Received: 16 November 2021 | Accepted: 12 June 2023

DOI: 10.1111/aman.13926

AMERICAN ANTHROPOLOGIST

RESEARCH ARTICLE

Evolving payoff currencies through the construction of causal theories

Ze Hong[1] | Joseph Henrich[2]

[1]Department of Sociology, University of Macau, Taipa, Macau, China

[2]Department of Human Evolutionary Biology, Harvard University, Cambridge, MA, United States

Correspondence
Ze Hong, Department of Sociology, University of Macau, Taipa, Macau, China.
Email: zehong@um.edu.mo

Abstract

Payoff-biased cultural learning has been extensively discussed in the literature on cultural evolution, but where do payoff currencies come from in the first place? Are they products of genetic or cultural evolution? Here we present a simulation model to explore the possibility of novel payoff currencies emerging through a process of theory construction, where agents come up with "channels" via which different cultural traits contribute to some ultimate payoff and use such "channels" as intermediate payoff currencies to make trait-updating decisions. We show that theory-building as a strategy is mostly favored when the noise associated with the ultimate-level payoff is high, selective pressures are strong, and the probability of arriving at the right theory is high. This approach provides insights into both the emergence of payoff currencies and the role of cognition for causal model building. We close by discussing the implications of our model for the broader question of causal learning in social contexts.

KEYWORDS
payoff-biased imitation, theory-building, cultural evolution

Resumen

El aprendizaje cultural sesgado en términos de recompensa ha sido extensivamente discutido en la literatura de la evolución cultural, ¿pero de dónde vinieron los valores de compra en primer lugar? ¿Son ellos productos de evolución genética o cultural? Aquí presentamos un modelo de simulación para explorar la posibilidad de unos valores de recompensa novedosos emergiendo a través de un proceso de construcción de una teoría, donde los agentes idearon "canales" a través de los cuales características culturales diferentes contribuyen a un último pago y usan tales "canales" como formas de pago intermedio para tomar decisiones de actualización de características. Mostramos que la construcción de la teoría como una estrategia es en su mayoría favorecida cuando el ruido asociado con el último nivel de pago es alto, las presiones selectivas son altas y la probabilidad de llegar a la teoría correcta es alta. Esta aproximación provee conocimiento tanto en la emergencia de medios de pago como en el rol de la cognición en la construcción de modelos causales. Concluimos discutiendo las implicaciones de nuestro modelo para la cuestión más amplia del aprendizaje causal en contextos sociales. [imitación sesgada pagada, construcción de teoría, evolución cultural]

INTRODUCTION

Two important and interrelated behavioral features that distinguish humans from other animals are extensive social learning across diverse behavioral domains (Herrmann et al., 2007) and substantial cumulative culture (Dean et al., 2014; Henrich and Tennie, 2017; Tennie, Call, and Tomasello, 2009). To account for the origin and evolution of these features, researchers have pointed to the evolution of various adaptive learning mechanisms (Henrich and McElreath, 2003; R. Kendal et al., 2018). Crucially, human culture is not only cumulative but also adaptive, in the sense that the outcomes of evolved learning biases should generally be fitness-enhancing (Richerson and Boyd, 2005).

Adaptive cultural learning requires that individuals acquire traits that on average increase their fitness. This means that processes that do not involve direct fitness evaluations, such as conformist transmission, which can be adaptive in spatially heterogeneous environments (Henrich and Boyd, 1998; Nakahashi, Wakano, and Henrich, 2012), and when observational errors are common (Henrich and Boyd, 2002), are not going to effectively spread novel adaptive cultural traits by themselves, and can sometimes even lead to population collapse under specific ecological conditions (Whitehead and Richerson, 2009). However, adaptive cultural evolution occurs when a combination of learning strategies is used, such as: (1) conformist learning combined with individual learning (Boyd and Richerson, 1995; Henrich and Boyd, 1998) or payoff-biased learning (Henrich and Boyd, 2002; Hong, 2022b); (2) "flexible learning" that enables individual learning to be more accurate or less costly (Boyd and Richerson, 1995); (3) "critical social learning," where individuals engage in social learning first, and if results are unsatisfactory, then deploy individual learning (Enquist, Eriksson, and Ghirlanda, 2007); and (4) "specialized hybrid learning," in which individuals deploy asocial learning in some contexts and social learning in other contexts (Boyd and Richerson, 1995; Kharratzadeh et al., 2017).

In contrast, purely payoff-biased cultural learning appears to be an attractive and parsimonious strategy because naïve learners can directly assess the payoff of specific behavioral variants and adopt the ones with higher payoff (Nakahashi, Wakano, and Henrich, 2012; Vale et al., 2017). However, this raises an important question: How do naïve learners assess payoffs? Most published work has not fully interrogated this problem. In genetic evolutionary models, payoff is often assumed to be some currency that correlates with genetic fitness (Boyd and Richerson, 1985; J. Kendal, Giraldeau, and Laland, 2009), and in experimental studies, payoff is often represented by some quantity that individuals with particular preference and cultural background care about, such as the caloric efficiency of hunting, money, reward tokens (Vale et al., 2017), or food items (Barrett, McElreath, and Perry, 2017). The more fundamental question is thus: What are the payoff currencies that humans have genetically and culturally evolved to care about? Food, for example, may be a currency that is evolutionarily significant, but both money and reward tokens are relatively recent human cultural inventions, and it seems unlikely that we have genetically evolved psychological tendencies to detect and react to them (Durham, 1991).

A large literature on cultural learning has examined the kinds of cues that learners use to figure out the who, what, and when of cultural learning, and there are reasons to suspect that some of these reliably emerge over human development and across societies (Chudek et al., 2012). Researchers have extensively examined model-based cues, such as competence, status (Henrich and Gil-White, 2001; Wood, Kendal, and Flynn, 2013), health, number of children (given age), family size, and physical well-being (Henrich, 2016). The underlying logic is that because these cues are statistically associated with higher fitness, learners who use these cues to acquire traits socially enjoy a fitness advantage, and therefore the psychological tendency of attending to these cues was selected for during our evolutionary history. These cues may be rather general, in the sense that they do not specify what traits to adopt, so individuals with high cue values are copied across a broad range of domains. However, such cues are often enriched by content-specific mechanisms that bias the attention of learners in fitness-relevant domains like food, sex, and gossip (reputation). The emergence and evolution of more fine-grained, culturally specific currencies, on the other hand, has gathered less attention.

In this article, we use the term "payoff currency" in a rather loose manner to denote the traits or characteristics that individuals value and use as proxies to make learning decisions. For example, hunting returns may have been an important payoff currency for much of our species' evolutionary time. Good hunters are typically admired and learned from, and they may benefit from such cultural learning as they accumulate prestige (Henrich and Gil-White, 2001). A notable feature of learning using specific, culturally learned currencies is that learners often need to associate payoff currencies with particular cultural variants. In other words, naïve individuals need not only evaluate payoffs from potential models but also decide what traits are relevant for particular payoffs. In the above example, if we are making an arrow with the goal of maximizing hunting return, should we copy the arrow design from the hunter who hits the target most frequently or the hunter who is the strongest? Alternatively, we could "play it safe" and copy the arrow design from the hunter who brings the most meat back to camp since hunting return is what we ultimately care about. Here, in associating aiming accuracy and arrow design, we are effectively constructing causal theories regarding *how* arrow design contributes to hunting return and are using aiming accuracy as an *intermediate* payoff currency to decide what arrow design to copy. Thus, once the causal theory is in place, individuals can bypass the ultimate payoff (hunting return) and assess the intermediate payoff (aiming accuracy) directly.

