and evolution and the Baldwin Effect [3, 23, 81, 85]. The question of whether adaptive learning accelerates or hinders genetic evolution has been much studied [44, 81, 85], but it is the genetic assimilation of learned behaviors [23, 58, 85] that is most relevant to life history evolution here. The essential idea is that if genetic changes can lead to individuals being born “closer” to a good learned behaviour, then the cost of learning that behaviour will be reduced, leading to a selection advantage for that individual. In that way, learned behaviors may gradually become innate. Simple transference of learned neural connection strengths to innate initial neural connection strengths has already been demonstrated in agent-based evolutionary simulations of the type proposed here [14]. Of course, more complicated behaviors will be relatively difficult to encode genetically, and less likely to become assimilated into the genotype, so worrying about the assimilation of many realistic memes will probably not be necessary. Moreover, good adult behaviors will not necessarily be appropriate for new-born individuals, in the same way that learned adult neural connection strengths can be sub-optimal for new-borns [14], and having optimal innate behaviors will be difficult in spatially or temporally changing environments [23], so some lifetime learning will still be required even when perfect genetic assimilation is possible. Nevertheless, genetic assimilation is likely to have been an important issue for the life history evolution of many species, so it will eventually need to be incorporated into the approach proposed in this paper. This may not be easy, but the proposed framework is general enough to accommodate it. Identifying simple-enough memes, or sequences of memes, that might be assimilated, and running simulations with general-enough genotypes to accommodate them, would be the obvious first step, though doing that for any but the simplest species is likely to be challenging. Crispo [23] discusses several species that may provide a profitable starting point for such an endeavor.

Further cost-related issues for future work involve the incorporation of more realistic indirect costs, such as those due to gestation or providing parental support [21], or allowing copying, or being an active teacher [36, 83]. Again, these will all be rather species and environment dependent, influenced by other factors such as the number of children being produced and protected by each individual [53, 60], whether the children are protected by individuals other than their own parents [23, 68], and what restrictions are placed on the protected children (like the no procreation example of Figures 13 and 14). Also relevant here will be the need for a clearer distinction between different types of learned behaviour, and related factors such as their costs, the distribution of the meme performance contributions, and the relative ease of copying particular types of memes.

It is also clear that while humans, and many other animals, do have long childhoods during which they cannot reproduce, there is more to it than simple restrictions imposed by their protecting parents [72, 78]. There are usually unrelated minimum ages for reproduction that are triggered by body weight and other innate factors. Moreover, child survival probabilities, in relation to the costs of pregnancy and child rearing, may also lead to performance related delays to reproduction that are independent of the performance-based selection process. It may simply be more cost effective to delay reproduction until parental performance allows a better chance of child survival. In modern humans, reproduction delays may also be more about acquiring enough wealth than enough learning, though those factors are not unrelated. So, although it is informative to see how learning factors alone can lead to the evolution of protection periods with associated delayed reproduction, in practice it will have co-evolved with various other factors, and it will be a challenge for the future to incorporate those additional factors into the models proposed here to simulate and understand the trade-offs and synergies between them.

Another crucial feature of many real populations is the presence of two distinct genders, potentially with different performance specifications, cooperation and reproduction strategies, selection pressures and mutation rates, and the requirement that parents be of different genders. This will inevitably render the simulations more complex, but the general idea is understood and many aspects have already been explored [74]. Likewise, the interaction and co-evolution of multiple species within a single simulation [31], and the ability of individuals of the same evolved species to cope and compete or cooperate in different or changing cultural situations or physical environments [58].

In conclusion, this paper has presented and tested a general flexible meme-based framework for simulating learning as a factor in life history evolution, explored the effect of various design choices and implementational variations, demonstrated its power for resolving the trade-offs between conflicting factors for some representative cases, and suggested a number of avenues for future research within it. Moreover, the resulting framework is general and powerful enough to enable improved simulations of aspects of social learning and cultural evolution unrelated to life history evolution. Hopefully this framework will stimulate further work in these areas, with future publications presenting the results it generates for more specific and realistic scenarios.

Acknowledgements

I thank James Borg who originally inspired me to introduce social learning into my earlier models and facilitated progress by inviting me to give regular talks on the results, and also the anonymous reviewers who provided invaluable suggestions for improving this paper. A preliminary version of this study was presented at the *Artificial Life XII Conference* [18].

