Gonzalez, M., Watson, R., & Bullock, S. (2017). Minimally sufficient conditions for the evolution of social learning and the emergence of non-genetic evolutionary systems. *Artificial Life*, 23(4), 493-517.
https://doi.org/10.1162/ARTL_a_00244

Peer reviewed version

Link to published version (if available):
[10.1162/ARTL_a_00244](https://doi.org/10.1162/ARTL_a_00244)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via MIT Press at https://www.mitpressjournals.org/doi/full/10.1162/ARTL_a_00244 . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Minimally sufficient conditions for the evolution of social learning and the emergence of non-genetic evolutionary systems

MIGUEL GONZALEZ^{*,**}, RICHARD WATSON^{**} AND SETH BULLOCK^{***}

^{*}Contact author. ^{**}Institute for Complex Systems Simulations, University of Southampton, UK SO17 1BJ

^{***}Department of Computer Science, University of Bristol, UK, BS8 1UB

E-mail: mgc1g11@soton.ac.uk (M.G.), r.a.watson@soton.ac.uk (R.W.), seth.bullock@bristol.ac.uk (S.B.)

Abstract

Social learning, defined as the imitation of behaviours performed by others, is recognised as a distinctive characteristic in humans and several other animal species. Previous work has claimed that the evolutionary fixation of social learning requires decision-making cognitive abilities that result in transmission bias (e.g., discriminatory imitation) and/or guided variation (e.g., adaptive modification of behaviours through individual learning). Here, we present and analyse a simple agent-based model which demonstrates that the transition from instinctive actuators (i.e., non-learning agents whose behaviour is hardcoded in their genes) to social learners (i.e., agents that imitate behaviours) can occur without invoking such decision-making abilities. The model shows that the social learning of a trait may evolve and fix in a population if there are many possible behavioural variants of the trait, if it is subject to strong selection pressure for survival (as distinct from reproduction), and if imitation errors occur at a higher rate than genetic mutation. These results demonstrate that the (sometimes implicit) assumption in prior work that decision-making abilities are required is incorrect, thus allowing a more parsimonious explanation for the evolution of social learning that applies to a wider range of organisms. Furthermore, we identify *genotype-phenotype disengagement* as a signal for the imminent fixation of social learners, and explain the way in which this disengagement leads to the emergence of a basic form of cultural evolution (i.e., a non-genetic evolutionary system).

Keywords: non-genetic evolution, animal culture, evolution of social learning, horizontal information transfer, imitation, survival selection.

1 The evolution of social learning

The evolution of social learning (i.e., the ability to acquire and replicate behavioural information expressed by others) has been identified as a pre-requisite for the emergence of culture [38, 31]. Research on the evolution of social learning generally investigates conditions where (1) social learners invade a population of individual learners (i.e., agents that can learn fit strategies at the expense of a fitness cost associated with trial-and-error learning), and (2) social learners increase the average fitness of a population [39, 15, 1, 5, 29]. In contrast, few models have investigated conditions where social learners can invade a monomorphic population of instinctive actuators (i.e., individuals whose behaviour is hardcoded in genes) [11, 6]. The rarity of these types of models reflects the consensus view that all forms of social learning in nature must have emerged in species where individual learning already existed [31, 38, 20]. In cases where instinctive actuators compete with social learners and individual learners, a mixed population of these learning strategies has no difficulty invading and entirely displacing instinctive actuators [1, 5, 6]. Furthermore, social learners alone are considered to be capable of displacing instinctive actuators as long as they have certain decision-making abilities [38, 15]. These abilities are grouped into two categories (1) transmission biases, where social learners bias their imitation towards certain (fitter) behaviours, and (2) guided variation, where imitated behaviours can be adaptively improved through a facultative form of individual learning [38].

Decision-making abilities allow social learners to accelerate the pace of cultural evolution far beyond that of relatively slow genetic evolution [32, 3]. With these abilities social learners can iterate, discriminate and selectively retain behaviours throughout a single lifetime, whereas these processes would require several generations of genetic evolution [38, 32]. However, it has been argued that if decision-making abilities and/or individual learners are not present, social learners are unable to invade and fixate in a population of instinctive actuators [11]. In 1983, Cavalli-Sforza and Feldman showed that any degree of horizontal (or oblique) imitation in social learners (i.e., imitating non-parental behaviours) translates into a fitness cost when compared to instinctive actuators and therefore concluded that unbiased (indiscriminate) imitation could not enable social learners to invade a population of instinctive actuators.

In this paper we challenge two fundamental ideas underlying prior work; (1) social learners require decision-making abilities (i.e., transmission bias and/or guided variation) in order to fix in a population of instinctive actuators [39, 15, 1, 5, 29, 38] and (2) unbiased social learners (i.e.,

indiscriminate social learners without decision-making abilities) can only invade a population of instinctive actuators if individual learners are also present [38, 11, 39, 1, 5]. In contrast to prior work, we investigate scenarios where: (1) the relative importance of selection for reproduction and selection for survival can be varied (also referred to as *fecundity* selection and *viability* selection, respectively), and (2) an extended strategy space can potentially be explored by either a genetic evolutionary system (in the form of instinctive actuators expressing genetically hard-wired behaviours that are subject to mutation) or a non-genetic evolutionary system (in the form of social learners expressing imitated behaviours that are subject to imitation error). Using simulation modelling, we show that when selection for survival (viability) is stronger than selection for reproduction (fecundity), and imitation error rates are higher than genetic mutation rates, social learners can irreversibly fix in a population of instinctive actuators. Furthermore, we identify a process termed genotype-phenotype disengagement that acts as a signal for the imminent fixation of social learners, and explain the way in which such a process leads to the emergence of a non-genetic evolutionary system.

Our work is framed as a theoretical demonstration that (1) the prior existence of individual learners is not required to explain the fixation of social learners (i.e., unbiased social learners can be systematically favoured by selection under the conditions we identify), and (2) decision-making abilities are likewise not necessary for social learners to invade a population of instinctive actuators.

The evolution of social learning in our model is explained by the relatively advanced exploration and exploitation capabilities of the invading non-genetic evolutionary system compared to those of the wild-type genetic system. Our model offers insight into the evolution of social learning and the emergence of non-genetic evolutionary systems in cases where individual learning is either prohibitively costly or where the behaviour to be learned is too complex for one individual to achieve. As discussed in the final section of this paper, our model can also be generalised to other examples of horizontal information transfer (i.e., other than social learning) and hence explain the emergence of alternative evolutionary systems (i.e., other than cultural evolution), for example; the evolution of behaviours that promote horizontal infection -over vertical infection- by symbiotic bacteria (e.g., through faecal matter consumption or close-range physical contact) and the consequent emergence of an evolutionary system which defines key traits of the host population but which information resides in the bacterial population.

Below we briefly discuss the relationship of social learning to culture before summarising prior

work on the two main ideas that are challenged in this paper and the assumptions that they depend on. In order to discuss these ideas we employ the concept of a *behavioural variant* (i.e., a specific behaviour obtained by social learning) [38]. Different behavioural variants may achieve different degrees of success at solving a challenge (ultimately translated into fitness scores). In this work we assume that the particular behavioural variant adopted by a social learner overrides its genetically encoded instinctive behaviour, and that different behavioural variants can be understood to be competing with each other for memory space in the nervous systems of social learners [38, 13].

1.1 Culture as a sophisticated non-genetic evolutionary system

Culture is a difficult concept to grasp, yet it is easy to recognise its importance to our species. We owe to it our extensive geographic distribution, and our capacity for environmental adaptation, innovation and civilisation, all of which have increased in magnitude and efficiency over a time period that is short enough to exclude genetic evolution as an explanation [8, 38]. The specific processes that gave rise to culture are still unclear, but logically the ability to imitate behaviours (i.e., social learning) must have evolved genetically before culture emerged [38].

Attempts to define culture are abundant in the literature [24, 19, 37, 39, 38]. Although many of these definitions encompass processes that we do not consider or include in our model, some view culture as a simple non-genetic evolutionary system maintained by social learning. More specifically, the theory of cultural evolution developed by Robert Boyd and Peter Richerson defines culture as a set of behavioural traits that are not the direct result of genetic expression but the product of an evolving pool of variants stored and imperfectly transmitted within and/or between overlapping generations by means of social learning [18, 2, 38].

To avoid confusion with other definitions of culture and frame our claims within the limitations of our simulation model we will use the generic term *non-genetic evolutionary system* to refer to all forms of evolutionary systems that emerge as a consequence of non-genetic information transfer between individuals. In nature, the existence of non-genetic evolutionary systems driven by social learning have been identified in many species [27, 26, 2]. Evidence of these systems is found in birds, where song learning occurs by means of imitation [21, 17]; chimpanzees, with well-known cases of tool use and material culture [30]; Cetaceans, imitating cooperative hunting practices and mating calls [36]; and others [25].