An obvious advantage of using intermediate payoff currencies to make learning decisions is that they may be less affected by noise compared to that in the ultimate payoff. For example, hunting returns in small-scale societies are subject to much uncertainly (Hill and Kintigh, 2009), and field anthropologists often use specific metrics, such as physical strength and aiming accuracy, to evaluate hunting ability (Stibbard-Hawkes, Attenborough, and Marlowe, 2018). However, constructing causal theories and using intermediate payoffs may not always be the best strategy for at least two reasons. First, constructing causal theories may be cognitively expensive, thus absorbing time and energy resulting in opportunity costs;

second, humans may not always be able to construct the right causal theory. In reality, figuring out the correct causal model can be difficult. What leads to hunting success? How does one stay healthy? How did our neighbor have so many children? These are hard problems that individuals may fail to solve during their lifetime. In fact, one important goal of modern science is to discover the mechanisms of natural and social phenomena (Hitchcock, 1992), and scientists still have trouble pinpointing the precise causal pathways despite our unprecedented measurement technologies and sophisticated statistical methods. Therefore, it is useful to explore the conditions under which theory construction and the use of intermediate payoffs are advantageous. In this article, we aim to address the following questions:

(1) When and how do intermediate payoff currencies arise via a theory-construction process in our efforts to elicit the relationship between outcomes of interest and relevant cultural variants?
(2) What are the conditions for theory construction to be a superior strategy compared to no theory construction? Once theories are formulated and transmitted, how do they stand up to individual observations over many trials over time?

To do so, we employ a formal modeling approach. In the following sections, we will describe a computational model, present the simulation results from two studies, and discuss the evolutionary implications for "theory-building" as a genetically evolved strategy as well as the possibility of culturally transmitted payoff currencies emerging from the theory-building process.

BASELINE MODEL

We build an agent-based simulation to model the temporal evolution of some ultimate payoff and the associated culturally transmitted trait values, with the exact associations specified under different "theory-building" conditions. We also consider "theory-building" as a genetically transmitted strategy and examine the conditions under which it may be favored by natural selection. Of course, this is not to say that there is no cultural input in individuals' propensity to construct theories: like most human traits, such habit of reasoning about the causal structure of the world surely requires some cultural conditioning. Rather, in this exercise we focus on the genetic component of "theory-building" in order to explore how natural selection at the genetic level may have shaped our "theory-building" propensity in the long run.

Life Cycle of Agents

We assume asexual reproduction and nonoverlapping generations, as is commonly done in evolutionary modeling (Boyd and Richerson, 1985; Day and Bonduriansky, 2011). In each generation, every agent inherits values for x_1 and x_2 from a parent and samples n individuals in the parental generation and randomly draws a sample from their parental generation (potential models) and selects the one with the largest relevant payoff as their model. The learner then compares the model's payoff with their own payoff: if the model's payoff is smaller than their own payoff, then no copying will occur; otherwise (the model's payoff larger than their own payoff), they will seek to copy the model's cultural trait, x_1 and x_2, but imperfect inference and copy errors will be introduced. Agents will then produce an offspring with the same x_1 and x_2 values they acquired through this process.

Payoff Setup

Suppose F, the "ultimate payoff," is a genetically evolved payoff currency that individuals use to evaluate and update cultural traits, here captured by x_1 and x_2. Why individuals choose to associate x_1 and x_2 instead of other traits with F has to do with evolved intuition and cultural transmission, the exploration of which lies beyond the scope of our current effort. In the most general form,

$$F = f(x_1, x_2) + \varepsilon \tag{1}$$

where x_1 and x_2 represent cultural traits that contributes to payoff F via function f, and ε represents factors other than x_1 and x_2 (including intrinsic stochasticity) that contribute to F. For mathematical convenience, we assume both x_1 and x_2 can be observed without error, and the ε is a normally distributed error term (hereafter referred to as "noise") with mean 0 and variance σ^2. To add some concreteness, assume that humans have a genetically evolved preference for better hunting returns (i.e., more meat beats less meat). Here, F would be average hunting returns per week (the amount of food a hunter brings back to camp), x_1 would represent arrow length, x_2 the time in early morning that a hunter departs, and ε captures stochastic, external factors that humans have no control over such as weather. Intuitively, equation (1) could capture how having an arrow closer to the right arrow length and departing at a time closer to the optimal contribute to hunting success.

To further analyze the evolutionary dynamics of the system, we need to specify the function f. Assuming that both x_1 and $x_2$contribute[1] to F in the form of a negative quadratic function, equation (1) becomes

$$\underbrace{F}_{\substack{\text{hunting}\\\text{return}}} = \underbrace{\beta_1 \cdot [(a - x_1) \cdot x_1]}_{\substack{\text{contribution due to}\\\text{arrow design}}} + \underbrace{\beta_2 \cdot [(b - x_2) \cdot x_2]}_{\substack{\text{contribution due to}\\\text{departure design}}} + \underbrace{\varepsilon}_{\text{externalfactors}} \tag{2}$$

We choose the negative quadratic function in our model to capture the intuition that values of x_1 and x_2 that maximize F are often not the extreme values (for example, the best arrow is neither too long nor too short). Here, the β coefficients are the relative contribution of different cultural traits to the ultimate payoff F, whose maximum value is obtained when $x_1 = \frac{a}{2}$ and $x_2 = \frac{b}{2}$ given that both $\beta_1$and β_2 are positive. In the simulation, both β coefficients as well as are a and b are set to 1 for simplicity, and thus equation (2) can be rewritten as

$$F = (1 - x_1) \cdot x_1 + (1 - x_2) \cdot x_2 + \varepsilon \tag{3}$$

Because the mean of ε is 0, F is maximized when $x_1 = 0.5$ and $x_2 = 0.5$, and the maximal expected value of F is $0.5 \cdot 0.5 + 0.5 \cdot 0.5 = 0.5$.

Intermediate Payoff as a Result of Causal Theory Construction

Humans are exceptionally good at detecting patterns and identifying correlations (Ayton and Fischer, 2004), even when observed patterns were produced by a random process (Fyfe et al., 2008). At the individual level, people may discover that arrow length co-varies with some other observable variable—aiming accuracy, which in turn co-varies with hunting return. Similarly, departing time may be discovered to co-vary with number of prey encounters, which predicts hunting returns as well. It is worth noting the identification of chains of correlations can also be biased by evolved or transmitted folk intuitions about the causal structure of the world (Carey, 2009). Regardless of the mechanisms, in discovering these *intermediate* variables, people effectively construct causal theories on how various cultural traits contribute to ultimate payoff. These intermediate variables (represented by R and T throughout this article) can then be used as payoff currencies to make trait-adoption decisions: instead of copying x_1 and x_2 values from the model with the largest F, naïve individuals adopt x_1 and x_2 separately from models with high values of respective intermediate payoffs. We consider three different theory conditions.

No Theory Condition

As previously stated, the "no theory" condition refers to a nondiscriminative copying strategy. There has been both theoretical reasoning (Henrich, 2016) and experimental studies (Derex et al., 2019) showing that explicit theory in the form of causal models is not always necessary for adaptive cultural evolution. A lot of human cultural know-how and technologies are causally opaque, and people often cannot articulate why they engage in certain cultural practices that are in fact fitness-enhancing (Henrich and Henrich, 2010; Placek, Madhivanan, and Hagen, 2017). For example, Henrich and Henrich (2010) found that many Fijian food taboos for pregnant women protect them from marine toxins, yet the locals often fail to explicitly specify the causal connection. Developmental research also shows that unlike chimpanzees, human children often copy causally irrelevant actions to achieve tasks in experimental settings (Lyons, Young, and Keil, 2007; Whiten et al., 2009). In our simulation, we model this condition by having agents update both x_1 and x_2 based on F. In other words, individuals are not assuming how x_1 and x_2 contribute to F; they simply notice some individuals possessing higher F values and copy their x_1 and x_2 values. In the concrete example, people pick the best hunter based on his hunting return and copy both his arrow design and dietary habits.

