

Research



Cite this article: Morgan TJH, Suchow JW, Griffiths TL. 2020 Experimental evolutionary simulations of learning, memory and life history. *Phil. Trans. R. Soc. B* **375**: 20190504. <http://dx.doi.org/10.1098/rstb.2019.0504>

Accepted: 12 February 2020

One contribution of 16 to a theme issue ‘Life history and learning: how childhood, caregiving and old age shape cognition and culture in humans and other animals’.

Subject Areas:

behaviour, cognition, evolution

Keywords:

learning, memory, childhood, coevolution, human evolution

Author for correspondence:

Thomas J. H. Morgan
e-mail: thomas.j.h.morgan@gmail.com

†These authors contributed equally to this work.

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.4978280>.

Experimental evolutionary simulations of learning, memory and life history

Thomas J. H. Morgan^{1,2,†}, Jordan W. Suchow^{3,†} and Thomas L. Griffiths³

¹School of Human Evolution and Social Change, Arizona State University, 900 S. Cady Mall, Tempe, AZ 85287, USA

²Institute of Human Origins, Arizona State University, 951 S. Cady Mall, Tempe, AZ 85287, USA

³Department of Psychology, UC Berkeley, Tolman Hall, Berkeley, CA 94720, USA

TJHM, 0000-0002-2848-8461

Humans possess an unusual combination of traits, including our cognition, life history, demographics and geographical distribution. Many theories propose that these traits have coevolved. Such hypotheses have been explored both theoretically and empirically, with experiments examining whether human behaviour meets theoretical expectations. However, theory must make assumptions about the human mind, creating a potentially problematic gap between models and reality. Here, we employ a series of ‘experimental evolutionary simulations’ to reduce this gap and to explore the coevolution of learning, memory and childhood. The approach combines aspects of theory and experiment by inserting human participants as agents within an evolutionary simulation. Across experiments, we find that human behaviour supports the coevolution of learning, memory and childhood, but that this is dampened by rapid environmental change. We conclude by discussing both the implications of these findings for theories of human evolution and the utility of experimental evolutionary simulations more generally.

This article is part of the theme issue ‘Life history and learning: how childhood, caregiving and old age shape cognition and culture in humans and other animals’.

1. Introduction

Humans are an unusual species: we inhabit an enormous variety of terrestrial ecosystems [1–3], with total human biomass exceeding that of any non-domesticated species [1]. Compared to other species, we build and use a wider variety of tools [4–6], exhibit complex cognitive abilities like theory of mind [7–10] and causal reasoning [11,12], and make more extensive use of cultural inheritance [13–16]. Across mammalian species, life history shows a strong relationship with body size [17], but taking body size into account, we have an unexpectedly long juvenile period, a long adult lifespan and high early fertility followed by an extended post-reproductive period [18,19]. Even in the details, human life history is unusual: we have a long gestation period, high birthweight and early weaning, to name but a few [20].

Much research proposes that these various features have coevolved. For instance, it has been argued that post-reproductive adults function to provision younger kin (such as juveniles or lactating women) with resources [21,22]. An alternative (but not necessarily competing) hypothesis is that human life history has been shaped by our dependence on cultural inheritance (see [23,24]). From this point of view, an evolutionary function of the human post-reproductive period is to increase the fitness of younger kin by transfer of *information*. Similarly, the juvenile period may have been extended or modified to facilitate learning ([24], also [25], though see [26]). In support of this, a study of teaching among Fijian villagers found that grandparents were the next most likely individuals (behind parents) to act as teachers [27]. Similarly, in the same populations, food taboos were preferentially learned from mothers, grandmothers and other elders [28].

In many cases, the evidence marshalled in support of co-evolutionary hypotheses such as these is theoretical. For instance, the coevolution of learning and memory [29] has been examined using a model in which individuals' fitness depends on solving a series of 'multi-armed bandit' problems, in which an individual must learn which of several options (arms) produces a reward in a given context (bandit). For example, an individual might learn which food resource to collect, contingent on the weather. In the model, learning and memory were each underpinned by a genetic locus. The learning gene determined the number of arms individuals assessed at each bandit, after which they would choose the arm associated with the greatest payoff. The memory gene determined the length of time over which individuals were able to recognize previously visited bandits and so repeat past decisions without engaging in further learning. Both learning and memory came at a cost to fitness, meaning they would evolve only if they enhanced decision-making. The results showed that under many conditions, learning and memory will coevolve, i.e. they will evolve as a pair, but not in isolation. This is because investing in information collection pays off only if the information can be accurately stored; similarly, investing in information storage pays off only if enough valuable information is stored for later use.

Other evidence is empirical. For instance, Buchsbaum *et al.* [30] argue for the coevolution of the extended juvenile period, pretend play and causal reasoning on the basis of empirical data showing that pretend play and causal reasoning are related in young children. Specifically, they observe a correlation between a child's ability to engage in accurate causal reasoning and their ability to engage in causally coherent pretend play. The argument is that this correlation reflects the intertwined evolutionary history of these traits: (i) the elongation of the human juvenile period gave more opportunities for play, (ii) the increase in play supported the development and evolution of causal reasoning, and (iii) the evolution of causal reasoning, in turn, increased the benefit of play and so enhanced the benefits of a longer juvenile period.

The conjunction of theoretical and experimental approaches to evolution has been highly successful, but even so, it has limitations. Theoretical work typically necessitates a number of assumptions in order to remain feasible. In the case of cognition, this implies a simplified depiction of the human mind. The traditional remedy for this is to conduct laboratory experiments with human participants to test the predictions of theory. But such experiments are limited to comparing predicted equilibria with human behaviour, and so provide only indirect evidence of the evolutionary dynamics that produced them.

Here, we present a series of 'experimental evolutionary simulations' concerning the coevolution of cognition and life history that provide empirical data not only about human behaviour, but also about the evolutionary dynamics that it produces. To do this, we recruit human participants as agents inside an evolutionary simulation. We consider the coevolution of the ability to gather information (learning), store information (memory) and an early period of information acquisition (childhood). Our goals are threefold. First, to test theoretical predictions concerning the coevolution of learning, memory and childhood. Second, to provide novel insights into the coevolution of these traits in human history. Third, to illustrate the feasibility of

conducting experimental evolutionary simulations with human participants as a means to test evolutionary hypotheses. We begin by providing an introduction to experimental evolutionary simulation before continuing with our specific experiments.

2. Experimental evolutionary simulations

An experimental evolutionary simulation can be most readily understood as a traditional theoretical simulation, except that human participants make decisions on behalf of simulated agents. As such, while the agents inhabit a simulated world, their decisions are not simulated but instead result from human psychology.

Theoretical simulations also include the genetic evolution of cognitive traits, with heritable simulated genes affecting simulated decision-making, phenotypes and fitness. Human participants, however, arrive with pre-existing genes and cognition which cannot be directly manipulated. Nonetheless, we can *mimic* the genetic evolution of human cognition by assigning each agent a simulated genotype that affects the task posed to the participant. For example, consider a simulation in which a mutant 'social learner' genotype invades a population of 'asocial learners' (as per [31]). All participants are required to solve a fitness-relevant problem, but their genome determines whether they receive social or asocial information; participants with the asocial allele have access only to environmental cues, while participants with the social allele can view the decisions of their groupmates. If social information improves decision-making, then the fitness of the social learners will be higher, and the social learning allele will spread via natural selection. This is the approach we use here, with simulated genes masking aspects of human cognition, and selection 'unmasking' them when they increase fitness.