Human culture is a sophisticated example of a non-genetic evolutionary system [38]. The dis-

tinctions between human and animal non-genetic evolutionary systems are outside the scope of the current paper, but some schools of thought object to classifying non-human, socially learned animal behaviours as examples of culture [25]. There is little contention, however, that all these instances share a general mechanism during their emergence, most certainly involving some form of social learning [38, 31, 25].

1.2 Can unbiased social learners invade populations of instinctive actuators?

Previous models describing evolutionary dynamics in populations of social learners and instinctive actuators have concluded that social learners performing any degree of unbiased horizontal imitation (i.e., imitation of behaviours from non-parental individuals chosen at random in a well mixed population) could not invade and fix in a population of instinctive actuators [11, 38]. The dynamical system model proposed by Cavalli-Sforza and Feldman reaches this conclusion with a population that has two types of strategies, imitators and instinctive actuators; and two possible behaviours, one behaviour being fitter than the other. Their model keeps track of three types of individuals; fit imitators (imitators that happen to copy a fit behaviour from the population), unfit imitators (imitators that happen to copy the unfit behaviour) and fit instinctive actuators (performing the fit behaviour which they inherited genetically). Cavalli-Sforza and Feldman prove that given these three types, populations would always converge to the fixation of instinctive actuators [11].

The results of the Cavalli-Sforza and Feldman model can be explained by considering the effect of unbiased horizontal imitation (referred to as “oblique” transmission in Cavalli-Sforza and Feldman (1983) generational model; where social learners imitate non-parental members of the previous generation). In their model, the strategy of an instinctive actuator (“be instinctive”) and its behaviour (“be fit” or “be unfit”) are both transmitted vertically (from its parents). The strategy of an imitator (“imitate”) is transmitted vertically, but its behaviour (“be fit” or “be unfit”) is transmitted largely horizontally, since an imitator will pick a random member of the population as a model. When strategies are evaluated in terms of their reproductive fitness, horizontal imitation suffers a penalty due to the reduced heritability of imitated behaviours. Since the offspring of an imitator will imitate any agent’s behaviour at random, the heritability of an imitator parents behaviour is broken whenever their offspring imitates horizontally rather than vertically. As such, in a population that exhibits any unfit behaviour at all, a lineage of fit instinctive actuators will

outperform a lineage of imitators, since whereas the offspring of the instinctive actuators tend to inherit their parents' fitter than average behaviour, the offspring of an imitator is, behaviourally, anyone's offspring, and hence tends to have average fitness.

Consequently, the claim is that social learners require something more than unbiased imitation in order to fix in a population of instinctive actuators; they require another source of selection. To solve this issue several theoretical models have been developed, including examples with transmission biases and guided variation [15, 31, 1, 5], spatial structure (i.e., where agents interact in an explicit space) [9], and population structure [35]. In this work, models in which social learners can discern between variants are considered transmission bias models. With this ability, social learners can imitate fitter variants more often, or alternatively, change their behavioural strategies by comparing the observed behaviour of several individuals [15]. Guided variation models assume imitators optimise behavioural traits after acquiring them [38, 1, 5], a trait which itself implies some degree of individual learning. Spatial and population structure models add assumptions that create constrained scenarios in which fitter individuals are copied more often than average [9, 35]. Here, we claim that none of these additions are required and that therefore the assumption that social learners must be endowed with decision-making abilities is unnecessary.

Most existing work is based on simple analytical models where (1) the effect of natural selection is only represented by differential reproduction (i.e., without considering its effects on survival), and/or (2) the strategy space is defined by a two-state environment-matching problem. When either or both of these assumptions are made, unbiased social learners (i.e., social learners with no decision making abilities) cannot invade a population of instinctive actuators [11, 29] and will only partially invade a population of individual learners until a mixed equilibrium is achieved [39].

In contrast, our model shows conditions under which unbiased social learners can fix in a population when (1) natural selection is separated into survival and reproductive components and (2) an extended strategy space (i.e., fitness landscape with many more than two genotypes) is explored in parallel by both the genetic and non-genetic evolutionary systems, subject to their respective mutation and imitation error rates. In this model, the non-genetic evolutionary system emerges when social learners with unbiased imitation successfully invade the population.

Previous models with two-state strategy spaces have sometimes included rates of mutation and imitation error [39, 38]. However, in a two-state model, these rates merely act as the probability of defective transmission, whereas in our extended strategy space they affect the balance of exploration

and exploitation in the evolutionary system. Therefore, although some two-state models have included mutation/imitation error rates, we claim they play a different role in our model.

1.3 The effect of individual learners

When considering the question of how social learning evolved, most evolutionary theorists have focused their research on finding conditions under which social learners can invade a population of individual learners [8, 4, 39, 38, 15, 31, 1]. In these models individual learners tend to discover a fitter-than-average behaviour for the current environmental state, but pay a fitness cost in learning this behaviour. This cost is generally associated with the process of trial-and-error [38, 31]. In these models the environment changes with a given frequency but individual learners always match their behaviour to the current environmental state [8, 4, 39, 38]. In a population of individual learners, a mutant social learner with unbiased imitation can easily increase in frequency because the social learner is spared the trial-and-error cost of individual learning, but can acquire, through imitation, the fit behaviour produced by it [39, 8, 4]. However, as the frequency of social learners increases so does the chance of imitating a behaviour that does not match the current environmental state (i.e., a population of only social learners cannot keep track of environmental changes). Therefore, the frequency of social learners will only increase to an equilibrium point where both individual and social learners have the same fitness value (i.e., a mixed population) [38, 39]. The fitness of individual learners is not density dependent. Hence, at equilibrium the mixed population has an overall fitness value equal to that of a monomorphic population of individual learners [39]. According to these models, decision-making attributes are thus required to explain the invasion of social learners beyond this mixed equilibrium (e.g., in extremis, displacing all but one individual learners) and an increase in the overall mean fitness of the population compared to a monomorphic population of individual learners [38, 39, 15, 31, 1].

In this prior work, social learners can invade (but not fix) because the presence of individual learners confers a benefit on social learners compared to instinctive actuators; social learners have access to the fitter-than-average behaviours found by individual learners but instinctive actuators do not. Therefore, the underlying source of fitter-than-average variants that allow social learners to invade when individual learners are present is similar to that found in models where social learners have decision-making abilities that improve on imitated behaviours (i.e., guided variation). In both cases the valuable information found by social learning is ultimately derived from some

form of individual learning (i.e., the individual ability to optimize behavioural fitness). In our model (to be explained in the next section), social learners fix due to a purely Darwinian process where behavioural variants (generated without bias) are the units of selection. In contrast, we can view previous models as Lamarckian adaptation processes, necessitating mechanisms of directed improvement acting on innate or acquired variants, in those cases via individual learning or guided variation.

Individual learners exist in current natural populations alongside social learners, and may have been present prior to the evolution of social learners [2]. But, as a theoretical exercise, we exclude individual learning from the following models in order to test the hypothesis that it is not required to explain the fixation of social learning and the emergence of a corresponding non-genetic evolutionary system. Accordingly, the model presented here also enables the investigation of mechanisms for the emergence of non-genetic/parallel evolutionary systems through the evolution of unbiased horizontal information transfer where individual learning may not be relevant (e.g., the emergence of an evolutionary system maintained by symbiotic bacteria which are transmitted horizontally in a population of hosts; see Subsection 4.4 in the discussion section).

2 The Model

Our individual-based simulation model represents each individual by two strings of bits. The first string represents an individual’s phenotype and the second its genotype. The fitness of an individual is determined by considering only the phenotypic information. In addition, each individual has a genetically inherited one-bit switch that determines whether it is an imitator or an instinctive actuator. If the value of this switch is 1, the individual is a social learning imitator and, rather than express a phenotype derived from its own genotype, it will express an “imitated” phenotype randomly selected from the population. If the value of the switch is 0, the individual is an instinctive actuator and its phenotype is simply a copy of its own genotype (for the purposes of our research questions, a one-to-one genotype-phenotype map is sufficient). For each imitator, the action of imitation takes place at birth and once it has happened an individual’s phenotype remains unchanged for its lifetime.

Each agent in the initial population is an instinctive actuator with an imitation switch value of 0 and each bit of its genotype set to 1 or 0 with equal probability. The phenotype of each agent in the initial generation is a perfect copy of its genotype. These initial conditions prevent mutation biases from imposing directed drift on the genotypes or phenotypes, even in the absence of selection pressure. For all the results shown here, unless stated otherwise, the length of the genotype and phenotype bit strings is 200 bits ($L = 200$), and the size of the population is 100 individuals ($N = 100$). Selection is defined by a Boltzmann-weighted sum of ones in the string of phenotypic bits (L_1). A string of all-1s (i.e., $L_1 = L$) represents the optimal solution.

Qualitatively similar results can be reproduced with larger populations and longer bit strings. Smaller populations can result in variation between simulation runs, but on average they also agree qualitatively with the results described here. Shorter bit-strings can produce different results, as they represent simpler search problems, and a sustained search and optimisation period of several generations is required in order to produce the results reported here. For example, the commonly used case of a single-bit (two-strategy) model will not reproduce our results. When only two-strategies are explored the instinctive actuators are likely to find the fittest of the two strategies quickly and hence prevent the key process of genotype-phenotype disengagement which we will explain in the next section.