Correct Theory Condition

This is the ideal situation, where agents are able to come up with the causal theory that matches reality exactly. Recall that in our simulation setup

$$F = (1 - x_1) \cdot x_1 + (1 - x_2) \cdot x_2 + \varepsilon \tag{4}$$

which can be rewritten as $F = R + T + \epsilon$, and the intermediate payoffs are thus $R = (1 - x_1) \cdot x_1$ and $T = (1 - x_2) \cdot x_2$. Individuals will then adopt x_1 from models with high intermediate payoff, R, and x_2 from models with high intermediate payoff, T. Concretely, this is to say that having correctly identified the causal pathways (aiming accuracy and physical strength), people are copying arrow length from the person who aims better and departure times in the morning from the person who encounters more prey that day. Note that in this way individuals completely bypass the noise factor ε, because the construction of theory provides more precise guidance regarding what to look for when deciding who to copy.

Incorrect or Incomplete Theory Condition

The constructed causal theory can be wrong (in the sense that it does not match reality) in many ways, and some theories are more wrong than others as they deviate further from the real causal model. We explore three ways for a causal theory to be wrong: (1) including irrelevant currencies, (2) missing relevant currencies, and (3) misconstructing intermediate payoffs, which means that maximization of intermediate payoff does not correspond to maximization of ultimate payoff. Here we discuss three possibilities as follows:

(1) Irrelevant causal currency:

Suppose all agents have an additional cultural trait x_3 that doesn't contribute to F. In addition to adopting x_1 and x_2 values, agents also adopt x_3 from the model with the highest F at no additional cost. With respect to evolutionary change in x_1, x_2, and F, this condition is *exactly the same* as the "no theory" condition and will not be explicitly modeled. Future work may fully explore this situation by including additional factors such as the cost of copying a particular trait.

(2) Missing relevant currencies:

In this scenario, agents have the right theory regarding x_1 ($R = (1 - x_1) \cdot x_1$) but ignore x_2 when updating traits. The overall causal model thus only includes x_1, and agents only adopt x_1 values from models with high values of R. Note that in this situation agents also avoid the noise factor ε, but not including x_2 means F will not reach its maximal value and x_2 will freely drift around its initial values.

(3) Misconstructing intermediate payoffs:

In our simple model, either intermediate payoff (R or T) may be inaccurately constructed. For simplicity, we only consider the misconstruction of one intermediate payoff. Suppose the causal model constructed is

$$F = R' + T + \varepsilon \tag{5}$$

where $R' = (a' - x_1) \cdot x_1$ and $T = (1 - x_2) \cdot x_2$. Here R' and T are the respective intermediate payoffs. While agents correctly identify T as one of the intermediate payoffs, R' is different from how x_1 contribute to F in reality ($a' \neq 1$; in the simulation $a' = 0.6$). Relating back to our hunting example, suppose people think that shooting speed (instead of aiming accuracy) is the channel via which arrow design contribute to hunting success and use that as the intermediate payoff. The arrow design that maximizes shooting speed does not necessarily maximize aiming accuracy, and since aiming accuracy is the real mechanism, using the misconstructed intermediate payoff will lead to suboptimal arrow design and thus reduced hunting success.

STUDY 1: THE CULTURAL EVOLUTIONARY TRAJECTORY OF TRAIT VALUES

The main results of our simulation are shown in Figure 1, which largely confirm the verbal argument that using intermediate payoffs via theory construction can significantly speed up adaptive change. In the correct theory condition, both cultural trait values (x_1 and x_2) reach their optimal value much faster than in the "no theory" condition, especially when the noise in the ultimate payoff F is large. For example, when the variance σ of the error term in F is 1 (most of the variation in F is explained by noise), agents in the "no theory" condition reach the optimal value (0.5 in benchmark parameter settings) after roughly 300 generations, whereas agents in "right theory" condition reach optimal value in 3 generations—so 100 times faster.

However, when agents possess inaccurate causal models, we don't see convergence to the optimal values for x_1 and x_2. Specifically, when agents possess an incomplete causal model, traits rapidly stabilize at nonoptimal values of F. By contrast, when causal models are missing relevant variables, cultural traits do not change much around their initial value (0.8 for x_1, marked in blue) and as a result the population mean payoff F does not reach its optimum either. Notice that the evolutionary trajectory of the "correct theory" and the two "incorrect theory" conditions are not affected by

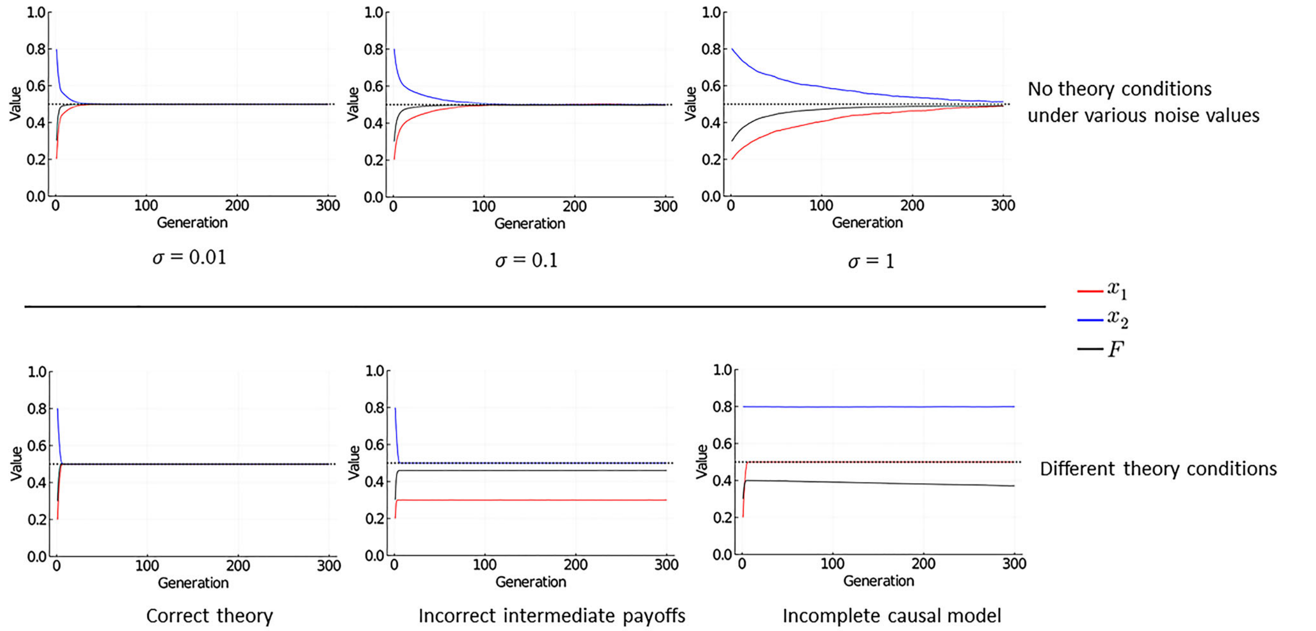

FIGURE 1 The evolutionary trajectory of two cultural trait values x_1 (red line) and x_2 (blue line) and the corresponding ultimate payoff F (black line) in 300 generations across different theory and noise conditions. Agent number (N) = 1000, and sample size $n = 5$. The starting values are 0.2 for x_1 (red) and 0.8 for x_2 (blue). σ represents the standard deviation of the normally distributed error term in F. In all simulation runs the optimal x_1 and x_2 value is 0.5, and the maximum F value is 0.5 as well. [This figure appears in color in the online issue]