The inclusion of simulated genes differentiates this approach from typical evolutionary experiments with human participants. Experimental designs such as the transmission chain and micro-society [32] have a rich history of being used to experimentally model human cultural change, including the cultural transmission of folk stories [33], stone knapping skills [34], arrow head designs [35], paper plane designs [32], computer code [36] and interpretations of Rorschach diagrams [37]. However, they do not include a genetic component. In this regard, our approach resembles experimental studies of genetic evolution in fruit flies [38] and bacteria [39], although the genetic component remains simulated and the participants are human. As such, our focus is not on genetic change *per se*, but rather how human behaviour drives genetic evolution. It is therefore critical to our approach that aspects of decision-making are assigned to human participants, allowing an agent's phenotype (and in turn, their fitness) to be jointly determined by both human cognition and simulated genes. As such, the evolutionary dynamics produced result from both the assumptions of the simulation and human psychology.

(a) Strengths

The primary strength of experimental evolutionary simulation is its ability to produce genetic evolutionary dynamics from human psychology. This can be contrasted with traditional theory that produces evolutionary dynamics but relies on agents with simulated psychologies, and traditional

experiments that involve human psychology but do not produce evolutionary dynamics. With the current approach, we can go further than comparing theoretical predictions about equilibria with experimental data. We can instead compare theoretical predictions about the entire evolutionary process with data collected from an evolving population of human participants. This is critical to questions about coevolution, which are fundamentally about the process of evolution, not just the outcome.

Moreover, by including real (unsimulated) human participants, the need to make assumptions about the psychology of simulated agents is removed. This is important because theory is not just about understanding how evolution works in the abstract (although this is certainly part of its value [40]), but also about trying to understand specific cases. For example, we have argued in support of gene–culture interactions between human cognition and stone tool technologies on the basis of theory showing that a model including gene–culture interactions produces dynamics that more closely resemble the archaeological record than do those of a model without such interactions [41]. In such cases, the gap between theory and historical reality created by assumptions may be problematic and the inclusion of human psychology in evolutionary simulations goes at least some of the way towards closing the gap.

A further benefit of experimental evolutionary simulations is that they naturally include inter-individual variation in decision-making to the extent that it exists in the human population recruited from. This can be contrasted with theory, where variation is often ignored under the assumption that it is simply noise, despite the fact that theory modified to include such variation shows it has evolutionary consequences. For instance, behavioural variation supports the evolution of cooperation [42,43] and may itself be an evolved adaptation [44].

(b) Limitations

As with any method, there are limitations to experimental evolutionary simulations. Moreover, while in some ways this approach combines the strengths of traditional theory and experiment, it also brings with it some of their weaknesses. For instance, as already mentioned, while real (unsimulated) human psychology is introduced into the evolutionary process, the rest remains simulated and so necessitates assumptions by the researcher. Moreover, the participants in the experiment may not be representative of our species as a whole, a problem faced by all experimental work [45].

There are also some limitations specific to experimental evolutionary simulations. For instance, although this approach aims to produce reliable experimental dynamics, it is currently not plausible to aim for the population sizes or number of repeat simulations typical of theory, which would collectively require millions of participants. As such, experimental evolutionary simulations will likely involve smaller populations and more constrained regions of parameter space than typical theory. They may also require preliminary theory to identify the population sizes and selection strengths required for reliable dynamics to be produced.

Nonetheless, modern experimental methods allow experimental simulations to involve numbers of participants far larger than in many traditional experiments (here we present data from 4800 participants). A downside of this is that such

experiments may be logistically challenging, and are also likely to be more expensive than smaller-scale studies. Despite this, online recruitment services and experimental software packages are reducing this burden. Here, we used the experimental automation software *Dallinger* (see dallinger.readthedocs.io) to conduct the experimental simulations, recruiting large numbers of participants from Amazon's Mechanical Turk.

A final limitation of this approach is that where it examines the evolution of cognition and behaviour, the psychology of participants cannot itself be directly manipulated. Rather, experimental manipulations of the task are used to mimic the evolution of cognitive abilities, by selectively displaying stimuli or forcing certain options on participants. The precise nature of these manipulations will vary across experiments (and may not always be desired), but it is important to note that they act only as a proxy for the evolutionary process of interest.

(c) Summary

Experimental evolutionary simulations combine elements of theory and experiment by recruiting participants to make decisions on behalf of agents within an otherwise theoretical simulation. This allows researchers to observe the evolutionary dynamics produced by human decision-making, and includes naturalistic human variation in the evolving population. Nonetheless, it has limitations and is intended to complement rather than replace existing methods.

3. Experiment 1

The first experiment verifies that our experimental approach produces robust evolutionary dynamics by examining whether the coevolution of learning and memory observed in theory [29] occurs when our experimental design constrains human behaviour to closely match theoretical assumptions. To verify coevolution is taking place, we conducted two parallel simulations, one in which memory could evolve and another in which it could not. We predicted that learning would evolve to greater levels when memory was also capable of evolving than when it could not. This experiment sticks close to theoretical assumptions and will serve as a baseline for comparison with experiments 2 and 3.

(a) Methods

We carried out two experimental evolutionary simulations. In each simulation, a population of 40 participants evolved for 40 generations (i.e. 1600 participants per simulation, each participant playing the role of a single agent). This population size and simulation length was chosen on the basis of preliminary theory that suggested it was sufficient to produce robust evolutionary dynamics.

Participants were recruited through Amazon's Mechanical Turk. While we did not collect demographic data, several studies have examined the demographics of Mechanical Turk workers and provide relevant information. The total worker population is estimated to be stable at over 100 000 people, although the effective population size may be smaller [46], with around 2000 workers active at any one time [47]. Most workers are US-based, with a sizable number of Indian workers, and the US and Indian sub-populations being most



Figure 1. Examples of the task across three different countries. (a) Upon arrival at a new country, participants are instructed to check a number of possible locations. (b) Checking a location reveals either the treasure or a red X. After finishing checking, all locations return to their original images and the participant must make a final decision. (c) If participants revisit a familiar country, they cannot check locations again and are asked to make a single decision immediately.

active at different times of day [47,48]. This may explain why participant characteristics and behaviour can also vary by time of day, as well as on weekends versus weekdays [49]. The worker population is diverse and somewhat comparable to the overall US population in terms of age, sex, education and socio-economic status, and is significantly more diverse than typical undergraduate populations [50]. Nonetheless, it is younger, more female, better educated and less rich than the US population as a whole [47,51]. Despite these differences, Mechanical Turk workers behave similarly to the general US population in many ways [52,53].

Upon being recruited, participants gave their consent, were given instructions, completed the experiment and were then debriefed and paid. Recruitment occurred one generation at a time: experimental slots were made available in batches of 40 and once all 40 participants had successfully taken part, the next batch was made available. This process repeated until all generations were complete. Recruitment and testing were approved by the Committee for Protection of Human Subjects at University of California, Berkeley (protocol ID 2015-12-8227). The experiment took around 5 min and participants were paid \$0.50 for taking part, with an additional performance-related bonus of up to \$0.50.