Selection is established by the joint action of a reproduction function and a death function, implementing reproductive selection and survival selection, respectively. Reproduction selects an

individual, i , from the population with a probability P_{ri} using a Boltzmann-weighted function of the number of 1s in the individuals phenotype (i.e., L_{1i}) normalised across the population (Eq. 1).

$$P_{ri} = \frac{e^{\frac{L_{1i}}{x_r}}}{\sum_{k=1}^N e^{\frac{L_{1k}}{x_r}}} \quad (1)$$

The death function selects an individual, j , from the population with a probability P_{dj} using a Boltzmann-weighted function of the number of 0s in the individuals phenotype (i.e., $L_{0j} = L - L_{1j}$). Thus, P_{dj} is the relative probability of dying, or anti-fitness (Eq. 2).

$$P_{dj} = \frac{e^{\frac{L - L_{1j}}{x_d}}}{\sum_{k=1}^N e^{\frac{L - L_{1k}}{x_d}}} \quad (2)$$

A key feature of this model is the ability to change the relative balance between reproductive selection and survival selection. Higher values for exponents x_r and x_d reduce the effect of the number of 1s in an individuals phenotype on the probabilities of reproduction and death, respectively. For lower values of x_r or x_d , small changes in behaviour have a larger impact on reproductive output and life expectancy, respectively.

On each iteration of the simulation, an individual is selected for reproduction by the reproduction function and an individual is selected for death by the death function. The genotype and imitation switch of the reproducing individual are copied and mutated, and replace the genotype and imitation switch of the dying individual. Genotype mutations are bit-flips, and occur with probability of μ_g per bit, except for the imitation switch which has a bit-flip probability of μ_c . If the imitation bit of the new individual has value 0, the new phenotype bit string is a one-to-one copy of the new genotype (not including the imitation switch). If the switch has value 1, a random phenotype from the population is copied (unbiased horizontal imitation) with a bit-flip error rate of μ_p per bit. The imitation switch can only be passed genetically and is not part of any phenotype. This separates the genetically inherited strategy of being an imitator or not (i.e., the imitation switch) from the string of phenotypic bits that express an individuals actual behavioural variant.

For simplicity, individuals reproduce asexually with no “crossover” during genetic replication or phenotypic imitation and there is no environmental change altering the optimum bit-string ($L_1 = L$), although a brief comment on the effect of the latter is presented in the discussion.

Changing the values of μ_g and μ_p controls rates of genetic mutation and imitation error, respectively. We fix $\mu_c = 0.01$ for all results reported here and vary μ_g and μ_p relative to this value. Figure 1 illustrates an algorithmic implementation of the model.

Figure 1: An algorithmic representation of the model's logic. (1) The selection function picks an individual from the population. (2) Its genotype string is copied including the imitation switch (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_p per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy of the individual's genotype. (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function.

3 Results

Simulations were run for a large range of mutation rates $[0.0, 0.015]$, imitation error rates $[0.0, 0.015]$, survival selection strengths $[1, 20]$, and reproductive selection strengths $[1, 20]$. Each simulation consisted of 5×10^4 iterations. The average proportion of 1s in the population’s genotypes, the average proportion of 1s in the population’s phenotypes, and the proportion of individuals with their imitation switch “on” (i.e., the proportion of social learners) were recorded. All simulation runs have a population size of $N = 100$ and a string length for phenotype and genotype strings of $L = 200$.

Where imitation error rate is larger than genetic mutation rate, $\mu_p > \mu_g$, simulations converge to one of three distinctive patterns (Figure 2): A) Both genotypes and phenotypes converge on the single optimum, while the proportion of imitators remains close to zero; B) Phenotypes and genotypes initially improve together, but after around two thousand iterations, imitators rapidly invade the population and only the phenotypes continue to improve to the optimum; C) imitators increase in frequency slowly and inconsistently (i.e., as if the proportion of imitators in the population was drifting), and eventually fix in the population; after fixation only the phenotypes continue to improve. In B and C, genotypes tend to converge to the average, containing approximately equal numbers of 1s and 0s. The de-correlation of genotype and phenotype information, accompanied by the fixation of imitators, is here referred to as *genotype-phenotype disengagement*. This pattern is considered to indicate the emergence of a non-genetic evolutionary system. The stochastic nature of the model makes convergence probabilistic. Multiple replicates are hence required to analyse outcomes for each parameterisation. The chances of a given run of the model converging to pattern A, B or C is affected by the strength of reproductive and survival selection (x_r and x_d , respectively - recall that a high value indicates lower selection strength), as well as the mutation rate and imitation error rate (μ_g and μ_p , respectively).

Scenarios in which the genetic mutation rate is larger than the imitation error rate, $\mu_g > \mu_p$, are not considered relevant to the evolution of non-genetic evolutionary systems in this work, since it is unrealistic that the incipient secondary system has more fidelity than the established process of genetic replication [13, 2, 38, 31].

Figure 2: A single representative simulation run for each of three convergence patterns, A, B and C. Each plot depicts the change over evolutionary time in the average proportion of 1s in the population's phenotypes (green) and genotypes (red), and the proportion of imitators in the population (black). (A) No disengagement: phenotype and genotype improve together and imitation frequency remains low. This case is encountered when genetic mutation is high and the imitation error rate offers not enough exploration advantage for a given set of reproductive and survival selection strengths (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.003$, $\mu_p = 0.005$). (B) Early disengagement: the non-genetic evolutionary system emerges rapidly in the simulation, indicated by a sharp rise in imitators and genotype-phenotype disengagement. Imitators fix and subsequent phenotype improvement is due to non-genetic evolution. (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0005$, $\mu_p = 0.005$). (C) Late disengagement: non-genetic evolution takes over after several iterations. Imitators invade slowly and with large fluctuations in frequency. This pattern occurs when mutation and imitation error rates are very low and/or similar to one another. (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0002$, $\mu_p = 0.0002$).

One hundred replicates for each combination of mutation rate, μ_g , and imitation error rate, μ_p , drawn from the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \dots, 50 \times 10^{-4}]$, and where $\mu_p > \mu_g$, were carried out with fixed values of survival and reproductive selection coefficients ($x_d = 1$, $x_r = 10$). The proportion of

replicates that converge to non-genetic evolution (where the proportion of imitators is above 0.95 by the end of the run) is presented in a two-dimensional heat-map in Figure 3. This map shows all combinations below the diagonal $\mu_p = \mu_g$, where imitation error is greater than genetic mutation.

Figure 3: Heat-map depicting how the tendency for social learners to evolve varies with the genetic mutation rate, μ_g , and imitation error rate, μ_p . For all points in this plot $\mu_p > \mu_g$. Blue regions indicate that the explorative advantage of imitation errors is not great enough for disengagement to occur early on during the optimisation process. The lower region shows cases where social learners invade and fix in the population. In all these cases the values for the selection exponents are set to $x_d = 1$ and $x_r = 10$ creating a selective environment with strong survival selection and relatively weak reproductive selection. For μ values larger than 0.005, the trends observed at the positive edge of this heat-map extrapolate within sensible limits. The constrained range has been selected to increase resolution on contrasting regions. Black arrows indicate points in parameter-space corresponding to the single runs shown in figure 2 (i.e., A, B and C).

Exploration of a range of survival and reproductive selection strengths is presented in the array of heat maps in Figure 4; each having the format explained in the previous paragraph. Combinations of selection coefficient values of 1, 5 and 10 for x_d and x_r were selected for their contrasting

results. Figure 4 can be interpreted as a four-dimensional parameter-space representation showing the distribution of probabilities for the fixation of social learners. For relatively strong survival selection, and relatively high imitation error rates, social learners fix following pattern B in Figure 2. For mutation and imitation error rates that are low and similar, social learners fix following pattern C in Figure 2. Given that our model does not feature decision-making behaviours, and that both error rates and survival selection pressure on the trait being evolved are not functions of previously acquired traits, we consider these results to demonstrate the evolution (and fixation) of unbiased social learning by virtue of its adaptive value alone.

Figure 4: Array of heat maps for different combinations of reproductive and survival selection coefficients. Within each plot, red areas represent parameter sets for which social learners fix through genotype-phenotype disengagement in a high proportion of simulation replicates. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same range of mutation and imitation error rates (i.e., where $0.005 \geq \mu_p > \mu_g \geq 0$).

Our model highlights two important requirements for the emergence of a non-genetic evolutionary system from unbiased social learning. First, a strong selection pressure on survival relative to selection pressure on reproduction, and, second, an extended strategy space combined with the

presence of mutation and imitation error rates. The culture-enhancing effects of these conditions have been discussed in existing literature [2, 38, 31]. However, their sufficiency for the origin of culture as a non-genetic evolutionary system has not been discussed.