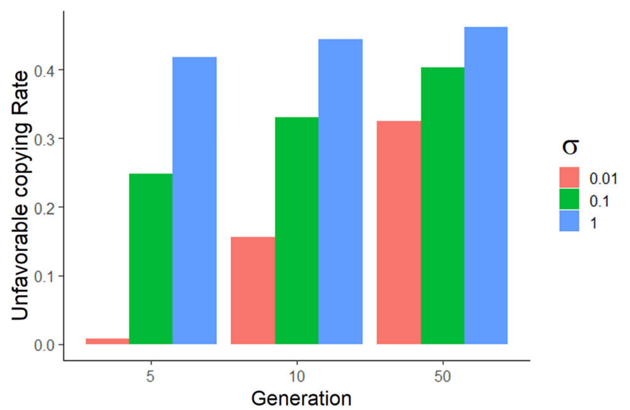

FIGURE 2 Unfavorable copying rate in "no theory" condition across different noise magnitudes at various generations. Unfavorable copying rate is calculated as the proportion of times an agent copied some model's trait values that yield a lower ultimate payoff F. Plotted rates are average values of 1000 simulation replications. Agent number (N) = 1000, and sample size $n = 5$. [This figure appears in color in the online issue]

noise at all, because by construction the noise factor is absent. In reality, intermediate payoff may also be noisy, yet the magnitude of their noise is likely to be much smaller than the overall noise involved in F.

The presence of noise makes it more likely for individuals to adopt trait values from models who appear successful due to chance. As a result, the trait values that naïve agents adopt can sometimes yield lower ultimate payoffs than their original trait values. Here we refer to instances where agents adopt trait values that lead to lower ultimate payoffs (in F) as "unfavorable copying." As shown in Figure 2, in the "no theory" condition, the magnitude of noise significantly influences unfavorable copying rates, especially in early generations when variation in ultimate payoffs (F) due to the two cultural traits is substantial. At generation 5, for example, over 40% of agents' trait-adoption decisions are "mistakes" when noise is high ($\sigma = 1$), compared to less than 1% when noise is low ($\sigma = 0.01$). In the "right theory" setting, in contrast, the unfavorable copying rate is 0 because noise is absent (not shown in graph). Unfavorable copying rates in various "wrong theory" conditions depend on how the constructed theory differs from reality. For example, in the "misconstruction of intermediate payoff" condition the unfavorable copying rate would be 0 for x_2 (as the intermediate payoff for x_2 is correctly identified) and roughly 50% for x_1 once all the agents' x_1 values reach the "false optimal" at 0.3 ($R' = (0.6 - x_1) \cdot x_1$).

TABLE 1 The fixation rate of "no theory" and "theory" strategies in finite populations by 300 generations. Values in cells are the ratio of proportion of time "theory" strategy reaches fixation over proportion of time "no theory" strategy reaches fixation. Simulation is repeated 500 times, with the initial population consisting of 250 individuals with "theory" strategies and 250 individuals with "no theory" strategies (Agent number (N) = 500, and sample size $n = 5$). Note that "strong," "medium," and "weak" here refer to the relative contribution of F to the overall fitness.

	Baseline Fitness					
	0 (strong selection)		*1 (medium selection)*		*10 (weak selection)*	
Magnitude of noise (σ)	*No theory*	*Theory*	*No theory*	*Theory*	*No theory*	*Theory*
1	*0.02*	*0.35*	*0.08*	*0.21*	*0.11*	*0.09*
0.1	*0.04*	*0.18*	*0.07*	*0.16*	*0.13*	*0.12*
0.01	*0.1*	*0.16*	*0.07*	*0.1*	*0.09*	*0.12*

Our simulation results show that the average ultimate payoff in the long run is the same for "no theory" and "right theory" conditions in homogenous populations. So when does "right theory" confer additional benefits over "no theory"? We suggest three possibilities. First, human perceptions are imperfect and can only notice stimulus difference above a certain threshold (Thurstone, 2017), which, when combined with stochasticity in payoff signals, may make detection of small technological improvement difficult or impossible.

Second, the "no theory" strategy almost always includes copying irrelevant traits. Though in our simulation this is modeled as costless, in realistic settings copying certain traits often incurs some nontrivial cost. For example, copying someone's clothing choice means one has to spend time and effort to make/buy new clothes. In principle, the more irrelevant traits one copies, the more cost one pays on average. The "theory" strategy, on the other hand, does not suffer this cost as much because by constructing theory and intermediate payoffs, individuals can be more precise regarding what to copy, which often means copying fewer traits.

Lastly, strategies that enable faster adaptive change may confer an evolutionary advantage in finite populations. Even though the long-term equilibrium outcome is the same, individuals with the "correct theory" reach the optimal trait value faster and thus achieve higher fitness before individuals with "no theory." In this scenario, individual fitness not only depends on their own genotype but also the distribution of culturally transmitted trait values in the population. When a substantial proportion of the population possesses nonoptimal trait values, individuals with more-efficient learning strategies may enjoy a relative fitness advantage. This is especially true when population size is small and selective pressure is strong (ultimate payoff significantly contributes to genetic fitness), because "no theory" individuals may be driven to extinction before they can reach the same fitness as "correct theory" individuals.

Table 1 shows the fixation rate of the "theory" (assuming individuals possessing "theory" strategies always arrive at the right theory) and "no theory" strategies under different parameter settings in a small, well-mixed[2] population. The initial population starts with an equal number of "theory" and "no theory" individuals, who update their trait values and reproduce according to the life cycle specified in section 2.1. The fitness of individuals is modeled as the sum of a fixed baseline fitness and the ultimate payoff F. When selection is weak and noise is small, both strategies have roughly equal chance of fixation, but when selective pressure is strong and noise is substantial, the "theory" strategy is much more likely to reach fixation. Of course, here we are ignoring the problem of incorrect theories.

STUDY 2: THE CULTURAL TRANSMISSION OF CAUSAL THEORIES AND PAYOFF CURRENCIES

Given the social nature of information flow in human societies, causal theories, once constructed, may be transmitted culturally among individuals in a community. If some individual in a population comes up with a theory and starts to use intermediate payoff for trait updating in a more selective manner and achieves high final ultimate payoff F, other individuals may notice it and copy their copying strategy. For example, in a population where everyone employs a "no theory" strategy to copy arrow length and departure time based on hunting return, an individual with the "correct theory" may stand out because of their higher hunting returns. As a result, their causal theory (in addition to their many other traits) may be copied, and other individuals in the community may start to use aiming accuracy as the payoff currency for arrow length and number of prey animals encountered as the payoff currency for departure time. This process could result in the spread of theories that confer fitness advantages. Wrong theories, on the other hand, will be selected against as individuals with causal theories that do not match reality will most likely end up with lower ultimate payoffs.