Within each simulation, each participant made decisions on behalf of a single simulated agent. Each agent 'lived' for 20 time periods, and in each time period, participants were asked to complete a single trial of a multi-armed bandit

task for their agent. The task was framed as a treasure hunt (figure 1); participants were told they were visiting a country (e.g. 'France') and were asked to decide which of the 10 possible locations in that country (e.g. 'The Louvre') contained the treasure. Only one location in each country contained treasure, and the treasure did not move within or between trials.

Agents were assigned a simulated genome consisting of two genes: a learning gene (L) and a memory gene (M), both of which were represented as integer values ($0 < L \leq 10$, $M \geq 0$). An agent's genome placed certain constraints on how their participant could interact with the task, although participants were not told anything about their agent's simulated genome or its effects.

The country associated with each trial was determined as follows: for each participant, at the start of the experiment, four countries were selected at random from a list of 20. For each trial, one of these four countries was selected at random as the country for that trial. However, if at any point a country was not visited within M trials, it was replaced with a previously unused country randomly selected from the initial list of 20 countries. For instance, a participant with $M=5$ may initially be assigned the four countries France, Senegal, Mongolia and Italy. If their first trial takes place in Italy, but Italy is not visited again across trials 2–6, it is 'forgotten' and the next time Italy is selected, it is replaced with a different country from the original list (excluding France, Senegal, Mongolia and Italy). This procedure mimics forgetting by

limiting how long information stored by participants remains useful; in the above example, information about Italy was rendered useless after it was not used for five trials. However, note that participants may still forget information across fewer trials (and they sometimes did, see below). As such, the value of M places a *cap* on memory, but does not guarantee information can be stored for that duration.

On their first visit to a country, participants were required to ‘check’ L locations for the treasure (L being the value of their agent’s learning gene). Participants checked a location by clicking on it, and checking accurately revealed whether or not the treasure was present at that location (figure 1*a,b*). After checking L locations, participants were asked to choose which location they thought had the treasure. On trials at a previously visited country, participants could not check any locations, and were required to make a decision based on what they remembered from previous trials (figure 1*c*). As such, the experimental design (including the value of L for each participant) entirely determined how much information participants collected on each trial, and participants could only decide *which* locations to check. These heavy constraints were used in order to check that our experimental simulations could produce known results when constraints effectively forced human behaviour to match theoretical expectations [29].

After completing all 20 trials, agents were assigned a fitness value f

$$f = 10 + 10N_s - N_c - 2M, \quad (3.1)$$

where N_s is the number of trials on which their participant chose the correct location of the treasure and N_c is the number of times their participant checked a location. Note that N_c and M are entirely determined by the experimental design, but N_s depends on both the experimental design and human behaviour (e.g. how good participants are at remembering the location of the treasure). Negative fitness values were set to 0 and all values were additionally squared to increase fitness differences and allow selection to proceed more rapidly. While this function produces selection far stronger than is estimated in wild populations, strong selection is a common assumption of evolutionary simulations in order to produce reliable dynamics over reasonable time scales.

After fitness was calculated, the next generation of agents (participants) was created (recruited). Agents inherited their genes via asexual reproduction from a single agent in the previous generation, selected with probability proportional to fitness. Given the fitness function, this implies that increasing M or L will decrease simulated reproductive success, unless it can be offset by increased decision-making accuracy by the participant. Genetic inheritance was subject to mutation: there was a 50% chance the value of each gene would be unchanged, a 25% chance that the value would increase by 1 and a 25% chance that it would decrease by 1, unless the mutation produced a value outside the permitted range (e.g. a negative value for M), in which case it was prevented. This mutation rate is higher than in wild populations, but was chosen to facilitate evolution over the short experimental time scales.

Participant performance affected their bonus, b , which was calculated as

$$b = \max\left(\min\left(\frac{10N_s - N_c}{200}, 0.5\right), 0\right). \quad (3.2)$$

The bonus calculation incentivizes participants to find the treasure. Note that, because checking locations decreases the bonus, participants who perform well despite having a low value for L are rewarded with a larger bonus than those who perform similarly but with a larger value of L .

(b) Analyses

We analysed the data with Bayesian models using Markov chain Monte Carlo methods via the *rjags* package in R [54]. Parameter estimates are provided as the median and 95% highest density interval of greater than 3000 independent samples generated from three chains using the Gelman–Rubin statistic (upper CI ≤ 1.01) to check for convergence.

For each experimental simulation, we separately modelled each agent’s learning capacity L and their memory capacity M , as a normally distributed variable ($N = 1600$ per simulation) with the following structure

$$\begin{aligned} L \text{ or } M &\sim N(\mu, \sigma^2) \\ \mu &= \beta_g \\ \beta_g &\sim N(0, 100) \\ \sigma^2 &\sim \text{Gamma}(0.001, 0.001) \end{aligned}$$

where β_g is a parameter to be estimated that takes independent values for each generation in the simulation.

To explore how participants’ learning behaviour changed across their lifespan and across generations, we modelled whether participants engaged in learning (i.e. checked any locations) on a given trial as a Bernoulli variable ($N = 32\,000$ per simulation) with the following structure

$$\begin{aligned} C &\sim \text{Bern}(p) \\ p &= e^{-\text{rate} \times (\text{trial} - 1)} \times (1 - \text{floor}) + \text{floor} \\ \log(\text{rate}) &= \beta_1 + \beta_2 \times (\text{generation} - 20) \\ \text{logit}(\text{floor}) &= \beta_3 + \beta_4 \times (\text{generation} - 20) \\ \beta_{1:4} &\sim N(0, 100) \end{aligned}$$

where $\beta_{1:4}$ are parameters to be estimated. The model assumes that the probability a participant engages in learning on the first trial is 1 (necessarily true, due to the experimental design) but that it then decays at a speed given by *rate* towards the value given by *floor*. Because both *rate* and *floor* are functions of generation number, this allows the relationship between age and behaviour to change across generations.

(c) Results

In the simulation where both learning and memory could evolve, both learning and memory increased (L : $\beta_{40} = 9.15$, [8.75, 9.55]; M : $\beta_{40} = 6.67$, [6.20, 7.13]), while the estimated standard deviations were 1.29, [1.25, 1.34] and 1.51, [1.45, 1.56], respectively. These closely match theoretical expectations (figure 2*a*). This similarity was expected because constraints on participants’ behaviour meant that they could fail to match theory only by (i) finding the treasure but not choosing it (occurred on 484/8590 relevant trials, 5.6%), (ii) not finding the treasure but choosing somewhere they knew the treasure was not (200/6639, 3%), or (iii) not choosing the treasure’s location when revisiting a country where they had previously found it (921/4595, 20%). Though all three behaviours occurred, only the third was common, and we suggest it likely resulted from participant memory errors or inattention. The modification of preliminary theory to include this error suggests that (at observed rates) it dampens the coevolution, but does not prevent it.