An extensive strategy space resulting in a potential for sustained evolutionary competition between instinctive actuators and social learners is a key feature of our agent-based approach that is not present in classical models [15, 31, 9, 35, 39]. Understanding the effect of this feature demands analysis of individual simulation runs. In the next subsection we explain the underlying principles that makes social learners more likely to invade a population of instinctive actuators when survival selection is stronger than reproductive selection. The following two subsections describe two specific mechanisms for the fixation of social learners and the emergence of a non-genetic evolutionary system in our model; these correspond to scenarios B and C in Figure 2. Both mechanisms co-occur in all simulations where social learners fix. However, as will be explained, the mechanism that drives scenario B is dominant for simulations where $\mu_p \gg \mu_g$, and the one that drives scenario C is dominant for simulations where $\mu_p \simeq \mu_g$. In the final subsection of our results, we show how the length of the problem sequence (i.e., the length of the agents phenotype and genotype bit strings) affects the likelihood of social learning evolving in a population.

3.1 Imitation error and survival selection can compensate for lack of reproductive selection pressure on social learners

The results shown in Figure 4 indicate that populations experiencing strong survival selection (i.e., x_d is relatively low) and weak reproductive selection (i.e., x_r is relatively high) tend to converge on social learning for a wider range of mutation and imitation error rates (i.e., μ_g and μ_p , respectively). In order to account for the distribution of results, we first explain here why relatively strong survival selection benefits social learners.

First, consider that selection for reproduction alone cannot favour unbiased social learners over instinctive actuators. In a population of instinctive actuators, selection ensures that the phenotype expressed by a newly added offspring tends to be fitter than the current population mean. By contrast, in a population of unbiased social learners the offspring of the fittest individuals (i.e., those that happen to have imitated the fittest behaviours) will imitate phenotypes chosen at random from the current population; hence the phenotypes expressed by newly added social learner offspring will tend to have the same fitness as the current population average. For this reason, it is generally

held that unbiased social learning causes a “regression to the mean” when compared to genetic reproduction [38, 11]. However, note that instinctive actuators will also tend to suffer from this regression to the mean to the extent that selection for reproduction is weak, since this is the extent to which parents will tend to be selected at random.

In contrast to selection for reproduction, selection for survival ensures that novel phenotypes with higher fitness than the population average are maintained for longer, regardless of whether instinctive actuators or imitators express them. By virtue of their longer than average persistence in the population such individuals have more opportunities to reproduce, but also have more chance of being imitated by unbiased social learners.

Instinctive actuators inherit their phenotype from their parents. Therefore, long-living instinctive actuators leave more copies of their phenotype than short-living instinctive actuators. Social learners, on the other hand, do not inherit the phenotype of their parents but imitate the phenotype of a randomly chosen member of the population. Nevertheless, long-living agents are more likely to be randomly selected for imitation before they are eliminated by the death function. Therefore, social learners also leave a number of phenotypic copies proportional to their relative longevity.

In our model, selection for survival affects instinctive actuators and social learners to the same extent, while selection for reproduction affects instinctive actuators exclusively.

In the absence of selection for reproduction, the only process mitigating against the regression to the mean for social learners is selection for survival and imitation error rates. When a population’s phenotypic quality is high, imitation errors are likely to be maladaptive, but when a population’s phenotypic quality is low, errors are more likely to cause imitators to express phenotypes with improved fitness. If the rate of imitation error is higher than that of genetic mutation, and phenotypic quality in the population is low, then social learners will tend to discover improved phenotypes at a higher rate than instinctive actuators.

Consequently, the invasion and fixation of social learners depends on whether or not they find fitter solutions than instinctive actuators through imitation errors. For certain combinations of selection pressure and error rate parameters, the exploration advantage of social learners (due to high imitation error and strong selection for survival) more than compensates for the slightly stronger regression to the mean that they suffer relative to instinctive actuators (which have additional “weak” selection for reproduction acting only on them). Under such conditions, our results show that social learners can invade populations of instinctive actuators. As we will see in the next subsection,

this invasion must occur during the early stages of evolutionary optimisation, as higher error rates become disadvantageous once the phenotypic quality of instinctive actuators has exceeded a certain threshold. However, if and when an invasion of social learners occurs, it is unlikely that instinctive actuators can ever recover.

3.2 When imitation error is much higher than genetic mutation, social learners fix during early stages of optimisation

In order to explain how this process works, we define two periods that tend to occur during a simulation. Figure 5 shows these periods during the first 10^4 iterations of a simulation run where social learners fix as per scenario B of Figure 2.

Figure 5: The first ten thousand iterations of a single run, with two distinct periods highlighted. The green dashed line represents the average proportion of 1s in the phenotypes, the red triangles do the same for the genotypes and the black line represents the proportion of imitators in the population. The first period is shaded in blue and extends from the beginning of the simulation to around iteration 3700 when the second period begins (shaded in red). During the first period, the exploration advantage of imitators must overcome the penalty arising from the lower heritability of their imitated behaviour or imitation will not evolve. If imitators do not achieve a high enough frequency during the first period, genetic evolution will have time to improve phenotypes to the extent that the high imitation error rate associated with social learning becomes disadvantageous, halting their invasion. The hypothetical limit α represents a threshold degree of phenotypic quality, above which the high imitation error rate of social learners produces more deleterious than beneficial errors. Consequently, if instinctive actuators achieve this threshold before social learners achieve genotype-phenotype disengagement, the non-genetic evolutionary system will not emerge. (Parameter Set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0005$, $\mu_p = 0.005$).

The first period runs from the start of the simulation, when the population comprises only instinctive actuators that map their own genotype into their phenotype and reproduce by replicating their genotype with a mutation rate μ_g . Individuals are selected to reproduce on the basis of their

phenotype, and also selected to die on the basis of the same bit string. With a probability μ_c , a new agent will mutate its inherited imitation switch from 0 to 1, and hence obtain its phenotype by imitation with error rate μ_p per bit. During this period two processes occur in parallel. First, a marginal number of imitators arise through mutation drift alone. Second, these imitators start exploring the strategy space with error rate μ_p , imitating strategies from other imitators and instinctive actuators alike. Under survival selection, these unbiased imitators are more likely to copy strategies that stick around for longer than strategies that are quickly selected out by the death function.

In the first iterations of this period, instinctive actuators are a large majority and their phenotypes improve due to selective pressure from both reproductive and survival selection. They pass their phenotype to their offspring via vertical genetic inheritance with high fidelity ($\mu_g < \mu_p$). By contrast, imitator phenotypes improve only in response to selective pressure for survival selection. This is because phenotypes expressed by imitators cannot increase their copy numbers by means of reproductive selection as their offspring do not inherit the parental phenotype, but instead imitate any model at random (i.e., there is no assumption of discriminatory, guided, or adaptive imitation in this model). By contrast, survival selection still affects imitators, as having a better phenotype reduces the chance of being selected by the death function, and therefore leads to a longer life. Longer-living agents have a higher chance of being copied by another imitator.

At this stage, imitation has an exploratory advantage, since imitation error is higher than genetic mutation (Figure 6). But social learners can only fix if this exploratory advantage overcomes the selection handicap that imitators suffer from. This is more likely to occur if the trait that we are considering has little effect on fertility but a large positive effect on longevity (i.e., $x_r \gg x_d$), as positive effects on fertility will increase the optimisation rate (through higher selective pressure) of the genetic system (maintained by instinctive actuators) compared to the non-genetic system (maintained by social learners).

Figure 6: Comparison between a simulation run where social learning is evolvable (plot A), and a simulation run where social learning is not evolvable (plot B). In plot A and B the green dotted line represents the proportion of ones in the phenotypes of the population and the red triangles the proportion of ones in the genotypes. The black line in plot A represents the frequency of social learners. For both runs $\mu_g = 0.0005$ and for run A $\mu_p = 0.005$. Selection strength coefficients are set to $x_r = 10$ and $x_d = 1$. Plot C shows the variance of the sum of 1s values for the phenotypes in the populations of each run. Run A has a consistently higher phenotypic variance than run B. This is explained by the fixation of social learners in run A, and the consequent emergence of the non-genetic evolutionary system in which phenotypes are imitated with an error rate higher than genetic mutation. The variance in simulation run A exhibits large spikes that arise when the offspring of a social learner mutates into an instinctive actuator and expresses a phenotype derived from its own genotype (which has not been subject to selection for as long as its ancestors have been imitators). These mutants are selected against, causing variance to recover rapidly.

For social learning to evolve and genotype-phenotype disengagement to take place as per scenario B in Figure 2, the imitator minority has to find better solutions at a rate that overcomes its disadvantage from having only one source of selective pressure; survival. For all cases where $\mu_p \gg \mu_g$, this needs to occur quickly. Once genetic evolution, guided by instinctive actuators, improves

the phenotype beyond a certain number of 1s (i.e., α value in Figure 5), the likelihood of imitators taking over the population drops to zero. At this point the average phenotype solution starts benefiting from lower mutation rates due to its proximity to the optimum, and the exploration advantage of social learners turns into a disadvantage. In Figure 3 and 4 social learners do not fix in simulations where μ_g is above a certain value precisely because the population evolves phenotypes beyond this threshold of number of 1s before social learners can rise in frequency. As a result of this process, the area with a high proportion of simulations converging to social learners is limited to a horizontal band at the bottom of the plot.