To explore the evolutionary dynamics of the genetically transmitted strategy of theory construction and the culturally transmitted theories, we modified our baseline model to simulate this process. For simplicity, we assume that there are two types of genetically transmitted strategies ("theory constructor" and "non-theory constructor"), and three types of theories ("no theory," "correct theory," and "incorrect theory") that may be culturally transmitted. Thus, each agent can be represented by a list of four elements: [theory-construction strategy, theory, x_1, x_2]. In each generation, theory constructors will attempt to create a theory if they do not have intermediate payoff currencies, with probability p of arriving at the right theory and $1 - p$ the incorrect theory. Each naïve agent initially inherits the theory strategy of their parent and will then sample the

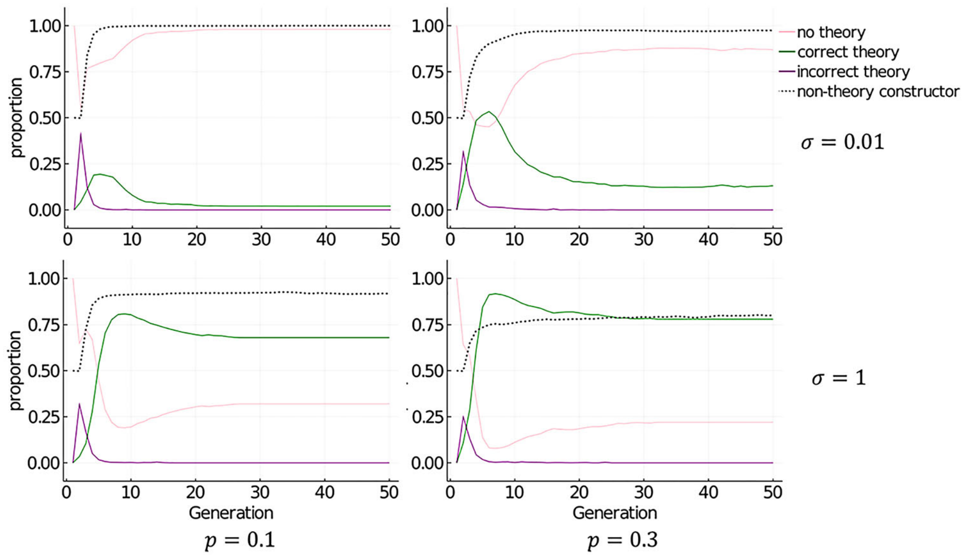

FIGURE 3 The evolutionary trajectories of various genetically and culturally transmitted strategies. Population size $N = 500$, number of models picked $n = 5$, cost of possessing the wrong theory $c = 0.5$. The initial frequencies of theory constructor and non-theory constructor in the population are the same (50%). The probability of theory constructors to arrive at the right theory is denoted by p. [This figure appears in color in the online issue]

parental generation for cultural models to learn from and compute their payoff according to their theories, as in the baseline model. The difference is that once the naïve agent has selected their model, they will not only adopt the model's trait values but also the theory that the model possesses. Importantly, they will pass on this theory regarding what currency to use for evaluating payoffs and selecting models that they acquire culturally to their offspring.

For this setup, note: first, for individuals with the "correct theory" or "incorrect theory" before trait adoption, they will evaluate their models using the intermediate payoffs by picking two models from the sample (one with the largest R, the other with the largest T). After trait updating, we assume that the focal individual will adopt the theory type of the R model if the model's R payoff is larger than that of the focal individual, and the theory type of the T model if the model's T payoff is larger than that of the focal individual. If both the R model and the T model have higher respective payoff, then the focal individual will adopt either model's theory with equal probability (i.e., 50%). In the Supplemental Material we analyze scenarios where individuals are allowed to revise their theory type and update their trait multiple times within their lifetime (one generation) and the general patterns (to be discussed below) do not qualitatively change. Secondly, for the sake of simplicity, we define the "incorrect theory" strategy as the same as "correct theory" strategy but with an additional cost c. This means that while the benefit of possessing the right theory is endogenously determined, the cost of possessing the wrong theory is fixed and externally imposed.

Figure 3 shows the evolutionary trajectories of the three culturally transmitted theory strategies and the genetically transmitted type "non-theory constructor" under different parameter combinations. We note the following observations. First, the frequency of non-theory constructor individuals always increases (see Supplemental Material Figure 2 for additional analysis of different initial non-theory constructor frequencies and p and σ values). This is expected because the correct theory is culturally transmitted, which means that the benefit of possessing it can be shared by both theory constructors and non-theory constructors, and there is always a cost of attempting to construct a theory (e.g., possibility of constructing the wrong theory). In a sense, non-theory constructors are "free riders" in that they obtain the correct theory via social learning but do not pay the cost of constructing theories (cf. Boyd, Richerson, and Henrich's [2011] discussion of imitation as a form of free riding). Interestingly, these "free riders" do not always reach fixation: when the magnitude of noise (σ) is large, we observe a stable polymorphic equilibrium of theory constructors and non-theory constructors. This is because possessing the correct theory is especially advantageous when noise is large (as shown previously) and thus may offset the cost, and theory constructors have higher probabilities of possessing the correct theory, as they will actively construct a theory when their current theory strategy is "no theory."

A second observation is that there is an obvious initial increase in the frequency of correct theories in the population, but after a few generations the frequency starts to decrease and reaches stable levels after ~20 generations. The reason is that in the beginning of the simulation, when x_1 and x_2 have not reached their optimal values in the population, the observed payoff difference between different individuals can be quite large, and those individuals with the correct theory (initially inherited from their parent) are more likely to obtain x_1 and x_2 values that are closer to their optimum and obtain higher genetic fitness. Thus, individuals with correct theories are more likely to produce offspring and also more likely to be picked as cultural models in this initial period. After x_1 and x_2 have reached their optimal values, however, the advantage conferred by possessing correct theories diminishes and because there is a cultural component to the transmission of theories (individuals will pick cultural models to update their

x_1, x_2 and theory strategies in the social learning phase), there is always some probability that the "no theory" strategy gets copied by individuals originally with correct theory due to payoff stochasticity. As such, a dynamic equilibrium of the frequency of culturally transmitted correct theory is reached, as seen in Figure 3.

These results indicate that when the ultimate payoff noise is large, "correct theory" regarding using intermediate payoff currency for trait-adoption decisions may reach substantial levels through cultural transmission, while the genetically transmitted strategy of theory-constructing may remain the minority in the population. Notably, the probability of constructing a correct theory needs not be very large; for example, when $p = 0.1$ and ff $= 1$, roughly 70% of the individuals possess the correct theory yet only 10% of the individuals are theory constructors. It is worth noting that our model assumes that the magnitude of noise and the probability of arriving at the right theory are independent. However, in more realistic scenarios, theory constructors may actively search for cultural traits that strongly correlate with the ultimate payoff. Their success in finding these traits may be negatively correlated with the magnitude of noise in the ultimate payoff, meaning that the probability of arriving at the right theory decreases as the noise magnitude increases. Future studies may further explore the evolutionary dynamics of this important possibility.

DISCUSSION

Despite the large literature on payoff-biased cultural learning, there has been little rigorous theorizing about the origins and cultural evolution of the payoff currencies that humans use to make trait-adoption decisions (for a notable exception, see primary/secondary character distinction in Boyd and Richerson, 1985). Economics tends to treat payoffs/utilities as exogenously given (Mas-Colell, Whinston, and Green, 1995), and evolutionary psychology primarily focuses on genetically evolved preferences (Cosmides, Tooby, and Barkow, 1995). Here, we present the possibility that humans may identify correlations between ultimate payoffs and some observed variable to generate "intermediate payoff" via a theory-construction process. Our simulation results show that theory construction as a strategy is most advantageous when (1) noise in genetically evolved payoffs is large, (2) selective pressures are strong, and (3) the probability of arriving at the right theory is large. The primary advantage of the "theory" strategy over the "no theory" strategy is that by identifying the channels via which cultural traits contribute to ultimate payoff F, individuals with "theory" strategy circumvent the noise associated with F and reach optimal trait values faster.