Figure 2. (a–c) Experimental evolutionary dynamics produced across experiments 1–3. Solid lines show the median model estimates of the mean value of each gene. The dark shaded area is the 95% HDI for the mean gene value. The light shaded area shows the median estimate of the standard deviation in the gene values. (a) In the first experiment, experimental dynamics closely match the ‘max-remember’ strategy in which participants check L locations at unfamiliar countries, 0 at familiar countries and perfectly remember the location of the treasure for M trials, although this largely reflects the experimental design. (b) In the second experiment, experimental dynamics closely match those of the ‘flex-remember’ strategy in which participants flexibly use their learning ability to minimize costs (stopping as soon as the treasure is found) and perfectly remember the location of the treasure for M trials, but not the ‘guess’ strategy. (c) In the third experiment, the experimental dynamics do not closely match any theoretical strategy. Learning evolves rapidly, as in the ‘flex-remember’ or ‘check-first’ (wherein participants double check the past location of the treasure at familiar countries) strategies, but memory evolves much more slowly, resembling the ‘max-remember’ strategy seen in experiment 1. (d–f) The evolution of childhood presented as how much learning participants engaged in across their lifespan (i.e. trials) for different generations. The solid lines are raw data averages, the semi-transparent lines show 100 model estimates of the most likely behaviour. (d) In the first experiment, participants evolve to decrease their probability of learning on later trials, resulting in an early ‘childhood’ followed by a later exploitation phase. (e) In the second experiment, participants initially engage in a small amount of learning but do so consistently across their lifespan; however, they evolve to do a great deal of learning on the early trials and then very little for the rest of their lifespan. Note that the model struggles to perfectly fit the raw data. This is because the model assumes steady, linear change in life history over generations while the evolutionary dynamics (see b) are clearly not linear. Nonetheless, the model does detect the evolution of childhood as discussed in the results. (f) In the third experiment, participants engage in a considerable amount of learning across their lifespan, and the amount of learning increases across generations. There is evidence that individuals learn less as their lifespan progresses, but the effect is very small, and there is minimal evidence that this is any more pronounced in later generations than in earlier generations.

Moreover, its effect decreases when there is inter-individual variation (i.e. when most errors are due to a minority of forgetful individuals).

In the simulation where memory was fixed to 0, L also increased (figure 2a, $\beta_{40} = 3.59$ [3.19, 4.04], s.d. = 1.36 [1.32, 1.41]), but was considerably lower than when memory was also allowed to evolve. This provides evidence that memory and learning were coevolving in the first simulation.

There was also a marked change in how much learning participants engaged in across their lifespan as the simulation progressed (figure 2d, $\beta_{1,4} = -0.92$, [−0.99, −0.86]; −0.04, [−0.05, −0.04]; −0.82, [−0.87, −0.77]; −0.10, [−0.11, −0.10]). In the first generation, there was only a modest decrease in the probability that a participant engaged in learning as their simulated lifespan progressed (probability of learning on final trial = 0.77, [0.76, 0.78]), whereas in the final generation, this decrease was much more severe (probability of learning on final trial = 0.09, [0.09, 0.10]). As such, the coevolution of learning and memory produced a life history where individuals engaged in a lot of learning over their first few

trials, but they were unlikely to engage in learning for much of the second half of their lifespan. In the simulation without memory, participants engaged in learning on every trial, so no such trajectory evolved. Note, however, that these results largely reflect the assumptions of the experimental design, but they will nonetheless serve as a baseline for comparison with subsequent experiments.

4. Experiment 2

The first experiment produced the predicted coevolution, but also placed heavy constraints on participants’ behaviour. As such, we conducted a second experiment to explore the evolutionary dynamics produced without these constraints.

(a) Methods

We carried out a single experimental evolutionary simulation in which both learning and memory could evolve, with the following differences: (i) Participants could check up to L

locations at *any* country, whether familiar or not. As such, the experimental design (including L) no longer determines how much information participants collect, but instead places a cap on how much information they collect. (ii) Because this decouples learning behaviour from learning capacity, the fitness costs of learning were split such that both learning capacity (L) and checking locations (determined by the participant) decreased fitness

$$f = 10 + 10N_s - 0.5N_c - 2M - 0.5L. \quad (4.1)$$

(iii) We increased the number of trials to 40 to increase the reliability of the evolutionary dynamics. (iv) In the light of these extra trials, we increased the base payment and potential bonus payment to \$0.60, where the bonus is calculated as

$$b = \max\left(\min\left(\frac{3(10N_s - N_c - 40)}{1000}, 0.6\right), 0.0\right). \quad (4.2)$$

(b) Analyses

Analyses of L and M were the same as in experiment 1. Rather than modelling whether participants engaged in learning on a given trial, we instead modelled the amount of learning participants engaged in (i.e. the number of locations checked, N_c) a given trial ($N = 32\,000$) with the following structure

$$N_c \sim \text{Poisson}(\mu)$$

$$\text{Log}(\mu) = \beta_1 + \beta_2 \times \text{generation} + (\beta_3 + \beta_4 \times \text{generation}) \times \text{trial} + \beta_{5,\text{participant}}$$

$$\beta_{1,4} \sim N(0, 0.01)$$

$$\beta_{5,1:1600} \sim N(0, T)$$

$$T \sim \text{Gamma}(0.01, 0.01)$$

(c) Results

The optimal strategy in this experiment is to check locations until you find the treasure, and then rely on memory (see electronic supplementary material, §1). Participant behaviour was broadly consistent with this, but less so when revisiting a country. At unfamiliar countries (14 730 trials), participants were most likely to search until they found (and then chose) the treasure (49%) or until they reached their limit of L locations (29% of trials). On the remaining trials, participants guessed without searching (13%); began searching, but then guessed anyway (5%); continued searching after finding the treasure (4%); or found the treasure, but choose a different location (less than 0.5%). When returning to a country where the treasure had previously been found (9113 trials), participants generally did not repeat their past decision without further searching (27%). Instead, they often rechecked the treasure's prior location (45%). Such behaviour may be optimal, given the fallibility of human memory (see memory errors in experiment 1, and electronic supplementary material, §2). Alternatively, it may reflect that participants did not trust the experimenters claim that the treasure did not move. Regardless, the evolutionary dynamics that resulted closely matched those produced by a theoretical simulation assuming the optimal strategy (figure 2b, L : $\beta_{40} = 8.82$, [8.42, 9.22], s.d. = 1.29, [1.25, 1.34]; M : $\beta_{40} = 6.09$, [5.54, 6.65], s.d. = 1.78, [1.72, 1.85]), suggesting that the rate of double-checking locations was small enough to only minimally change evolutionary dynamics.

As in the first experiment, a period of childhood emerged in the sense that, by the final generation, participants engaged in a

lot of learning over their first few trials, but much less across the rest of their simulated lifespan (figure 2e). Across generations, the amount of learning on the first trial increased ($\beta_1 = 0.97$, [0.93, 10.1]; $\beta_2 = 0.025$, [0.022, 0.029]), but the change in the amount of learning across trials decreased ($\beta_3 = -0.034$, [-0.035, -0.033]; $\beta_4 = -0.0019$, [-0.0020, -0.0018]). In sum, the coevolution between learning, memory and childhood remained robust to giving participants increased freedom in their behaviour, lending support to the hypothesis that learning, memory and childhood have coevolved in human history.