References

1. Arbib, M. A., editor (2003). *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press.
2. Aunger, R. A. (2002). *The Electric Meme: A New Theory of How We Think*. New York, NY: Simon and Schuster/Free Press.
3. Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30, 441-451.
4. Belew, R. K. (1990). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4, 11-49.
5. Best, M. L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7, 289-306.
6. Blackmore, S. (1999). *The Meme Machine*. Oxford, UK: OUP.
7. Bogin, B. (1997). Evolutionary hypotheses for human childhood. *Yearbook of Physical Anthropology*, 40, 63-89.
8. Borenstein, E. & Ruppin, E. (2003). Enhancing autonomous agents evolution with learning by imitation. *AISB Journal*, 1, 335-347.
9. Boyd, R. & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago, IL: University of Chicago Press.
10. Boyd, R. & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16, 125-143.
11. Boyd, R. & Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy of Science*, 88, 73-93.
12. Brodie, R. (1996). *Virus of the Mind: The New Science of the Meme*. Seattle, WA: Integral Press.
13. Briga, M., Pen, I. & Wright, J. (2012). Care for kin: Within-group relatedness and allomaternal care are positively correlated and conserved throughout the mammalian phylogeny. *Biology Letters*, 8, 533-536.
14. Bullinaria, J. A. (2003). From biological models to the evolution of robot control systems. *Philosophical Transactions of the Royal Society of London A*, 361, 2145-2164.
15. Bullinaria, J. A. (2007). Using evolution to improve neural network learning: Pitfalls and solutions. *Neural Computing and Applications*, 16, 209-226.
16. Bullinaria, J. A. (2007). Understanding the emergence of modularity in neural systems. *Cognitive Science*, 31, 673-695.
17. Bullinaria, J. A. (2009). Lifetime learning as a factor in life history evolution. *Artificial Life*, 15, 389-409.
18. Bullinaria, J. A. (2010). Memes in artificial life simulations of life history evolution. In: *Proceedings of Artificial Life XII Conference (Alife XII)*, 823-830. Cambridge, MA: MIT Press.
19. Byrne, R. W. & Russon, A. E. (1998). Learning by imitation: A hierarchical approach. *Behavioral and Brain Sciences*, 21, 667-684.
20. Clutton-Brock, T. H. (1991). *The Evolution of Parental Care*. Princeton, NJ: Princeton University

Press.

21. Clutton-Brock, T. H., Albon, S. D. & Guinness F. E. (1989). Fitness costs of gestation and lactation in wild mammals. *Nature*, 337, 260-262.
22. Cooper, J. (2007). *Cognitive Dissonance: 50 Years of a Classic Theory*. London, UK: Sage.
23. Crispo, E. (2007). The Baldwin effect and genetic assimilation: Revisiting two mechanisms of evolutionary change mediated by phenotypic plasticity. *Evolution*, 61, 2469-2479.
24. Custance, D. M., Whiten, A. & Bard, K. A. (1995). Can young chimpanzees (pan troglodytes) imitate arbitrary actions? Hayes & Hayes (1952) revisited. *Behaviour*, 132, 837-859.
25. Dawkins, R. (1976). *The Selfish Gene*. Oxford, UK: OUP.
26. Ehn, M. & Laland, K. (2012). Adaptive strategies for cumulative cultural learning. *Journal of Theoretical Biology*, 301, 103-111.
27. Eiben, A. E. & Smith, J. E. (2015). *Introduction to Evolutionary Computing*. Berlin, Germany: Springer.
28. Enquist, M., Eriksson, K. & Ghirlanda, S. (2007). Critical social learning: A solution to Rogers's paradox of non-adaptive culture. *American Anthropologist*, 109, 727-734.
29. Gadgil, M. & Bossert, W. (1970). Life history consequences of natural selection. *American Naturalist*, 104, 1-24.
30. Galef, B. G. & Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioScience*, 55, 489-499.
31. Gandon, S., Agnew, P. & Michalakis, Y. (2002). Coevolution between parasite virulence and host life-history traits. *American Naturalist*, 160, 374-288.
32. Gold, B. J., Pomplun, M., Rice, N. J. & Sekuler, R. (2008). A new way to quantify the fidelity of imitation: Preliminary results with gesture sequences. *Experimental Brain Research*, 187, 139-152.
33. Greenfeld, N. & Kuznicki, J. T. (1975). Implied competence, task complexity, and imitative behavior. *Journal of Social Psychology*, 95, 251-261.
34. Harvey, P. H. & Zammuto, R. M. (1985). Patterns of mortality and age at first reproduction in natural populations of mammals. *Nature*, 315, 319-320.
35. Harwerth, R. S., Smith, E. L., Duncan, G. C., Crawford, M. L. J. & von Noorden, G. K. (1986). Multiple sensitive periods in the development of the primate visual system. *Science*, 232, 235-238.
36. Hawkes, K., O'Connell, J. F., Blurton Jones, N. G., Alvarez, H. & Charnov, E. L. (1998). Grandmothering, menopause, and the evolution of human life histories. *Proceedings of the National Academy of Sciences USA*, 95, 1336-1339.
37. Helton, W. S. (2008). Expertise acquisition as sustained learning in humans and other animals: Commonalities across species. *Animal Cognition*, 11, 99-107.
38. Henrich, J. & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19, 215-241.
39. Henrich, J. & Boyd, R. (2002). On modeling cultural evolution: Why replicators are not necessary for cultural evolution. *Journal of Cognition and Culture*, 2, 87-112.
40. Henrich, J., Boyd, R. & Richerson, P. J. (2008). Five misunderstandings about cultural evolution.