From a strict optimization perspective, the payoff currency people should use to update trait values is genetic fitness. However, humans cannot really "see" fitness, and the number of surviving offspring can be a messy measure of fitness due to paternity uncertainty (Buss, 1996; Strassmann et al., 2012), changing environments, and other noisy factors. Therefore, natural selection may select for psychological mechanisms to enable individuals to "come up with theories" in which intermediate payoffs become the cues that humans use to make trait-adoption decisions. Of course, as mentioned previously, there is no guarantee that attempts of theory construction always produce the right theory or even a good theory. At the population level, although payoff-biased transmission based on genetically evolved preferences would disfavor wrong theories, humans may construct and transmit theories for other reasons, such as content bias (Henrich and McElreath, 2003): traditional theories of illness (Murdock, 1980) and principles of magic (Hong, 2022a) and divination (Hong and Henrich, 2021) may spread in human societies partly because of their intuitive plausibility. Additionally, constructing the right theory in certain domains may be objectively difficult and require cumulative progress; thus, if the initial construction of theories is more likely to be "wrong," then it might present a significant barrier for eventually arriving at the right theory. As shown in our simulation, having the wrong theory is worse than not having any theory at all.

The substantial amount of cultural variation in what people value and pursue suggests that many existing payoff currencies may be culturally transmitted. Individuals may notice what other people pay attention to or be explicitly told what is valuable or desirable, and these culturally evolved payoff currencies may spread in the population via biased transmission. Therefore, our model accounts for not only the near-universal nondiscriminant copying but also the possibility that populations in different socio-ecological environments may culturally evolve different payoff currencies that are locally adaptive. Importantly, during cultural evolutionary time, individuals may forget why these payoff currencies are in place, in a way similar to how the original rationale for rules and norms often get lost in the process of cultural transmission (Chibnik, 1981; Hong, forthcoming).

Our simulation results suggest that "nondiscriminant copying" ("non-theory" strategy) using genetically evolved cues and "specific copying" ("theory" strategy) using culturally evolved currencies may coexist depending on the characteristics of different domains, such as the strength of selection (the extent to which "ultimate payoffs" contribute to overall fitness), the intrinsic uncertainty/noise, and the difficulty of coming up with good theories in various domains. Experimentally, it has been shown that causal understanding is not necessary for technological improvement (Derex et al., 2019), while other work suggests that experimental subjects do possess some causal knowledge in mechanical systems (Osiurak et al., 2021). More recently, fieldwork among the Hadza reveals that local bow makers often have incomplete causal understanding of their bow design and mechanical properties, meaning that they possess the correct theories in certain aspects (e.g., how bow thickness affects arrow speed) but not others (e.g., cross-section shape and energy storage). Crucially, over 80% of the bow makers reported that their causal theories about bow design and mechanics were obtained culturally while 10% reported through personal experience (Harris, Boyd, and Wood, 2021). Such patterns correspond to our simulation results nicely: while causal theories are produced by a few individuals, such theories, once produced, may be culturally transmitted among individuals and potentially enable adaptive cultural evolution. Finally, we note that individuals in large-scale, industrialized societies may have increased use of explicit theories in addition to nondiscriminant copying, possibly due to the specialization of cognitive labor and increased social complexity (Lutz and Keil, 2002; Kominsky, Zamm, and Keil, 2018).

Although the present article focuses on the evolution of "payoff currencies," the key logic of our model may be extended to the more general topic of human causal inference. In a sense, "intermediate payoff currencies" are the products of human individuals trying to understand how various factors contribute to some (often fitness-relevant) outcomes. Much research in cognitive psychology has shown that causal reasoning is a fundamental cognitive ability (Pearl, 2000) with deep evolutionary roots (Blaisdell et al., 2006). In particular, human individuals are able to reason about the causal relationships for a chain of events (e.g., A→B→C) (Shultz, Pardo, and Altmann, 1982), and children as young as four years old are shown to be able to correctly make causal chain inferences (German and Nichols, 2003). More notably, our causal models of the world often come from observing and interacting with others (Legare, Sobel, and Callanan, 2017). This means that while the capability to construct causal theories likely has a strong genetic component (though see Heyes, 2018), much of our causal learning is socially influenced. In our model, the causal theories in the form of intermediate payoff currencies are transmitted among individuals, which affects what we value and use as proxies to make learning decisions; in reality, understanding the causal structures of the world is such a key task of our cognitive life that we often actively seek causal explanations in addition to passive observations. Much of this line of work focuses on children and while there is some ongoing debate regarding the role of active teaching in traditional societies (Lancy, 2010; Terashima and Hewlett, 2016), it is a recognized fact that human individuals can and do obtain important causal knowledge from others, though such knowledge is not necessarily factually correct.[3]

In our simulation, theory construction has been modeled as a genetically transmitted strategy that can invade or be invaded by mutation. This type of idealized "one-locus two-alleles" model is quite common in population biology (Karlin and Liberman, 1975), which aims to examine the evolutionary dynamics of distinct strategies in analytically convenient settings. We wish to point out that the biological reality of causal model building and causal learning is surely much more complex than on/off control by individual genes, and further studies may utilize more gradualist assumptions in modeling human cognitive evolution (e.g., via long-term social scaffolding).[4]

We believe much more work is needed to elucidate the origin and evolution of payoff currencies that humans heavily rely upon in both individual decision-making and social learning settings, as well as the broader question of how causal theories transmit among human individuals and how they stand up to subsequent observational learning both within developmental and evolutionary time. The present study on culturally evolved and transmitted payoff currencies would not only expand the current literature on payoff-biased transmission but may also be an important first step toward a more general account of how humans construct and transmit causal theories and how these theories shape human cognitive evolution.

ORCID

Ze Hong https://orcid.org/0000-0002-5343-3008

ENDNOTES

[1] The formulation here assumes an additive contribution from the two factors to the final outcome, while in some occasions it might make sense to view it as multiplicative (e.g., the number of predators encountered and tool design). The major difference between $f(x_1) \cdot g(x_2)$ and $f(x_1) + g(x_2)$ is when x_1 and x_2 change sign or are zero. For example, if $f(x_1)$ is negative then $g(x_2)$ can be arbitrarily large yet F will always be negative. In our hunting case, we are thinking that x_1 and x_2 are always sensibly positive and nonzero, in which case if we only consider the factors that contribute to the final outcome (excluding the error term) and take the logarithm of both sides of the equation $F = f(x_1) \cdot g(x_2)$, we get $log(F) = log(f(x_1)) + log(g(x_2))$, which is again an additive formation, with the only difference being the scale of the factor currency units. Therefore, changing from the additive formulation to the multiplicative formulation will not qualitatively change the modeling results.

[2] Here, "well-mixed" means we assume that any two individuals interact with the same probability, a common assumption in theoretical ecology and evolution (Nowak, Tarnita, and Antal, 2010).

[3] Importantly, factually incorrect causal understanding of the world could still give rise to adaptive behaviors. Divination, for example, has been suggested to effectively serve as randomizing devices, which increase hunting returns as hunters decide the directions in which to seek prey (Moore, 1957).