The second period, shaded in red in Figure 5, starts from the point at which the higher rate of imitation error turns into a disadvantage compared to the mutation rate of genetic replication. If fixation of imitators has already been achieved during the first period, then the genotype has been disengaged and is no longer subject to selection. In this case, a non-genetic evolutionary system has emerged, and only the phenotype will continue to improve with imitation error as its source of variation and survival selection pushing it towards the optimum. This non-genetic evolutionary system manifests as a “pool” of variants maintained by unbiased imitation alone.

An alternative way to understand this mechanism and the importance of the relative difference between imitation error rate and genetic mutation rate is by imagining two hypothetical populations under the same degree of survival selection and in the absence of reproductive selection; one population of unbiased imitators ($P_{imitate}$) and one population of instinctive actuators ($P_{inherit}$).

When both populations are initialised with random genotypes/phenotypes (i.e., populated on average with 50% 1s and 50% 0s), the accumulation of fit alleles in $P_{inherit}$ is limited by the mutation rate, while the accumulation of fit alleles in $P_{imitate}$ is limited by the imitation error rate. If these two rates are equal, while imitators in $P_{imitate}$ can copy any strategy currently in the population, this does not give them an advantage or disadvantage compared with $P_{inherit}$, since the strategies available to be copied in $P_{imitate}$ are no more diverse than the ones that are available to reproduce in $P_{inherit}$.

If the imitation error rate in $P_{imitate}$ were higher than the genetic mutation rate in $P_{inherit}$, then $P_{imitate}$ would initially have an advantage and would accumulate fit alleles at a higher rate than $P_{inherit}$. If, conversely, the mutation rate in $P_{inherit}$ were higher than the error rate in $P_{imitate}$ (an unrealistic assumption), then $P_{inherit}$ would have an advantage and would accumulate fit alleles faster than $P_{imitate}$.

However, when $P_{imitate}$ and $P_{inherit}$ are both initialised with very fit solutions, the population suffering the highest error rate would be disadvantaged since errors would tend to degrade the fit solutions more often. At this point, if imitation error rate were lower than mutation rate it would always pay $P_{inherit}$ to swap to the higher fidelity mechanism. But for $P_{imitate}$, even if mutation rate were much lower than imitation error rate, it would only benefit them to become instinctive actuators if their genotypes were as fit as their phenotypes. But during the time that $P_{imitate}$ have been imitators, their genotypes have been under no selection pressure, and will therefore have drifted towards the mean, which means that $P_{imitate}$ is unlikely to give up imitation even when its high error rate is counter-productive. This asymmetry creates a “ratchet” in the system which ensures that even when the conditions that allowed social learners to successfully invade instinctive actuators have changed, instinctive actuators cannot recover.

In Figure 7 we test this explanation by comparing runs that are initialised with average quality phenotypes with runs that are initialised with higher quality phenotypes. As expected, higher initial phenotypic quality reduces the region of the parameter space associated with convergence to social learning, but this region is not extinguished entirely. Two areas in which social learning still evolves can be identified: 1) where mutation rate is very low; 2) where the difference between mutation rate and imitation error rate is low.

The first case is associated with the phenotype-genotype disengagement account offered above. In extremis, when the genetic mutation rate is zero, social learning is the only source of phenotypic variation and social learners are expected to invade even when their high imitation error rate tends to be disadvantageous for cases where μ_g is very low (but not zero). If, under these conditions, the optimal phenotype were discovered through social learning, the optimal individual would not be selected to switch to the perfect fidelity of genetic inheritance since genotype-phenotype disengagement would have ensured its genotype would be of lower quality than existing phenotypes.

However, the second case (associated with the lower left-hand corner and lower portion of the diagonal of heat map B) is due to a second mechanism that will be discussed in the next subsection.

Figure 7: Heat maps depicting the influence of initial phenotypic quality on the evolution of social learning. In heat map A the initial population of instinctive actuators have average phenotypic quality (i.e., on average 50% of phenotypic bits are in position 1). In heat map B the population is initialised with higher quality phenotypes (i.e., on average 87.5% of phenotypic bits are in position 1). Red areas indicate a high proportion of simulations converging to social learning (i.e., where social learners fix). Simulations with low genetic mutation rates and relatively high imitation error rates (i.e., areas indicated by black arrows) converge to social learning in plot A but not in plot B. In both maps the selection exponents are set to $x_d = 1$ and $x_r = 10$. Increasing the initial phenotypic quality further does not qualitatively alter the results shown here.

As further proof of our rationale and to explain the horizontal (i.e., independent of μ_p) transition at $\mu_g = 0.001$ in our original heat map results (see Figure 3), we present Figure 8, where our standard simulation results in map A are compared against results for simulations where the population is initialised with a 0.5 proportion of imitators (i.e., 50% of the initial population have their imitation switch in position 1) and each individual phenotype is initialised with exactly half of their bits in position 1. These modifications produce two complementary effects that eliminate the relative advantage of instinctive actuators and the genetic evolutionary system over imitators and the non-genetic evolutionary system at the start of the simulation run according to our explanation (i.e., it levels the field between both evolutionary systems, at initial conditions); 1) they increase the probability for social learners to form chains of imitation (i.e., consecutive events of imitators imitating other imitators) from the beginning of the simulation (i.e., half of the population is already a social learner) and 2) they decrease the probability that instinctive actuators evolve phenotypes with sum of 1s values above the discussed α threshold in the first couple of hundred

iterations (i.e., as they would do by evolving the fittest individual phenotype from the initial high-variance population for simulations in map A; where each phenotype bit is initialised with equal chance of being 1 or 0).

In accordance with our explanation, large regions of the parameter space converge to social learning in map B, including areas where μ_g values are above 0.001 (i.e., for simulations that do not converge to social learning in map A). Moreover, as we use larger μ_p values (i.e., large imitation error rates), social learners are able to fix in simulations with increasing values of μ_g (notice the inclined boundary between red and blue regions in map B in contrast to the horizontal boundary at $\mu_g = 0.001$ in map A). Such a pattern is the result of the non-genetic evolutionary system having an equal competitive advantage at initial conditions, and therefore imitators being more likely to find and fix fitter behavioural variants in the early stages of optimisation. As imitation error is set to higher values, the non-genetic evolutionary system maintained by social learners can outcompete genetic systems with higher mutation rates (but never higher than imitation error rates) as per scenario B in Figure 2. Once social learners fix, and the non-genetic evolutionary system is instantiated, the process of genotype-phenotype disengagement hampers latter reinvasions by instinctive actuators.

Figure 8: Comparison between plot A, showing our standard simulation results where populations are initialised with all individuals as instinctive actuators and phenotype strings bits set to 1 or 0 with equal probability; and map B, where simulations are initialised with half the population as social learners and all individual phenotypes with exactly half of their bits as 1s. In map B social learners fix in areas where the imitation error rate is high, and the mutation rate is above 0.001. For increasing imitation error rate values (μ_p) social learners fix for higher values of mutation rate μ_g in map B. Contrasting areas are indicated with black arrows. In both maps the selection exponents are set to $x_d = 1$ and $x_r = 10$.

3.3 Social learners can also fix due to drift and the irreversibility of genotype-phenotype disengagement

In our model, when social learners fix, this fixation is irreversible. Once imitators have invaded the population and caused genotype-phenotype disengagement, genotypes cannot re-engage in genetic evolutionary adaptation, because the information contained in these bit strings drifts at random and is no longer subject to selective pressure (i.e., since it is not expressed in phenotypes). Therefore, once the population starts evolving phenotypes via the non-genetic evolutionary system, it is unlikely to go back to genetic evolution. In scenario C in Figure 2 this simple process alone (i.e., with no significant assistance from advantageous high error rates) can fix social learners in the long term.

Scenario C in Figure 2 occurs under conditions where both genetic mutation (μ_g) and imitation error (μ_p) are very low. In these simulations, tens of iterations can pass between mutation (or imitation error) events that introduce fitter variants in the population. Consequently, not only do social learners have a large window of opportunity to fix (as per the mechanism explained in

the previous subsection), but also the number of imitators can slowly increase as a result of an asymmetry in the effect of imitation switch mutations.

To understand this mechanism, we must remember that since social learner genotypes are not expressed they are shielded from selection and therefore tend to accumulate deleterious mutations. Consequently, the offspring of a long line of imitators is likely to be less fit if it mutates into an instinctive actuator, whereas the same is not true if the offspring of an instinctive actuator mutates into a social learner. Due to the asymmetry between these two mutation events, social learners can fix without any significant explorative advantage.