5. Experiment 3

Theory suggests that without a proportionate change in behaviour, the coevolution of learning and memory is disrupted by unpredictable environmental change [29] (see also [26]). This is because environmental change renders previously collected information out of date and so reduces the payoffs associated with long-term information storage. One possible solution is for learners to assume that the environment can change and check previously collected information, which may be a routine part of human behaviour [55]. However, accommodation to such so-called restless bandits is typically imperfect, and so it is unclear how reliably participants' behaviour will produce the coevolution in an unpredictable environment. To determine this, in our final experiment, we conducted one further simulation in which the environment was unpredictable.

(a) Methods

The experimental design was the same as in the prior experiment with the following changes. (i) For each country, on every trial, the location of the treasure was randomized with probability 0.4, with all locations being equally likely to be chosen. (ii) The fitness function was changed as follows:

$$f = 10 + 10N_s - N_c - 4M - 4L. \quad (5.1)$$

These changes were made following preliminary theory that suggested, under these conditions, the evolutionary dynamics produced by different candidate behaviours are markedly different (figure 2c).

In this experiment, the near-optimal strategy (see electronic supplementary material, §3) is to keep checking locations until the treasure is found, and to recheck this location upon returning, continuing to search if it has moved. This strategy supports the coevolution of learning and memory, even when many other strategies do not (figure 2c).

(b) Analyses

We conducted the same analyses as in the second experiment.

(c) Results

Participant behaviour was only partially consistent with the near-optimal strategy of checking the most recent location of the treasure when revisiting a familiar country. As before, when participants arrived at an unfamiliar country (40 875 trials), they usually checked at least one location (89%). On these trials ($N = 36\,271$), participants were likely to continue checking until they found the treasure or ran

out of checks, stopping early on 1962 trials (5%). When participants located the treasure ($N = 22\,005$), they typically stopped checking (19 936 trials, 91%). When participants returned to a country at which they had previously located the treasure ($N = 13\,904$), they immediately checked its previous location on 41% of trials, but committed to the past location without checking on only 4.2% of trials. Unexpectedly, on the majority of such trials, participants initially checked a location other than the reward's previous location (53%). This discrepancy may well reflect memory errors; such errors were produced at lower rates in experiments 1 and 2, and experiment 3 is more complicated and so errors may be more likely. Alternatively, some participants may have concluded the environment was sufficiently unpredictable that they should just ignore their remembered information entirely, an inefficient yet simple strategy.

The evolutionary dynamics produced do not closely match those of any strategy considered theoretically (figure 2c). Learning increased rapidly ($\beta_{40} = 9.30$ [8.97, 9.63]). However, memory showed only a modest increase ($\beta_{40} = 3.25$ [2.77, 3.74]), more than expected, assuming participants do not monitor for environmental change (where memory is selected against), but less than that expected with the optimal strategy. Accordingly, the mismatch between human behaviour and the optimal strategy suppressed the evolution of memory. Nonetheless, there may still have been a coevolutionary interaction because the modest increase in memory likely relied on the evolution of learning.

Unlike the prior two experiments, a period of childhood only minimally emerged. Instead participants continued to engage in a large amount of learning across their lifetimes, although by the later generations, they typically engaged in more learning in the first trial than in later trials (figure 2f). The amount of learning on the first trial increased over generations ($\beta_1 = 1.11$, [1.07, 1.25]; $\beta_2 = 0.030$, [0.027, 0.033]), but the change in the amount of learning across trials was very small ($\beta_3 = -0.004$, [-0.005, -0.003]), with only weak evidence for a meaningful interaction with generation ($\beta_4 = -0.000$, [-0.000, 0.000]). In sum, the coevolution between learning, memory and childhood was significantly dampened by environmental unpredictability.

6. Discussion

We have presented the results of a series of experimental evolutionary simulations examining the coevolution of learning, memory and childhood. Across three experiments, we (i) verified that our experiments produce robust evolutionary dynamics by showing that when human behaviour is constrained to match theoretical expectations then the resulting evolutionary dynamics closely match theory too; (ii) found that even when assumptions about participant behaviour are relaxed, the coevolution between learning and memory still occurs, including the evolution of childhood; and (iii) found that environmental unpredictability suppresses the evolution of memory and the emergence of childhood.

Collectively, these results support existing theory on this topic [29] by reinforcing the conclusion that learning and memory have coevolved across human evolutionary history. Moreover, they provide evidence that this coevolution would also have produced a period of intense learning in early life history. While this relationship between learning and life

history is part of many theories of human evolution (e.g. the cognitive niche [56,57], cultural niche [41,58], collective brain [59,60] and cultural drive [61] hypotheses), there has been little direct modelling work (though see [26,62,63], for models of the coevolution of cognition and life history more generally, and [64,65] for models of when individual versus social learning should be scheduled across the lifespan), with most supporting evidence coming from large-scale comparative studies, which find a correlation between learning and childhood across species [66,67]. Thus, experimentally simulating the coevolution of learning, memory and childhood provides a direct form of evidence that was previously unavailable. However, this work places important caveats on this hypothesis, showing that environmental unpredictability can disrupt the coevolution by impeding the evolution of memory and, in turn, childhood. This is because environmental change prevents information from remaining useful indefinitely, reducing the benefit to information storage such that it is favoured by selection only if memory is sufficiently accurate (as was the case in the hypothetical check-first strategy, but not in actual human behaviour).

One outstanding concern is that although these results suggest that environmental unpredictability can disrupt the evolution of childhood, humans (who have an extended juvenile period [18,19]) evolved during a period of unusually high environmental variability [68–70]. Key here is to consider the scale of the environmental variation in question. In our experiment, the treasure had a probability of 0.4 of moving on any given trial. Given that there are four possible destinations for each trial, this makes total environmental change (i.e. all four treasure items moving) probable every four trials or so (57% chance after 4 trials, 98% chance after 10 trials). Given that agents lived for 40 trials, and assuming a human lifespan of 70+ years, this corresponds to significant environmental change roughly every decade. How does this compare to what is known about environmental variation over human history?

The past 2 Myr have continued the general post-Mesozoic trend of an increasingly cool, dry and unstable climate. In particular, an approximately 41 ky temperature oscillation started around 1.5 Ma, shifting to a higher amplitude 100 ky cycle around 650 kya [68]. However, these cycles are extremely slow compared to human lifespans. Did more rapid variability exist too? One limitation is that the resolution of current methods is insufficient to identify small-scale variation in the distant past, but finer details are visible over more recent history. Such analyses find abundant evidence of sudden climate fluctuations during the two most recent glacial periods (i.e. the most recent 250 ky), with temperature changes of up to 7°C occurring on time scales comparable to human lifespans [69,71,72]. Moreover, historical data concerning the 'little ice age' (lasting from roughly 1300–1850 CE) show extensive variation on even shorter time scales that significantly disrupted agriculture, fishing and human populations more generally [73]. Collectively, this suggests that submillennial (and even centennial) environmental change was potentially widespread in human evolutionary history. Whether the rate of change approached the decadal change simulated in our experiment remains unclear due to limitations of current methodology, but it is perhaps less likely. Regardless, our data support the prediction that such extremely rapid change was rare as otherwise it may have prevented the evolution of human cognition and life history.