[4] For instance, see Godfrey-Smith's (2020) proposal for a gradualist account of the evolution of human subjective experience and Planer and Sterelny's (2021) account of the evolution of language.

REFERENCES CITED

Ayton, Peter, and Ilan Fischer. 2004. "The Hot Hand Fallacy and the Gambler's Fallacy: Two Faces of Subjective Randomness?" *Memory and Cognition* 32(8): 1369–78. https://doi.org/10.3758/BF03206327.

Barrett, Brendan J., Richard L. McElreath, and Susan E. Perry. 2017. "Pay-off-Biased Social Learning Underlies the Diffusion of Novel Extractive Foraging Traditions in a Wild Primate." *Proceedings of the Royal Society B: Biological Sciences* 284(1856). https://doi.org/10.1098/rspb.2017.0358.

Blaisdell, Aaron P., Kosuke Sawa, Kenneth J. Leising, and Michael R. Waldmann. 2006. "Causal Reasoning in Rats." *Science* 311(5763): 1020–22. https://doi.org/10.1126/science.1121872.

Boyd, Robert, and Peter J. Richerson. 1985. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Boyd, Robert, and Peter J. Richerson. 1995. "Why Does Culture Increase Human Adaptability?" *Ethology and Sociobiology* 16(2): 125–43. https://doi.org/10.1016/0162-3095(94)00073-G.

Boyd, Robert, Peter J. Richerson, and Joseph Henrich. 2011. "The Cultural Niche: Why Social Learning Is Essential for Human Adaptation." *Proceedings of the National Academy of Sciences* 108(S2): 10918–25.

Buss, David M. 1996. "Paternity Uncertainty and the Complex Repertoire of Human Mating Strategies." *American Psychologist* 51(2): 161–62. https://doi.org/10.1037/0003-066x.51.2.161.

Carey, Susan. 2009. *The Origin of Concepts*. New York: Oxford University Press.

Chibnik, Michael. 1981. "The Evolution of Cultural Rules." *Journal of Anthropological Research* 37(3): 256–68.

Chudek, Maciej, Sarah Heller, Susan Birch, and Joseph Henrich. 2012. "Prestige-Biased Cultural Learning: Bystander's Differential Attention to Potential Models Influences Children's Learning." *Evolution and Human Behavior* 33(1): 46–56. https://doi.org/10.1016/j.evolhumbehav.2011.05.005.

Cosmides, Leda, John Tooby, and Jerome H. Barkow. 1995. "Introduction: Evolutionary Psychology and Conceptual Integration." In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press. https://doi.org/10.4324/9781410608994.

Day, Troy, and Russell Bonduriansky. 2011. "A Unified Approach to the Evolutionary Consequences of Genetic and Nongenetic Inheritance." *The American Naturalist* 178(2): E18–E36. https://doi.org/10.1086/660911.

Dean, Lewis G., Gill L. Vale, Kevin N. Laland, Emma Flynn, and Rachel L. Kendal. 2014. "Human Cumulative Culture: A Comparative Perspective." *Biological Reviews* 89(2): 284–301. https://doi.org/10.1111/brv.12053.

Derex, Maxime, Jean-François Bonnefon, Robert Boyd, and Alex Mesoudi. 2019. "Causal Understanding Is Not Necessary for the Improvement of Culturally Evolving Technology." *Nature Human Behaviour* 3(5): 446–52.

Durham, William H. 1991. *Coevolution: Genes, Culture, and Human Diversity*. Stanford, CA: Stanford University Press.

Enquist, Magnus, Kimmo Eriksson, and Stefano Ghirlanda. 2007. "Critical Social Learning: A Solution to Rogers's Paradox of Nonadaptive Culture." *American Anthropologist* 109(4): 727–34. https://doi.org/10.1525/aa.2007.109.4.727.

Fyfe, Sophie, Claire Williams, Oliver J. Mason, and Graham J. Pickup. 2008. "Apophenia, Theory of Mind and Schizotypy: Perceiving Meaning and Intentionality in Randomness." *Cortex* 44(10): 1316–25. https://doi.org/10.1016/j.cortex.2007.07.009.

German, Tim P., and Shaun Nichols. 2003. "Children's Counterfactual Inferences about Long and Short Causal Chains." *Developmental Science* 6(5): 514–23.

Godfrey-Smith, Peter. 2020. "Gradualism and the Evolution of Experience." *Philosophical Topics* 48(1): 201–20.

Harris, Jacob A., Robert Boyd, and Brian M. Wood. 2021. "The Role of Causal Knowledge in the Evolution of Traditional Technology." *Current Biology* 31(8): 1798–803. https://doi.org/10.1016/j.cub.2021.01.096.

Henrich, Joseph. 2016. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton, NJ: Princeton University Press.

Henrich, Joseph, and Claudia Tennie. 2017. "Cultural Evolution in Chimpanzees and Humans." In *Chimpanzees and Human Evolution*, edited by M. Muller, R. Wrangham, and D. Pilbeam, 645–702. Cambridge, MA: The Belknap Press of Harvard University Press.

Henrich, Joseph, and Francisco J. Gil-White. 2001. "The Evolution of Prestige: Freely Conferred Deference as a Mechanism for Enhancing the Benefits of Cultural Transmission." *Evolution and Human Behavior* 22(3): 165–96. https://doi.org/10.1016/S1090-5138(00)00071-4.

Henrich, Joseph, and Natalie Henrich. 2010. "The Evolution of Cultural Adaptations: Fijian Food Taboos Protect against Dangerous Marine Toxins." In *Proceedings of the Royal Society B: Biological Sciences* 277(1701): 3715–24. https://doi.org/10.1098/rspb.2010.1191.

Henrich, Joseph, and Richard McElreath. 2003. "The Evolution of Cultural Evolution." *Evolutionary Anthropology* 12(3): 123–35. https://doi.org/10.1002/evan.10110.

Henrich, Joseph, and Robert Boyd. 1998. "The Evolution of Conformist Transmission and the Emergence of Between-Group Differences." *Evolution and Human Behavior* 19(4): 215–41. https://doi.org/10.1016/S1090-5138(98)00018-X.

Henrich, Joseph, and Robert Boyd. 2002. "On Modeling Cognition and Culture Representations." *Journal of Cognition and Culture* 2(2): 87–112. https://doi.org/10.1163/156853702320281836.

Herrmann, Esther, Josep Call, María Victoria Hernández-Lloreda, Brian Hare, and Michael Tomasello. 2007. "Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis." *Science* 317(5843): 1360–66. https://doi.org/10.1126/science.1146282.

Heyes, Cecilia. 2018. *Cognitive Gadgets*. Cambridge, MA: Harvard University Press.

Hill, Kim, and Keith Kintigh. 2009. "Reports: Can Anthropologists Distinguish Good and Poor Hunters? Implications for Hunting Hypotheses, Sharing Conventions, and Cultural Transmission." *Current Anthropology* 50(3): 369–78. https://doi.org/10.1086/597981.

Hitchcock, Christopher. 1992. "Causal Explanation and Scientific Realism." *Erkenntnis* 37(2): 151–78.

Hong, Ze. Forthcoming. "The Cognitive Origin and Cultural Evolution of Taboos in Human Societies." *Journal of the Royal Anthropological Institute*.

Hong, Ze. 2022a. "A Cognitive Account of Manipulative Sympathetic Magic." *Religion, Brain & Behavior* 12(3): 254–70.

Hong, Ze. 2022b. "Combining Conformist and Payoff Bias in Cultural Evolution." *Human Nature* 33(4): 463–84.