In Figure 9 we prove our rationale by comparing three different heat maps for; (A) our original simulation model, (B) simulations where phenotypes are initialised with a large proportion of 1s (i.e., 0.875 rather than 0.5), and (C) Lamarckian simulations in which the genotype of a social learner is updated to be a perfect copy of the phenotype that it achieved through social learning. Comparisons between map (A) and (B) (i.e., same plots from Figure 7) serve to highlight areas where a large proportion of simulations still converge to social learning despite phenotypes having passed the critical threshold discussed in the previous section. The remaining red areas in plot (B) correspond to cases where; 1) scenario C in Figure 2 occurs ($\mu_g \simeq \mu_p$) and 2) genetic mutation is non-existent ($\mu_g = 0$).

Forcing social learners to keep their genotype string as a copy of their phenotype prevents genotype-phenotype disengagement in map C of Figure 9. Under such a condition social learners do not fix for any area in the map. In these simulations genotype and phenotype effectively become a single string. Social learners are however still distinct from instinctive actuators, as the former obtains its phenotype through horizontal transmission (i.e., non-parental imitation), whereas the latter does it through vertical genetic inheritance.

In plot E of Figure 9, we show a single representative run where the prevention of genotype-phenotype disengagement (i.e., by keeping phenotypes and genotypes identical within individuals) makes the invasion of social learners reversible (compare to plot D in Figure 9). In such a situation it is unlikely that individual simulations would converge to the fixation of social learners (i.e., would have a final frequency of social learners that is higher than 0.95 after 5×10^4 iterations), instead we observe a typical “drifting” pattern where social learners fluctuate in frequency throughout the simulation run.

The results shown in map C and plot E of Figure 9 serve to illustrate our rationale with

regards to the fixation of social learners through the slow drifting mechanism explained in this subsection. This mechanism depends on a fundamental asymmetry between social learners and instinctive actuators, where the former is likely to irreversibly fix in a population even when having no significant advantage in comparison to the latter.

Figure 9: Heat maps depicting (A) baseline results from the standard version of the model, and (B) the influence on the evolution of social learning of increased initial phenotypic quality and (C) the effect of preventing genotype-phenotype disengagement. In heat maps A and C the initial population of instinctive actuators have average phenotypic quality (i.e., on average 50% of phenotypic bits are in position 1). In heat map B the population is initialised with higher quality phenotypes (i.e., on average 87.5% of phenotypic bits are in position 1). In heat map C genotype-phenotype disengagement is prevented by setting the genotype of a social learner to be a copy of the phenotype that it obtained through imitation. Red areas indicate that a high proportion of simulations converge to social learning. Simulations with low genetic mutation rates and relatively high imitation error rates (i.e., areas indicated by black arrows) converge to social learning in plot A but not in plot B. Simulations with very low mutation rate or low mutation and imitation error rates converge to social learning in plot B but not in plot C. Plots D and E show individual simulation runs where $\mu_g = \mu_p = 2 \times 10^{-4}$ from maps A and C, respectively (indicated by black arrows). The Lamarckian inheritance in C prevents genotype-phenotype disengagement, ensuring that the invasion of imitators is reversible, preventing the fixation of social learners through the “drift + ratchet” mechanism that is effective in plots A and B (compare plots D and E). In all maps and plots the selection exponents are set to $x_d = 1$ and $x_r = 10$.

3.4 The occurrence and irreversibility of genotype-phenotype disengagement depends on the dimensionality of the problem space

To the extent that the dimensionality of the problem to be solved is low (i.e., the extent to which L , the length of phenotypes and genotypes, is small), the likelihood of the genetic evolutionary system finding the optimal solution before social learners invade, increases. For low L , the number of bit flip mutations required for an initial random phenotype to become the optimum sequence is smaller, and hence the number of iterations to reach this optimum is also small. Social learners are less likely to reach significant numbers during shorter periods of optimisation, as their high imitation error rate (compared to genetic mutation) is only advantageous in the initial stages of evolution, where solutions are far from the optimum.

In the previous two subsections we have stated that once genotype-phenotype disengagement occurs a reinvasion by instinctive actuators is unlikely. In a disengaged population, genotypes are not subject to selection and hence tend to encode phenotypes that are of lower fitness than the population mean. However, where L is small, genotypes under no selection pressure may occasionally come to encode relatively high fitness phenotypes simply as a result of neutral drift. This becomes extremely unlikely as L becomes large.

When 1) the genotype of a social learner in a disengaged population encodes a phenotype that is fitter than the populations current phenotypic mean; and 2) this genotype is expressed as a phenotype as a result of a mutation to the imitation switch, it is possible for instinctive actuators to invade a population that has undergone genotype-phenotype disengagement. Furthermore, if the rate of genetic mutation (μ_g) is lower than the imitation error rate (μ_p), reinvasion under these conditions is likely to occur during the late stages of optimisation (i.e., when phenotypes are high quality); as the high fidelity of genetic transmission becomes advantageous compared to more error-prone imitation.

In Figure 10, we present supportive results for these claims. Plot A in Figure 10 shows a standard simulation run where irreversible genotype-phenotype disengagement occurs in a simulation where $L = 50$. Plot B in Figure 10 shows a simulation run where instinctive actuators invade after disengagement has occurred in a simulation where $L = 6$. Plot C in Figure 10 shows the tendency for genotype-phenotype disengagement to persist (i.e., the proportion of one hundred simulation replicates where the frequency of social learners is above 0.95 after 5×10^4 iterations) for simulations with different values of L (black line). Also in plot C, the average proportion of 1s in the phenotypes

and genotypes of the populations after 5×10^4 iterations, for 100 replicate simulations, is shown in green and red, respectively. This plot is in agreement with our explanation and demonstrates that for increasing values of L the likelihood of social learners to irreversibly fix also increases.

The results in Figure 10 also serve to support the claim that an extended strategy space is a key element of our simulation model, and that results shown in this paper could not be obtained with the classical two-state strategy space approach (i.e., where $L = 1$) [11, 39, 38].

Figure 10: The effect of problem length, L , on the likelihood of persistent genotype-phenotype disengagement is shown using two representative simulation runs (plots A and B) and a summary of multiple runs (plot C). Plot A shows the result of a simulation run with problem size $L = 50$, where irreversible genotype-phenotype disengagement occurs during the early stages of optimisation. Plot B shows the results for a representative simulation run in which the length of the problem size is $L = 6$. For this value of L genotype-phenotype disengagement is reversible. Plot C shows the tendency for genotype-phenotype disengagement to persist (i.e., the number of simulations where the proportion of social learners is above 0.95 after 5×10^4 iterations) in 100 replicate runs for each value of L between 1 and 100 (black line), and the corresponding averaged proportion of 1s for the phenotype and genotype (i.e., the average proportion of 1s for the phenotype and genotype after 5×10^4 iterations, for 100 simulation runs, for each value of L between 1 and 100) (green and red lines, respectively). For all simulations shown in this figure $\mu_g = 1/(100L)$ and $\mu_p = 1/L$. The strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 10$.

4 Discussion

4.1 Instantiating a non-genetic evolutionary system

An evolutionary system, whether genetic or non-genetic, must instantiate the fundamental elements of heritability, variation and selection. Here we reflect on the mechanisms that provide these elements in the model presented here, and in prior models of non-genetic evolutionary systems. The process of social learning implicitly introduces the first two elements when behaviour is *imperfectly* transmitted from one individual to the next [38]. The balance between mutation and inheritance is as fundamental in these systems as it is in any evolutionary process. On the one hand, too much mutation precipitates an error catastrophe in which offspring are unable to improve on existing parental variants [14, 12]. On the other, excessive fidelity retards evolution when little variation exists for selection to act upon [12]. The consensus within the field is that social learning has a higher error rate than genetic replication during transmission, but that this error rate is not high enough to prevent adaptation [13, 2, 38, 31].

In most models, selection is introduced by decision-making processes such as *transmission bias* and/or *guided variation*. Transmission bias introduces an explicit form of selection by equipping social learners with a direct or indirect bias towards imitating variants with higher reproductive fitness [38, 15]. Similar mechanisms include frequency dependent bias, also referred to as conformity [16] and social status bias [38]. Once a cultural system is established, these two types of biases can evolve without being proxies for fitness bias [38], but these biases lead to the evolution of social learning only if they act as indirect forms of fitness biased imitation. Guided variation introduces a hidden form of selection within every imitator. Models including guided variation assume individuals select amongst the variants that they have copied and improve them, so that the behaviours that they express are fitter than the ones that they imitated originally [38, 31, 1, 5].