Although our experiments have focused on human evolution, the theory on which they are based [29] is not, and so

these results are likely applicable to non-human species also. Coupled with supporting comparative data [67], it is plausible that memory, learning and life history have coevolved across multiple species. Nonetheless, our results were generated using human participants and similar experiments with other species may produce different evolutionary dynamics. While using populations of non-human animals at the scale required here may be a challenge, large-scale cultural evolutionary experiments have been conducted in baboons [74] and so it may be possible to replicate these experiments with other species.

The results presented here also support 'experimental evolutionary simulations' as a means to study human evolution. Our first experiment validated the approach by replicating theoretical work. However, participants did not behave entirely in line with theoretical expectations and in the final experiment, this was sufficient to cause the resultant dynamics to diverge from theoretical predictions. This highlights another value of experimental simulations, as they can generate evolutionary dynamics unanticipated by theory and impossible via traditional experimental methods. In this case, a theoretical model that assumed the optimal strategy (double-checking previous treasure locations in the case of environmental change) would have concluded that the coevolution was robust even to the level of environmental unpredictability included in our experiment. However, by using fallible human participants instead, we found that the coevolution was less robust than that the optimal strategy predicts, with memory suppressed and little evidence of the evolution of childhood.

We do not suggest that experimental evolutionary simulations should be considered a replacement for either traditional theoretical or experimental work. Rather they are complementary, bridging the gap between the evolutionary dynamics produced by theoretical work and the descriptions of human behaviour produced by experimental work. Indeed, the experiments presented here have several limitations that could be addressed with additional theory, experiments or experimental simulations. One is that while learning and memory are associated with underlying genes, the evolution of childhood in our study is a purely emergent consequence of the decision-making of adult participants. Indeed, our definition of childhood for the purposes of this work is very limited, focusing solely on the timing of learning. This impoverished description of childhood contrasts with the more wide-ranging changes observed across human development (e.g. the emergence of cognitive abilities such as theory of mind), as well as the notion of dedicated childhood adaptations being the direct products of evolution [30,75]. Further simulations could include additional genes that directly affect the life history of the participants in the experiment, for instance, nullifying the fitness effects of a certain number of initial trials to mimic a protected period of childhood where low-cost learning via 'play' is possible.

The use of adult participants in an experiment investigating childhood is also potentially problematic because the genuine psychology of children differs from that of adults and so its inclusion may produce different evolutionary dynamics. For instance, learning may not be as front-loaded if it is inefficient during early childhood. There are ways to work around this issue, although they are not simple. For instance, experimental manipulations could mask aspects of adult psychology in an attempt to mimic child psychology during the early trials. Alternatively, each agent could be

assigned both an adult and a child participant to make decisions at the appropriate stages of the agent's life.

One might ask, given the use of adult participants, whether our experiments speak more to periods of learning in adult life, as opposed to childhood learning specifically. Indeed, some general lessons may be drawn: if fitness depends on effective learning at any point in the life cycle of an organism, then a rich learning capacity may evolve, along with a memory capacity sufficient to store the information for as long as it is relevant. However, in our experiments, information collection was not possible prior to the first experimental trial; similarly, information use was not possible after the final experimental trial. As such, the hard boundaries of the experimental task correspond better to birth and death than they do to the start and end of a period of learning in adult life where information collected at some earlier time may be relevant, and the information learnt can continue to be used well into the future.

Another limitation of these experiments is that our implementation of the memory gene is an imperfect attempt at masking human memory. Because people can readily store information for far longer than the duration of our experiment, and because we were unable to force participants to forget specific items, the memory gene merely rendered useless information participants had stored beyond their imposed memory limit. Nonetheless, the information likely was retained and may have interfered with the acquisition of additional information. This illustrates a general limitation of this approach: while it can manipulate the experimental design to mask certain cognitive abilities, it cannot actually remove the abilities themselves.

A final limitation of the experiments here is the omission of many traits believed to be relevant to human evolutionary history. For instance, overall lifespan was held constant, and cultural inheritance between participants was not possible. The latter is likely important because it is believed that cultural evolution, plus its interactions with genetic evolution, have had significant consequences for human cognitive evolution [41,58,76,77] and further work could include communication between participants.

Across three experiments, we have presented data on the coevolution of learning, memory and childhood. Collectively, we find that the coevolution of cognition and life history is supported by human decision-making, but can be prevented by rapid environmental change. Crucially, we collected these data from experimental evolutionary simulations, a novel approach in which we can observe not only human behaviour, but also the evolutionary dynamics it produces. By including human psychology in the evolutionary process, such an approach has the potential to bridge the gap between theory and experiment thereby allowing us to draw firmer conclusions about human evolution.

Ethics. Recruitment and testing were approved by the Committee for Protection of Human Subjects at University of California, Berkeley (protocol ID 2015-12-8227).

Data accessibility. Data from all experiments are freely available: <https://github.com/thomasmorgan/phil-trans-2019-data>.

Authors' contributions. T.J.H.M. and T.L.G. designed the experiments. T.J.H.M. and J.W.S. created the experimental software. T.J.H.M. and J.W.S. created and executed the experiment. T.J.H.M. analysed the data. All authors wrote the manuscript.

Competing interests. We declare we have no competing interests.

Funding. This work was funded by NSF grant nos 1456709 and 1408652.