Hong, Ze, and Joseph Henrich. 2021. "The Cultural Evolution of Epistemic Practices." *Human Nature* 32(3): 622–51.

Karlin, Samuel, and Uri Liberman. 1975. "Random Temporal Variation in Selection Intensities: One-Locus Two-Allele Model." *Journal of Mathematical Biology* 2(1): 1–17.

Kendal, Jeremy, Luc Alain Giraldeau, and Kevin Laland. 2009. "The Evolution of Social Learning Rules: Payoff-Biased and Frequency-Dependent Biased Transmission." *Journal of Theoretical Biology* 260(2): 210–19. https://doi.org/10.1016/j.jtbi.2009.05.029.

Kendal, Rachel L., Neeltje J. Boogert, Luke Rendell, Kevin N. Laland, Mike Webster, and Patricia L. Jones. 2018. "Social Learning Strategies: Bridge-Building between Fields." *Trends in Cognitive Sciences* 22(7): 651–65. https://doi.org/10.1016/j.tics.2018.04.003.

Kharratzadeh, Milad, Marcel Montrey, Alex Metz, and Thomas R. Shultz. 2017. "Specialized Hybrid Learners Resolve Rogers' Paradox about the Adaptive Value of Social Learning." *Journal of Theoretical Biology* 414:8–16. https://doi.org/10.1016/j.jtbi.2016.11.017.

Kominsky, Jonathan F., Anna P. Zamm, and Frank C. Keil. 2018. "Knowing When Help Is Needed: A Developing Sense of Causal Complexity." *Cognitive Science* 42(2): 491–523. https://doi.org/10.1111/cogs.12509.

Lancy, David F. 2010. "Learning from Nobody: The Limited Role of Teaching in Folk Models of Children's Development." *Childhood in the Past* 3(1): 79–106.

Legare, Cristine H., David M. Sobel, and Maureen Callanan. 2017. "Causal Learning Is Collaborative: Examining Explanation and Exploration in Social Contexts." *Psychonomic Bulletin & Review* 24:1548–54.

Lutz, Donna J., and Frank C. Keil. 2002. "Early Understanding of the Division of Cognitive Labor." *Child Development* 73(4): 1073–84. https://doi.org/10.1111/1467-8624.00458.

Lyons, D. E., A. G. Young, and F. C. Keil. 2007. "The Hidden Structure of Overimitation." *Proceedings of the National Academy of Sciences* 104(50): 19751–56. https://doi.org/10.1073/pnas.0704452104.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Moore, Omar Khayyam. 1957. "Divination—A New Perspective." *American Anthropologist* 59:69–74. https://doi.org/10.1525/aa.1957.59.1.02a00060.

Murdock, George Peter. 1980. *Theories of Illness a World Survey*. Pittsburgh, PA: University of Pittsburgh Press.

Nakahashi, Wataru, Joe Yuichiro Wakano, and Joseph Henrich. 2012. "Adaptive Social Learning Strategies in Temporally and Spatially Varying Environments: How Temporal vs. Spatial Variation, Number of Cultural Traits, and Costs of Learning Influence the Evolution of Conformist-Biased Transmission, Payoff-Biased Transmission, and Individual Learning." *Human Nature* 23:386–418. https://doi.org/10.1007/s12110-012-9151-y.

Nowak, Martin A., Corina E. Tarnita, and Tibor Antal. 2010. "Evolutionary Dynamics in Structured Populations." *Philosophical Transactions of the Royal Society B: Biological Sciences* 365(1537): 19–30.

Osiurak, François, Salomé Lasserre, Julie Arbanti, Joël Brogniart, Alexandre Bluet, Jordan Navarro, and Emanuelle Reynaud. 2021. "Technical Reasoning Is Important for Cumulative Technological Culture." *Nature Human Behaviour* 5(12): 1643–51. https://doi.org/10.1038/s41562-021-01159-9.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Placek, Caitlyn D., Purnima Madhivanan, and Edward H. Hagen. 2017. "Innate Food Aversions and Culturally Transmitted Food Taboos in Pregnant Women in Rural Southwest India: Separate Systems to Protect the Fetus?" *Evolution and Human Behavior* 38(6): 714–28. https://doi.org/10.1016/j.evolhumbehav.2017.08.001.

Planer, Ronald, and Kim Sterelny. 2021. *From Signal to Symbol: The Evolution of Language*. Cambridge, MA: MIT Press.

Richerson, Peter J., and Robert Boyd. 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago: University of Chicago Press.

Shultz, Thomas R., Seymour Pardo, and Esther Altmann. 1982. "Young Children's Use of Transitive Inference in Causal Chains." *British Journal of Psychology* 73(2): 235–41.

Stibbard-Hawkes, Duncan N. E., Robert D. Attenborough, and Frank W. Marlowe. 2018. "A Noisy Signal: To What Extent Are Hadza Hunting Reputations Predictive of Actual Hunting Skills?" *Evolution and Human Behavior* 39(6): 639–51. https://doi.org/10.1016/j.evolhumbehav.2018.06.005.

Strassmann, Beverly I., Nikhil T. Kurapati, Brendan F. Hug, Erin E. Burke, Brenda W. Gillespie, Tatiana M. Karafet, and Michael F. Hammer. 2012. "Religion as a Means to Assure Paternity." *Proceedings of the National Academy of Sciences of the United States of America* 109(25): 9781–85. https://doi.org/10.1073/pnas.1110442109.

Tennie, Claudio, Josep Call, and Michael Tomasello. 2009. "Ratcheting Up the Ratchet: On the Evolution of Cumulative Culture." *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1528): 2405–15. https://doi.org/10.1098/rstb.2009.0052.

Terashima, Hideaki, and Barry S. Hewlett. 2016. *Social Learning and Innovation in Contemporary Hunter-Gatherers*. Berlin: Springer.

Thurstone, Louis L. 2017. "Psychophysical Analysis." In *Scaling: A Sourcebook for Behavioral Scientists*. New York: Routledge. https://doi.org/10.4324/9781315128948-6.

Vale, Gillian L., Emma G. Flynn, Jeremy Kendal, Bruce Rawlings, Lydia M. Hopper, Steven J. Schapiro, Susan P. Lambeth, and Rachel L. Kendal. 2017. "Testing Differential Use of Payoff-Biased Social Learning Strategies in Children and Chimpanzees." *Proceedings of the Royal Society B: Biological Sciences* 284(1868): 20171751. https://doi.org/10.1098/rspb.2017.1751.

Whitehead, Hal, and Peter J. Richerson. 2009. "The Evolution of Conformist Social Learning Can Cause Population Collapse in Realistically Variable Environments." *Evolution and Human Behavior* 30(4): 261–73. https://doi.org/10.1016/j.evolhumbehav.2009.02.003.

Whiten, Andrew, Nicola McGuigan, Sarah Marshall-Pescini, and Lydia M. Hopper. 2009. "Emulation, Imitation, Over-Imitation and the Scope of Culture for Child and Chimpanzee." *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1528): 2417–28. https://doi.org/10.1098/rstb.2009.0069.

Wood, Lara A., Rachel L. Kendal, and Emma G. Flynn. 2013. "Whom Do Children Copy? Model-Based Biases in Social Learning." *Developmental Review* 33(4): 341–56. https://doi.org/10.1016/j.dr.2013.08.002.

How to cite this article: Hong, Ze, and Joseph Henrich. 2023. "Evolving payoff currencies through the construction of causal theories." *American Anthropologist* 1–12. https://doi.org/10.1111/aman.13926