Human Nature, 19, 119-137.

41. Henrich, J. & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22, 165-196.
42. Henrich, J. & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12, 123-135.
43. Higgs, P. G. (2000). The mimetic transition: A simulation study of the evolution of learning by imitation. *Proceedings of the Royal Society B: Biological Sciences*, 267, 1355-1361.
44. Hinton, G. E. & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495-502.
45. Hoppitt, W. J. E., Brown, G. R., Kendal, R., Rendell, L., Thornton, A., Webster, M. M. & Laland, K. N. (2008). Lessons from animal teaching. *Trends in Ecology & Evolution*, 23, 486-493.
46. Horner, V. & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (pan troglodytes) and children (homo sapiens). *Animal Cognition*, 8, 164-181.
47. Huber, L., Range, F., Voelkl, B., Szucsich, A., Viranyi, Zs. & Miklosi, A. (2009). The evolution of imitation: What do the capacities of non-human animals tell us about the mechanisms of imitation? *Philosophical Transactions of the Royal Society B*, 364, 2299-2309.
48. Hurley, S. & Chater, N., editors (2005). *Perspectives on Imitation: From Neuroscience to Social Science*. Cambridge, MA: MIT Press.
49. Isler, K. & van Schaik, C. P. (2009). The expensive brain: A framework for explaining evolutionary changes in brain size. *Journal of Human Evolution*, 57, 392-400.
50. Johnson, M. H. (2005). Sensitive periods in functional brain development: Problems and prospects. *Developmental Psychobiology*, 46, 287-92.
51. Kaplan, H., Hill, K., Lancaster, J. & Hurtado, A. M. (2000). A theory of human life history evolution: Diet, intelligence, and longevity. *Evolutionary Anthropology*, 9, 156-185.
52. Kendal, R. L., Coolen, I., van Bergen, Y. & Laland, K. N. (2005). Trade-offs in the adaptive use of social and asocial learning. *Advances in the Study of Behavior*, 35, 333-379.
53. Kennedy, G. E. (2005). From the ape's dilemma to the weanling's dilemma: Early weaning in its evolutionary context. *Journal of Human Evolution*, 48, 123-145.
54. Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences*, 38, e31.
55. Klug, H. & Bonsall, M. B. (2009). Life history and the evolution of parental care. *Evolution*, 64, 823-835.
56. Laland, K. N. (2004). Social learning strategies. *Animal Learning and Behavior*, 32, 4-14.
57. Lamb, M. E., Bornstein, M. H. & Teti, D. M. (2002). *Development in Infancy: An Introduction*. Mahwah, NJ: Lawrence Erlbaum Associates.
58. Lande, R. (2009). Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. *Journal of Evolutionary Biology*, 22, 1435-1446.
59. Leigh, S. R. (2004). Brain growth, life history and cognition in primate and human evolution.