1. Hill K, Barton M, Hurtado AM. 2009 The emergence of human uniqueness: characters underlying behavioral modernity. *Evol. Anthropol.* **18**, 187–200. (doi:10.1002/evan.20224)
2. Vitousek PM, Mooney HA, Lubchenco J, Melillo JM. 1997 Human domination of Earth's ecosystems. *Science* **277**, 494–499. (doi:10.1007/978-0-387-73412-5_1)
3. Boyd R. 2018 *A different kind of animal*. Princeton, NJ: Princeton University Press.
4. Basalla G. 1988 *The evolution of technology*. Cambridge, UK: Cambridge University Press.
5. Roffman I, Savage-Rumbaugh S, Rubert-Pugh E, Ronen A, Nevo E. 2012 Stone tool production and utilization by bonobo-chimpanzees (*Pan paniscus*). *Proc. Natl Acad. Sci. USA* **109**, 14 500–14 503. (doi:10.1073/pnas.1212855109)
6. Whiten A, Goodall J, McGrew WC. 1999 Cultures in chimpanzees. *Nature* **399**, 15–18. (doi:10.1038/21415)
7. Whiten A, Erdal D. 2012 The human socio-cognitive niche and its evolutionary origins. *Phil. Trans. R. Soc. B* **367**, 2119–2129. (doi:10.1098/rstb.2012.0114)
8. Penn DC, Povinelli DJ. 2007 On the lack of evidence that non-human animals possess anything remotely resembling a 'theory of mind'. *Phil. Trans. R. Soc. B* **362**, 731–744. (doi:10.1098/rstb.2006.2023)
9. Robalino N, Robson A. 2012 The economic approach to 'theory of mind'. *Phil. Trans. R. Soc. B* **367**, 2224–2233. (doi:10.1098/rstb.2012.0124)
10. Heyes C. 2014 False belief in infancy: a fresh look. *Dev. Sci.* **17**, 647–659. (doi:10.1111/desc.12148)
11. Buchsbaum D, Gopnik A, Griffiths TL, Shafto P. 2011 Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition* **120**, 331–340. (doi:10.1016/j.cognition.2010.12.001)
12. Gardiner AK. 2014 Beyond irrelevant actions: understanding the role of intentionality in children's imitation of relevant actions. *J. Exp. Child Psychol.* **119**, 54–72. (doi:10.1016/j.jecp.2013.10.008)
13. Richerson PJ, Boyd R. 2005 *Not by genes alone: how culture transformed human evolution*. Chicago, IL: University Of Chicago Press.
14. Rendell L, Whitehead H. 2001 Culture in whales and dolphins. *Behav. Brain Sci.* **24**, 309. (doi:10.1017/S0140525X0100396X)
15. Whiten A. 2011 The scope of culture in chimpanzees, humans and ancestral apes. *Phil. Trans. R. Soc. B* **366**, 997–1007. (doi:10.1098/rstb.2010.0334)
16. Laland KN, Galef Jr BG. 2009 *The question of animal culture*. Cambridge, MA: Harvard University Press.
17. Promislow DEL, Harvey PH. 1990 Living fast and dying young: a comparative analysis of life-history variation among mammals. *J. Zool.* **220**, 417–437. (doi:10.1111/j.1469-7998.1990.tb04316.x)
18. Kaplan H. 1997 The evolution of the human life course. In *Between zeus and the salmon: the biodemography of longevity* (eds KW Wachter, CE Finch), Washington, DC: National Academy Press.
19. Kaplan H, Hill K, Lancaster J, Hurtado AM. 2000 A theory of human life history evolution: diet, intelligence, and longevity. *Evol. Anthropol.* **9**, 156–185. (doi:10.1002/1520-6505(2000)9:4<156::AID-EVANS>3.0.CO;2-7)
20. Thompson ME, Ellison PT. 2017 Fertility and fecundity. In *Chimpanzees and human evolution* (eds MN Muller, RW Wrangham, DR Pilbeam). Harvard, MA: The Belknap Press of Harvard University.
21. Hawkes K, O'Connell JF, Blurton Jones NG. 1997 Hadza women's time allocation, offspring provisioning, and the evolution of long postmenopausal life spans. *Curr. Anthropol.* **38**, 551–577. (doi:10.1086/204646)
22. Hawkes K, Coxworth JE. 2013 Grandmothers and the evolution of human longevity: a review of findings and future directions. *Evol. Anthropol.* **22**, 294–302. (doi:10.1002/evan.21382)
23. Gurven MD, Davison RJ, Kraft TS. 2020 The optimal timing of teaching and learning across the life course. *Phil. Trans. R. Soc. B* **375**, 20190500. (doi:10.1098/rstb.2019.0500)
24. Richerson PJ, Boyd R. 2020 The human life history is adapted to exploit the adaptive advantages of culture. *Phil. Trans. R. Soc. B* **375**, 20190498. (doi:10.1098/rstb.2019.0498)
25. Jones NB, Marlowe FW. 2002 Selection for delayed maturity: does it take 20 years to learn to hunt and gather? *Hum. Nat.* **13**, 199–238. (doi:10.1007/s12110-002-1008-3)
26. Deffner D, McElreath R. 2020 The importance of life history and population regulation for the evolution of social learning. *Phil. Trans. R. Soc. B* **375**, 20190492. (doi:10.1098/rstb.2019.0492)
27. Kline MA, Boyd R, Henrich J. 2013 Teaching and the life history of cultural transmission in Fijian villages. *Hum. Nat.* **24**, 351–374. (doi:10.1007/s12110-013-9180-1)
28. Henrich J, Henrich N. 2010 The evolution of cultural adaptations: Fijian food taboos protect against dangerous marine toxins. *Proc. R. Soc. B* **277**, 3715–3724. (doi:10.1098/rspb.2010.1191)
29. Kerr B, Feldman MW. 2003 Carving the cognitive niche: optimal learning strategies in homogeneous and heterogeneous environments. *J. Theor. Biol.* **220**, 169–188. (doi:10.1006/jtbi.2003.3146)
30. Buchsbaum D, Bridgers S, Skolnick Weisberg D, Gopnik A. 2012 The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Phil. Trans. R. Soc. B* **367**, 2202–2212. (doi:10.1098/rstb.2012.0122)
31. Rogers AR. 1988 Does biology constrain culture. *Am. Anthropol.* **90**, 819–831. (doi:10.1525/aa.1988.90.4.02a00030)
32. Caldwell CA, Millen AE. 2008 Experimental models for testing hypotheses about cumulative cultural evolution. *Evol. Hum. Behav.* **29**, 165–171. (doi:10.1016/j.evolhumbehav.2007.12.001)
33. Bartlett FC. 1932 *Remembering: a study in experimental and social psychology*. Cambridge, UK: Cambridge University Press.
34. Morgan TJH *et al.* 2015 Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nat. Commun.* **6**, 1–8. (doi:10.1038/ncomms7029)
35. Mesoudi A, O'Brien MJ. 2008 The cultural transmission of Great Basin projectile-point technology I: an experimental simulation. *Am. Antiq.* **73**, 3–28. (doi:10.1017/S0002731600041263)
36. Miu E, Gulley N, Laland KN, Rendell L. 2018 Innovation and cumulative culture through tweaks and leaps in online programming contests. *Nat. Commun.* **9**, 1–8. (doi:10.1038/s41467-018-04494-0)
37. Jacobs RC, Campbell DT. 1961 The perpetuation of an arbitrary tradition through several generations of a laboratory microculture. *J. Abnorm. Soc. Psychol.* **62**, 649–658.
38. Mery F, Kawecki TJ. 2002 Experimental evolution of learning ability in fruit flies. *Proc. Natl Acad. Sci. USA* **99**, 14 274–14 279. (doi:10.1073/pnas.222371199)
39. Lenski RE, Travisano M. 1994 Dynamics of adaptation and diversification: a 10 000-generation experiment with bacterial populations. *Proc. Natl Acad. Sci. USA* **91**, 6808–6814. (doi:10.1073/pnas.91.15.6808)
40. Servodio MR, Brandvain Y, Dhole S, Fitzpatrick CL, Goldberg EE, Stern CA, Van Cleve J, Yeh DJ. 2014 Not just a theory—the utility of mathematical models in evolutionary biology. *PLoS Biol.* **12**, 1–5. (doi:10.1371/journal.pbio.1002017)
41. Morgan TJH. 2016 Testing the cognitive and cultural niche theories of human evolution. *Curr. Anthropol.* **57**, 370–377. (doi:10.1086/686531)
42. McNamara JM, Barta Z, Houston AI. 2004 Variation in behaviour promotes cooperation in the Prisoner's Dilemma game. *Nature* **428**, 745–748. (doi:10.1038/nature02432)
43. McNamara JM, Leimar O. 2010 Variation and the response to variation as a basis for successful cooperation. *Phil. Trans. R. Soc. B* **365**, 2627–2633. (doi:10.1098/rstb.2010.0159)
44. Dall SRX, Houston AI, McNamara JM. 2004 The behavioural ecology of personality: consistent individual differences from an adaptive perspective. *Ecol. Lett.* **7**, 734–739. (doi:10.1111/j.1461-0248.2004.00618.x)
45. Henrich J, Heine SJ, Norenzayan A. 2010 Most people are not WEIRD. *Nature* **466**, 29. (doi:10.1017/S0140525X0999152X)
46. Stewart N, Ungemach C, Harris AJL, Bartels DM, Paolacci G, Chandler J. 2015 The average laboratory samples a population of 7,300 Amazon