Our results show that when survival selection is stronger than reproductive selection, the strategy-space exploration advantage of social learners can effectively offset the reduced selective pressure on their emerging non-genetic system. The separation of survival and reproductive components of selection is a key feature of this work and contrasts with the classical assumption that survivability is merely a proxy for reproductive fitness. This feature has been included in previous analytical work by McElreath and Strimling [29] where an extension of Rogers' basic model evaluates conditions under which vertical imitation is favoured over horizontal imitation. The authors

show that vertical imitation (i.e., imitating ones own parents) is favoured over horizontal imitation (i.e., imitating non-parental agents) when reproductive selection is stronger than survival selection [29]. However, social learners in McElreath and Strimling’s model (i.e., both vertical and horizontal imitators) rely on the presence of individual learners in order to invade; as they do in the original Rogers model (i.e., by imitating fit behaviours from individual learners without themselves having to pay the trial-and-error cost of learning) [39]. In both models, therefore, social learners are only able to invade until they reach an equilibrium frequency in a mixed population [29, 39]. The models proposed by McElreath and Strimling, and Rogers, do not include an extended strategy space nor genetic mutation or imitation error rates [29, 39]. Only when (1) we differentiate between reproductive and survival selection types and (2) we consider these rates in an extended strategy space (i.e., not a simple two-strategy model), can unbiased social learners irreversibly fix in a population of instinctive actuators.

4.2 Limitations

Our model is an extreme simplification of a natural process. It offers, however, a theoretical insight into minimal conditions for the fixation of social learners, conditions that can be considered sufficient for a basic non-genetic evolutionary system to emerge [27, 26, 2]. Our results suggest that any population encountering a new environmental challenge can fix social learners provided that potential solutions in the strategy space enhance survivability more than fertility, and that variation introduced during social transmission is higher than that of genetic inheritance, consequently, affording social learners an exploration advantage. In our results, disengagement should not be interpreted as the complete substitution of entire genomes by non-genetic information. After all, social learning cannot evolve without a biological substrate. We rather consider the genotypes and phenotypes in our model to represent solutions to specific challenges that can either be solved behaviourally through imitated traits or by genetically encoded traits. In our model, once genotype-phenotype disengagement has occurred, behavioural variants are unaffected by the genetic evolutionary system. This is, of course, a simplification. In natural populations the indirect effect of countless other genes (i.e., genes not directly expressing behaviours for a particular challenge) can have an effect on the evolution of behavioural variants [38, 31].

In nature, the map between genetic information and the expressed phenotype is incredibly complex [33]. We are yet to untangle the intricate relations between genes and their developmental

environment, and the potential exploration capabilities of gene networks (i.e., their ability to explore and affect the fitness landscape), which might be greater than current theory has accounted for [42]. Nevertheless, we assume complex behaviours are harder to find genetically than by social learning. The argument is that solutions might be out of reach for genomes due to further developmental constraints and contingencies. While our model does not account for any of these complexities explicitly, it assumes that they would add more constraints on strategy-space exploration for instinctive actuators than for social learners, and therefore including them would more likely increase rather than decrease the chances of social learning evolving.

Our simulation model shows that a high imitation error rate can be detrimental to the evolution of social learners when behaviours are close to the optimum. Some analytical models explicitly include a fixed fitness cost for social learners as a way to represent this imitation error rate (i.e., the “cost” of social learning) [31]. Here, we show this high error rate can be both an enabler and an inhibitor of the fixation of social learners, depending on the optimisation problem and the value of the imitation error rate compared to genetic mutation. In our model no further costs associated with social learning are introduced. All costs related to the development and maintenance of the physiological hardware that allows imitation are considered negligible compared to the fitness value produced by the imitated behaviour itself. All agents compete for the same pool of resources and confront the same challenges when doing so. When population structure, multiple pools of resources, and optimal foraging theory are considered the fixation of unbiased social learners might not happen in the way described here.

In all our simulations the value of μ_c has been set to 0.01. This means, on average, once in one hundred iterations an instinctive actuator turns into an imitator or vice versa. The argument for fixing this rate independent from μ_g is one of simplicity, not artefact. Analyses of our model show that a critical density of imitators is required to start chains of imitation (imitators copying imitators), as sequential events of high-error transmission are required to start a non-genetic evolutionary system that is competitive compared to the genetic one. The lower the chance of these chains forming, the less likely it is for social learners to fix, even under otherwise favourable conditions. Fixing the mutation rate for the imitation switch ensures a constant “background” density of imitators is created by mutation alone, independent of the fidelity of both genetic and social transmission, and that this density is resistant to stochastic fluctuations throughout the simulation. This simplification does not decrease the validity of our analyses, as the critical initial density

of imitators is still marginally small. It is also not the only way of increasing the tendency for imitative chains to form. We hypothesise the same effect can be achieved with $\mu_c = \mu_g$ by imposing a simple spatial population structure and restraining imitation to only take place locally.

4.3 Increasing the long-term overall mean fitness of the population

Our model explores minimal conditions for the evolution of unbiased social learning using an extended strategy space defined by the “onemax” problem on a string of length L . For all parameterisations where social learners fix, the process of genotype-phenotype disengagement prevents instinctive actuators from re-invading the population, therefore maintaining the long-term stability of the non-genetic evolutionary system over its genetic counterpart. However, the genetic system on its own (i.e., a population with no social learners) can also find the optimal solution in this space eventually. It is not an explicit claim of our model that unbiased social learning can increase the long-term overall mean fitness of the population compared to a population of instinctive actuators. Instead, we merely show that by increasing the overall mean fitness *only* during the early stages of optimisation (i.e., blue region in Figure 5) social learners can irreversibly fix.

In theory processes like the *Baldwin Effect* where individuals gradually canalise learned behaviours into instincts [40] could cause the reinvasion of instinctive actuators and the displacement of social learners in the long term. However, this process is unlikely to occur in our model as the rate of optimisation of the non-genetic evolutionary system (i.e., the system driven by social learners) keeps socially learned behaviours far fitter than any instinctive behaviour. After disengagement, the large fitness difference between socially learned behaviours and instincts creates a scenario where no genetic lineage of instinctive actuators could be maintained for long enough to reinvade. Genotypes in the population of social learners accumulate a large number of random mutations. Therefore, whenever one of these agents’ offspring mutates into an instinctive actuator, their expressed behaviour is sure to be outcompeted by the behaviours of social learners. The relatively high rate at which genotypes accumulate mutations, compared to the rate at which they are exposed to any selective pressure, effectively makes it impossible for instinctive actuators to reinvade after genotype-phenotype disengagement has occurred.

Significant effort in evolutionary theory has been focused on explaining the long term advantage of social learning (i.e., how social learners can increase the overall mean fitness of a population) [39, 38]. As we mentioned, our model does not explicitly show that the fixation of unbiased social

learners increases the overall mean fitness of a population. However, in the following paragraphs we suggest two ways in which this could be achieved; (1) considering strategy spaces more complicated than the “onemax” problem and (2) including a rate of environmental change.

With the intention of keeping our fitness function as simple as possible, we have focused on the “onemax” problem in this paper. However, fitness functions can be far more complicated than this [12]. In particular there can be functions where no monotonically increasing paths between local and global optima exists [12, 23]. In these cases populations can get stuck at local optima; unable to “jump” the gap formed by intervening sequences with lower-than-average fitness (e.g., by mutation) and reach a gradient leading to sequences with higher fitness values.

Previous authors have suggested imitation can aggravate this situation by reducing the overall variation in the population [28], a claim that is only true if social learners are assumed to imitate with perfect fidelity. In our model we explicitly include an imitation error rate that is greater than the genetic mutation rate [38]. Therefore, the non-genetic system maintained by unbiased social learning would have a greater chance of crossing low-fitness gaps between optima in a “rugged” landscape [12, 43]. This ability would tend to increase the long-term mean fitness of a population of social learners compared to a population of instinctive actuators, as the former could reach optima inaccessible to the latter.

Another scenario in which unbiased social learners can increase the overall mean fitness of a population (i.e., compared to a monomorphic population of instinctive actuators) exists when a moderate rate of environmental change is included. Our fitness function considers a string of all-ones as an optimal sequence (i.e., phenotypic fitness varies inversely with the hamming distance between a phenotype and the string of all-ones). An equivalent function can be produced if we consider fitness to vary inversely with the hamming distance to some arbitrary bit string and consider this string to be the optimal solution. Given this equivalence, periodically altering the optimal solution throughout the optimisation process and recalculating the corresponding fitness values after each change can represent a rate of environmental change (without this change implying a new *type* of problem being “solved”).

Under a moderate rate of environmental change we hypothesise that neither of the evolutionary systems (i.e., genetic and non-genetic) would be able to drive the average phenotype of the population to the optimum as the rate of environmental change would constantly move this target. If the rate of environmental change keeps the average phenotype value within the blue region of Figure

5 (i.e., the region where social learners find solutions faster than instinctive actuators) unbiased social learners (with their higher imitation error rate) could increase the overall mean fitness of the population. Therefore, under this scenario, the evolution and fixation of social learners would improve adaptiveness at a population level when compared to a monomorphic population of instinctive actuators subjected to the same rate of environmental change. The idea that moderate rates of environmental change can favour the evolution of social learning is in agreement with conclusions derived from existing models [22, 41, 34, 10, 7].