- American Journal of Primatology*, 62, 139-164.
60. Leutenegger, W. (1979). Evolution of litter size in primates. *American Naturalist*, 114, 525-531.
 61. Lindblom, J. & Ziemke, T. (2003). Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior*, 11, 79-96.
 62. Marriott, C., Parker, J. & Denzinger, J. (2010). Imitation as a mechanism of cultural transmission. *Artificial Life*, 16, 21-37.
 63. Marsh, L. E., Ropar, D. & Hamilton, A. F. d. C. (2014). The social modulation of imitation fidelity in school-age children. *PLoS ONE*, 9(1), e86127.
 64. McElreath, R. & Strimling, P. (2008). When natural selection favors imitation of parents. *Current Anthropology*, 49, 307-316.
 65. Miller, N. E. & Dollard, J. (1941). *Social Learning and Imitation*. New Haven, CT: Yale University Press.
 66. Nehaniv, C. L. & Dautenhahn, K., editors (2009). *Imitation and Social Learning in Robots, Humans and Animals*. Cambridge, UK: CUP.
 67. Offerman, T. & Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of Economic Behavior and Organization*, 34, 559-575.
 68. Reidman, M. L. (1982). The evolution of alloparental care and adoption in mammals and birds. *Quarterly Review of Biology*, 86, 255-261.
 69. Rendell, L. et al. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, 328, 208-213.
 70. Richerson, P. J. & Boyd, R. (1992). Cultural inheritance and evolutionary ecology. In: E. A. Smith and B. Winterhalder, editors, *Evolutionary Ecology and Human Behaviour*, 61-91. Chicago, IL: Aldine de Gruyter.
 71. Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90, 819-831.
 72. Roff, D. A. (2002). *Life History Evolution*. Sunderland, MA: Sinauer Associates.
 73. Rosenbaum, M. E. & Tucker, I. F. (1962). The competence of the model and the learning of imitation and non-imitation. *Journal of Experimental Psychology*, 63, 183-190.
 74. Sanchez-Velazco, J. & Bullinaria, J. A. (2003). Sexual selection with competitive/co-operative operators for genetic algorithms. In: *Proceedings of the IASTED International Conference on Neural Networks and Computational Intelligence (NCI 2003)*, 191-196. IASTED/ACTA Press.
 75. Shultz, T. R. & Lepper, M. R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review*, 103, 219-240.
 76. Smaldino, P. E., Newson, L., Schank, J. C. & Richerson, P. J. (2013). Simulating the evolution of the human family: Cooperative breeding increases in harsh environments. *PLOS ONE*, 8(11), e80753.
 77. Stearns, S. C. (1989). Trade-offs in life history evolution. *Functional Ecology*, 3, 259-268.
 78. Stearns, S. C. (1992). *The Evolution of Life Histories*. Oxford, UK: OUP.
 79. Stearns, S. C. & Crandall, R. E. (1981). Quantitative predictions of delayed maturity. *Evolution*, 35, 455-463.
 80. Striedter, G. F. (2005). *Principles of Brain Evolution*. Sunderland, MA: Sinauer.

81. Sznajder, B., Sabelis, M. W. & Egas, M. (2012). How adaptive learning affects evolution: Reviewing theory on the Baldwin effect. *Evolutionary Biology*, 39, 301-310.
82. Taylor, C. & Jefferson, D. (1993). Artificial life as a tool for biological inquiry. *Artificial Life*, 1, 1-13.
83. Thornton, A. & Raihani, N. J. (2008). The evolution of teaching. *Animal Behaviour*, 75, 1823-1836.
84. Tomasello, M., Davis-Dasilva, M., Camak, L. & Bard, K. (1987). Observational learning of tool-use by young chimpanzees. *Human evolution*, 2, 175-183.
85. Weber, B. H. & Depew, D. J. (2003). *Evolution and Learning: The Baldwin Effect Reconsidered*. Cambridge, MA: MIT Press.
86. Werker, J. F. & Tees, R. C. (2005). Speech perception as a window for understanding plasticity and commitment in language systems of the brain. *Developmental Psychobiology*, 46, 233-51.
87. Whiten, A. (2000). Primate culture and social learning. *Cognitive Science*, 24, 477-508.
88. Zentall, T. (2001). Imitation in animals: Evidence, function, and mechanisms. *Cybernetics and Systems*, 32, 53-96.