- mechanical Turk workers. *Judgm. Decis. Mak.* **10**, 479–491.
47. Difallah D, Filatova E, Ipeirotis P. 2018 Demographics and dynamics of mechanical Turk workers. In *WSDM 2018—Proc. 11th ACM Int. Conf. Web Search Data Mining, February 2018, Marina Del Rey, CA, USA*, pp. 135–143.
 48. Ipeirotis PG. 2010 Demographics of Mechanical Turk. NYU Working Paper No. CEDER-10-01.
 49. Arechar AA, Kraft-Todd G, Rand DG. 2017 Turking overtime: how participant characteristics and behavior vary over time and day on Amazon mechanical Turk. *J. Econ. Sci. Assoc.* **3**, 1–11. (doi:10.1007/s40881-017-0035-0)
 50. Buhrmester M, Kwang T, Gosling SD. 2011 Amazon's mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspect. Psychol. Sci.* **6**, 3–5. (doi:10.1177/1745691610393980)
 51. Ross J, Zaldivar A, Irani L, Tomlinson B. 2010 Who are the Turkers? Worker demographics in Amazon mechanical Turk. *Chi Ea* **2010**, 2863–2872. (doi:10.1145/1753846.1753873)
 52. Levay KE, Freese J, Druckman JN. 2016 The demographic and political composition of mechanical Turk samples. *SAGE Open* **6**. (doi:10.1177/2158244016636433)
 53. Goodman JK, Cryder CE, Cheema A. 2013 Data collection in a flat world: the strengths and weaknesses of mechanical Turk samples. *J. Behav. Decis. Mak.* **26**, 213–224. (doi:10.1002/bdm.1753)
 54. Plummer M, Stukalov A, Denwood M. 2016 *rjags: Bayesian Graphical Models using MCMC*. <https://cran.r-project.org/web/packages/rjags/rjags.pdf>.
 55. Navarro DJ, Newell BR, Schulze C. 2016 Learning and choosing in an uncertain world: an investigation of the explore-exploit dilemma in static and dynamic environments. *Cogn. Psychol.* **85**, 43–77. (doi:10.1016/j.cogpsych.2016.01.001)
 56. Pinker S. 2010 The cognitive niche: coevolution of intelligence, sociality, and language. *Proc. Natl Acad. Sci. USA* **107**, 8993–8999. (doi:10.1073/pnas.0914630107)
 57. Barrett HC, Cosmides L, Tooby J. 2007 The hominid entry into the cognitive niche. In *The evolution of mind: fundamental questions and controversies* (eds SW Gangestad, JA Simpson), pp. 241–248. New York, NY: The Guilford Press.
 58. Boyd R, Richerson PJ, Henrich J. 2011 The cultural niche: why social learning is essential for human adaptation. *Proc. Natl Acad. Sci. USA* **108**, 10 918–10 925. (doi:10.1073/pnas.1100290108)
 59. Henrich J. 2016 *The secret of our success*. Princeton, NJ: Princeton University Press.
 60. Muthukrishna M, Henrich J. 2016 Innovation in the collective brain. *Phil. Trans. R. Soc. B* **371**, 20150192. (doi:10.1098/rstb.2015.0192)
 61. Laland KN. 2017 *Darwin's unfinished symphony*. Princeton, NJ: Princeton University Press.
 62. Kaplan HS, Robson AJ. 2002 The emergence of humans: the coevolution of intelligence and longevity with intergenerational transfers. *Proc. Natl Acad. Sci. USA* **99**, 10 221–10 226. (doi:10.1073/pnas.152502899)
 63. Kaplan HS, Lancaster JB, Robson A. 2003 *Embodied capital and the evolutionary economics of the human life span. In Life span: evolutionary, ecological, and demographic perspectives* (eds JR Carey, S Tuljapurkar). New York, NY: The Population Council, Inc.
 64. Aoki K, Wakano JY, Lehmann L. 2012 Evolutionarily stable learning schedules and cumulative culture in discrete generation models. *Theor. Popul. Biol.* **81**, 300–309. (doi:10.1016/j.tpb.2012.01.006)
 65. Lehmann L, Wakano JY, Aoki K. 2013 On optimal learning schedules and the marginal value of cumulative cultural evolution. *Evolution* **67**, 1435–1445. (doi:10.1111/evo.12040)
 66. Schuppli C, Graber SM, Isler K, van Schaik CP. 2016 Life history, cognition and the evolution of complex foraging niches. *J. Hum. Evol.* **92**, 91–100. (doi:10.1016/j.jhevol.2015.11.007)
 67. Street SE, Navarrete AF, Reader SM, Laland KN. 2017 Coevolution of cultural intelligence, extended life history, sociality, and brain size in primates. *Proc. Natl Acad. Sci. USA* **114**, 7908–7914. (doi:10.1073/pnas.1620734114)
 68. Hodell DA, Channell JET. 2016 Mode transitions in Northern Hemisphere glaciation: co-evolution of millennial and orbital variability in Quaternary climate. *Clim. Past* **12**, 1805–1828. (doi:10.5194/cp-12-1805-2016)
 69. Johnsen SJ *et al.* 1992 Irregular glacial interstadials recorded in a new Greenland ice core. *Nature* **359**, 311–313. (doi:10.1038/359311a0)
 70. Richerson PJ, Boyd R. 2013 Rethinking paleoanthropology: a world queerer than we supposed. In *Evolution of mind, brain, and culture* (eds G Hatfield, H Pittman), pp. 263–302. Philadelphia, PA: University of Pennsylvania Press.
 71. Martrat B *et al.* 2004 Abrupt temperature changes in the Western Mediterranean over the past 250,000 years. *Science* **306**, 1762–1765. (doi:10.1126/science.1101706)
 72. Martrat B, Grimalt JO, Shackleton NJ, de Abreu L, Hutterli MA, Stocker TF. 2007 Four climate cycles of recurring deep and surface water destabilizations on the Iberian margin. *Science* **317**, 502–507. (doi:10.1126/science.1229223)
 73. Fagan B. 2001 *The Little Ice age: how climate made history 1300–1850*. New York, NY: Basic Books.
 74. Claidière N, Smith K, Kirby S, Fagot J. 2014 Cultural evolution of systematically structured behaviour in a non-human primate. *Proc. R. Soc. B* **281**, 20141541. (doi:10.1098/rspb.2014.1541)
 75. Csibra G, Gergely G. 2011 Natural pedagogy as evolutionary adaptation. *Phil. Trans. R. Soc. B* **366**, 1149–1157. (doi:10.1098/rstb.2010.0319)
 76. Heyes C. 2018 *Cognitive gadgets*. Cambridge, MA: Harvard University Press.
 77. Laland KN, Odling-Smee J, Myles S. 2010 How culture shaped the human genome: bringing genetics and the human sciences together. *Nat. Rev. Genet.* **11**, 137–148. (doi:10.1038/Nrg2734)