4.4 Our work as a general model for the evolution of horizontal information transfer

Our conclusions not only have implications for the specific case of horizontal imitation, but also for the evolution of horizontal information transfer in general. For example, a mechanistically analogous scenario to the one described in this paper could occur in a population of hosts and symbionts. In this analogous scenario, the spread of symbiotic bacteria from one host to another would be equivalent to the spread of variants in social learning systems. We could imagine members of the host population having a “switch” with two states; one in which they obtain a useful enzyme by inheriting an enzyme-producing bacterial strain with a low mutation rate from their parents, and another in which they rely on the horizontal acquisition of this bacterial symbiont in order to produce the enzyme for them. This acquisition process could take many forms, from generic close-range interactions to specific types of transmission, such as sharing food or ingesting faecal matter. If horizontally transmitted bacteria are subject to higher mutation rates and shorter replication times compared to the parentally inherited bacterial strains (e.g., by virtue of the stress of being outside the host), this would ensure an exploratory advantage for horizontal transmission analogous to the relatively high error rate associated with imitation in social learners.

In this scenario our model predicts that horizontal contagion by symbiotic bacteria, and the associated behaviour that enhances the chances of contagion, can evolve in the host population if the enzyme extends the lifespan of the hosts (hence the chances of infecting others) more than it increases their reproductive output (i.e., survival selection greater than reproductive selection). The evolution of horizontal contagion also depends on the relative mutation rates of horizontally infecting strains and vertically inherited ones, with the former needing to be large enough such that the exploration advantage of horizontal transmission is able to displace vertical transmission during

early stages of the optimisation process (i.e., during the “first period” described in Figure 5).

5 Conclusion

Our model demonstrates that a basic non-genetic system can emerge in a population when selection pressure for survivability is stronger than selection pressure for reproduction. An extended strategy space (where exploration over a large sequence-space is required) distinguishes our approach from previous models and is essential for understanding the exploration advantage of social learning versus genetic inheritance. Analysis of our results leads to a consistent explanation for the emergence of a non-genetic evolutionary system where phenotype and genotype disengage, with the former evolving exclusively by social learning that is unbiased (i.e., non-critical/non-discriminatory/unguided). During this process of disengagement the imitator minority must be able to offset its lower selection pressure with its higher error rate, a condition that is facilitated when the contribution of the evolved trait to survivability is more important than its contribution to reproductive fecundity. Our simulation model offers a very simple framework for the emergence of non-genetic systems and serves as a tool for future research extensions. This simplicity, in particular removing the need for decision-making abilities such as those that result in transmission bias and/or guided variation, lowers the minimal number of traits that a species must have in order to evolve social learning and therefore broadens its potential application.

6 Acknowledgments

The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work. This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1). No data sets are associated with this publication.

References

1. Aoki, K., Wakano, J. Y., & Feldman, M. W. (2005). The emergence of social learning in a temporally changing environment: A theoretical model. *Current Anthropology*, 46(2), 334–340.
2. Avital, E., & Jablonka, E. (2005). *Animal traditions*. Cambridge: Cambridge University Press.

3. Bell, A. V., Richerson, P. J., & McElreath, R. (2009). Culture rather than genes provides greater scope for the evolution of large-scale human prosociality. *Proceedings of the National Academy of Sciences USA*, 106(42), 17671–17674.
4. Blute, M. (1987). Biologists on sociocultural evolution: A critical analysis. *Sociological Theory*, 5(2), 185–193.
5. Borenstein, E., Feldman, M. W., & Aoki, K. (2007). Evolution of learning in fluctuating environments: When selection favors both social and exploratory individual learning. *Evolution*, 62(3), 586–602.
6. Borg, J. M., & Channon, A. (2012). Testing the variability hypothesis: The adoption of social learning in increasingly variable environments. In Adami, C., Bryson, D. M., Ofria, C., & Pennock, R. T. (Eds.), *Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems* (pp. 317–324). Cambridge, MA: MIT Press.
7. Boyd, R., & Richerson, P. (1996). Why culture is common but cultural evolution is rare. *Proceedings of the British Academy*, 88(1), 73–93.
8. Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
9. Boyd, R., & Richerson, P. J. (1988). An evolutionary model of social learning: The effects of spatial and temporal variation. In Zentall, T., & Galef, B. G. (Eds.), *Social Learning: Psychological and Biological Perspectives* (pp. 29–48). Mahwah, NJ: Lawrence Erlbaum Associates.
10. Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16(1), 125–143.
11. Cavalli-Sforza, L., & Feldman, M. W. (1983). Cultural versus genetic adaptation. *Proceedings of the National Academy of Sciences USA*, 80(16), 4993–4996.
12. De Jong, K. (2002). *Evolutionary computation: A unified approach*. Cambridge, MA: MIT Press.
13. Dugatkin, L. A. (2000). *The imitation factor: evolution beyond the gene*. New York: NY Free Press.

14. Eigen, M., & Schuster, P. (1977). The hypercycle; a principle of natural self-organization: Emergence of the hypercycle. *Naturwissenschaften*, 64(11), 541–565.
15. Enquist, M., Eriksson, K., & Ghirlanda, S. (2008). Critical social learning: A solution to Rogers' paradox of nonadaptive culture. *American Anthropologist*, 109(4), 727–734.
16. Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behaviour*, 19(4), 215–241.
17. Heyes, C. (1994). Social learning in animals: Categories and mechanism. *Biological Reviews of the Cambridge Philosophical Society*, 69(2), 207–231.
18. Holden, C., & Mace, R. (1997). Phylogenetic analysis of the evolution of lactose digestion in adults. *Human Biology*, 69(5), 605–628.
19. Ingold, T. (1986). *Evolution and social life*. Cambridge: Cambridge University Press.
20. Jablonka, E., & Lamb, M. J. (2005). *Evolution in four dimensions: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. Cambridge, MA: MIT press.
21. Jenkins, P. (1978). Cultural transmission of song patterns and dialect development in a free-living bird population. *Animal Behaviour*, 26(1), 50–78.
22. Kameda, T., & Nakanishi, D. (2002). Cost-benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behaviour*, 23(5), 373–393.
23. Kauffman, S. A. (1993). *The origins of order*. Oxford: Oxford University Press.
24. Kroeber, A. L. (1948). *Anthropology: Race, language, culture, psychology, pre-history*. New York: Harcourt, Brace and World.
25. Laland, K., & Galef, B. G. (Eds.) (2009). *The question of animal culture*. Cambridge, MA: Harvard University Press.
26. Laland, K. N., & Hoppitt, W. (2003). Do animals have culture? *Evolutionary Anthropology*, 12(3), 150–159.
27. Laland, K. N., & Janik, V. M. (2006). The animal cultures debate. *Trends in Ecology and Evolution*, 21(10), 542–547.

28. Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
29. McElreath, R., & Strimling, P. (2008). When natural selection favors imitation of parents. *Current Anthropology*, 49(2), 307–316.
30. McGrew, W. (1998). Culture in nonhuman primates? *Annual Review of Anthropology*, 27(1), 301–328.
31. Mesoudi, A. (2011). *Cultural evolution: How darwinian theory can explain human culture and synthesize the social sciences*. Chicago, IL: University of Chicago Press.
32. Perreault, C. (2012). The pace of cultural evolution. *PLOS ONE*, 7(9), e45150.
33. Pigliucci, M. (2010). Genotype–phenotype mapping and the end of the ‘genes as blueprint’ metaphor. *Philosophical Transactions of the Royal Society B*, 365(1540), 557–566.
34. Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, 328(5975), 208–213.
35. Rendell, L., Fogarty, L., & Lanland, K. (2010). Rogers’ paradox recast and resolved: population structure and the evolution of social learning. *Evolution*, 64(2), 534–548.
36. Rendell, L., & Whitehead, H. (2001). Culture in whales and dolphins. *Behavioral and Brain Sciences*, 24(2), 309–382.
37. Richerson, P., & Boyd, R. (2002). Culture is part of human biology: Why the superorganic concept serves the human sciences badly. In Goodman, M., & Moffat, A. S. (Eds.), *Probing Human Origins* (pp. 59–85). Cambridge, MA: American Academy of Arts and Sciences.
38. Richerson, P. J., & Boyd, R. (2006). *Not by genes alone: How culture transformed human evolution*. Chicago, IL: University of Chicago Press.
39. Rogers, A. (1988). Does biology constrain culture? *American Anthropologist*, 90(4), 819–831.
40. Simpson, G. G. (1953). The Baldwin effect. *Evolution*, 7(2), 110–117.

41. Wakano, J., & Aoki, K. (2006). A mixed strategy model for the emergence and intensification of social learning in a periodically changing environment. *Theoretical Population Biology*, 70(4), 486–497.
42. Watson, R., & Szathmary, E. (2016). How can evolution learn? *Trends in Ecology and Evolution*, 31(2), 147–157.
43. Weissman, D. B., Desai, M. M., Fisher, D. S., & Feldman, M. W. (2009). The rate at which asexual populations cross fitness valleys. *Theoretical Population Biology*, 75(4), 286–300.