Figure 1. Initial simulation results when only the protection period is allowed to evolve, and the linear individual learned performance improvement with learning rate δ_i stops with probability $\rho\delta_i$ at a random performance level in the range $[0,100]$, for $\rho \in \{0, 0.02, 0.05, 0.1\}$. The mean performance levels exhibit the required ρ dependent peaks when plotted as a function of the expected learning time $100/\delta_i$, and the emergent protection periods depend only on the expected learning times.

Figure 2. Initial simulation results when the learning rate δ_i and protection period are both allowed to evolve. The mean evolved expected learning time (ELT) $100/\delta_i$ and protection period both increase with the learning task difficulty parameter ρ (left). The median learning performance increases with age differently for the evolved protection period (Ev.) and three fixed periods (1, 10, 20), for $\rho = 0.04$ and evolved δ_i (right).

Figure 3. Meme-based simulated evolution of the imitation rate α_i (left), and the corresponding change in average individual performance throughout evolution (right), for 20 runs of the baseline imitation-only models with $M = 500$ memes and fixed relatively small brain capacities $B_i = 100$.

Figure 4. Average numbers of good and bad memes known by individuals throughout evolution (left), and individual performance against age in a typical population after the evolution has stabilized (right), for 20 runs of the baseline imitation-only simulation with $M = 500$ memes and fixed relatively small brain capacities $B_i = 100$.

Figure 5. Average numbers of good and bad memes known by individuals throughout evolution, for imitation only of individuals' own parents (left), and imitation selection based only on performance-independent cultural values (right), for 20 runs of the imitation-only simulation with $M = 500$ memes and fixed relatively small brain capacities $B_i = 100$.

Figure 6. Evolution of average meme counts for imitation-only simulations with fixed brain capacities B large enough to accommodate a complete set of memes, for $B = M/2 = 250$ (top-left) and $B = M/2 = 100$ (top-right). Imitating high-performing rather than random individuals decreases the proportion of bad memes, shown for $B = M/2 = 250$ (bottom-left). Decreasing the meme injection rate increases the proportion of bad memes, shown for $B = M/2 = 100$ (bottom-right).

Figure 7. Typical imitation-only simulation in which the brain sizes B_i evolve alongside the imitation rates α_i , for $M = 200$ memes. A “mimetic transition” occurs after about 400,000 years, at which point the population average brain size B_i (top-left), imitation frequency $\alpha_i \phi B_i$ (top-right), performance (bottom-left) and meme counts (bottom-right) suddenly increase dramatically and then level off into a stable state.

Figure 8. Evolution of average individual performance (left) and meme counts (right) for copying fidelity $f = 0.99$ (i.e., 1% errors) with $M = 500$ memes and fixed relatively small brain capacities $B_i = 100$. There is a clear degradation in performance compared with the corresponding perfect copying case shown in Figures 3 and 4.

Figure 9. Evolution of the average meme counts for various levels of copying fidelity, when the copied individuals are chosen according to their performance, rather than randomly as for Figures 3, 4 and 8.

Figure 10. Variation in average evolved performances (left), imitation rates α_i (middle) and meme counts (right) as the copying fidelity worsens, for performance-based imitation as in Figure 9.

Figure 11. Evolution of average meme counts with performance-based imitation for three alternative forms of 1% copying errors: to a random worse meme ($x \rightarrow [-1, x]$), to a random bad meme ($x \rightarrow [-1, 0]$), and to a meme with performance contribution one unit less ($x \rightarrow x - 1$).

Figure 12. Evolution of average learning rate δ_i (left), and average numbers of good and bad memes known by individuals (right), for 20 runs of the meme-based simulation with direct learning only.

Figure 13. Evolution of average learning rate δ_i (left), and average protection period (right), for 20 runs of the meme-based simulations with direct learning only and evolved protection periods during which procreation is not allowed.

Figure 14. Average evolved direct learning rate δ_i and imitation rate α_i across the final populations for the cases of No Protection, Protection, and Protection with PWP, as a function of fixed copying fidelity f . Also shown are the corresponding evolved protection periods for the standard Protection case.

Figure 15. Average final population performances (left) and average final adult population performances (right) for the cases of No Protection, Protection, and Protection with PWP, as a function of the fixed copying fidelity f for the simulations of Figure 14. (For clarity, the graphs have different scales and the error bars represent standard errors rather than standard deviations.)