



Original Article

Strategic social learning and the population dynamics of human behavior: the game of Go

Bret Alexander Beheim^{a,*}, Calvin Thigpen^b, Richard McElreath^c^a Department of Anthropology, University of New Mexico, Albuquerque, NM 87131, USA^b Institute for Transportation Studies, University of California, Davis, CA 95616, USA^c Department of Anthropology, University of California, Davis, CA 95616, USA

ARTICLE INFO

Article history:

Initial receipt 12 September 2013

Final revision received 4 April 2014

Keywords:

Cultural evolution

Game of Go

Social learning

Game theory

ABSTRACT

Human culture is widely believed to undergo evolution, via mechanisms rooted in the nature of human cognition. A number of theories predict the kinds of human learning strategies, as well as the population dynamics that result from their action. There is little work, however, that quantitatively examines the evidence for these strategies and resulting cultural evolution within human populations. One of the obstacles is the lack of individual-level data with which to link transmission events to larger cultural dynamics. Here, we address this problem with a rich quantitative database from the East Asian board game known as Go. We draw from a large archive of Go games spanning the last six decades of professional play, and find evidence that the evolutionary dynamics of particular cultural variants are driven by a mix of individual and social learning processes. Particular players vary dramatically in their sensitivity to population knowledge, which also varies by age and nationality. The dynamic patterns of opening Go moves are consistent with an ancient, ongoing arms race within the game itself.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

One of the greatest contributions of evolutionary theory is the ability to answer “why” questions about living things, explaining the existence of their intricate and beautiful designs as adaptations via natural selection. This success extends to study of our own species; we now know a great deal about our behavior, psychology and, more indirectly, our ancestry, as a consequence of evolutionary theory (Barkow, Cosmides, & Tooby, 1992; Laland & Brown, 2002). Many pathological or idiosyncratic human social phenomena, for example, can be understood as a mismatch between the environment our minds and bodies are “expecting” (subsistence foraging and horticulture) and the one many of us live in (industrial food markets and sedentary lifestyles).

It is less clear, though, how much evolutionary theory helps us understand the dynamics of human history, replete with its rapid changes in technology, institutions and ideas. Work in the last three decades has approached this problem by treating culture as a system of inheritance, akin to but distinct from genetic transmission (Durham, 1992; Richerson & Boyd, 2005). Mathematical models have shown that such systems will exhibit Darwinian adaptive dynamics even if the traits under study are not discrete, particulate

replicators like genes, and even if they are not created through random mutation (Boyd & Richerson, 1985; Henrich & Boyd, 2002).

As with genetic evolution, the driving processes are organized within a set of general categories we call “forces,” but unlike natural selection, genetic drift, meiotic drive, etc., the forces of cultural evolution are largely based on the details of human psychology. What humans find easy or appealing to learn, what cues and associations they find salient, and how readily certain cultural traits encourage their own social transmission. As a result, much research in cultural evolution has concentrated on human social and individual learning (Mesoudi, Whiten, & Laland, 2006), and has shown that simple learning heuristics, such as “when in doubt, copy what the majority is doing” are plausible explanations for how humans efficiently identify fitness-enhancing behaviors (McElreath et al., 2008).

These learning processes, in aggregate, change the population distribution of cultural variants in a recursive fashion, and in turn alter the environment for future social learning. Thus, by applying what Ernst Mayr called “population thinking” to human culture, it is possible to develop mechanistic theories of how socially transmitted technologies, beliefs and behaviors spread and decline over time (Shennan, 2009).

Though these learning models have been supported by laboratory and field experiments in which participants are given controlled access to potentially valuable social information (Baum, Richerson, Efferson, & Paciotti, 2004; Caldwell & Millen, 2008; Efferson, Takezawa, & McElreath, 2007; Mesoudi, 2011), studying cultural evolution in outside controlled settings is a much harder task. Michel

* Corresponding author. Department of Anthropology, University of New Mexico, Albuquerque NM 87131, USA.

E-mail address: babeheim@unm.edu (B.A. Beheim).

et al. (2011) use Google book archives to trace word use frequencies over the last several centuries but, lacking information about *who* is using those words and under what circumstances, they are unable to explain these trends beyond phenomenological descriptions and qualitative associations to major historical events. Ideally, if we had individual-level phenotypic measures for an entire population over many years, we could identify and measure the driving forces of cultural change, as has been accomplished in non-human evolutionary systems (e.g. Grant & Grant, 2002; Ozgul et al., 2009).

In what follows, we present an analysis of such a high-resolution ($n = 31,137$) archive of cultural patterns: records from the East Asian game of Go. Over many decades of professional play, we can see particular strategies and counterstrategies be innovated, copied, decline, and vanish altogether. Our results provide strong evidence that, on the average, these trends are the result of strategic social learning decisions by professional Go players as they adopt and abandon variant openings. This learning model, in aggregate, is the basis for an apparent evolutionary arms race that is largely responsible for the oscillations in Go openings among professional players. In this way, we are able to link individual-level strategic learning to aggregate population dynamics over much longer timescales.

2. Historical dynamics in the game of Go

The East Asian board game known in the West as Go is, by players, one of the most popular games in the world, and certainly one of the oldest living games. Originating in China between 2000 and

3000 years ago (where it is called *weiqi*), it has since diffused to Korea (there the game is called *baduk*) and Japan (*go* or *igo*). Through the Japanese, it became widely known in the West in the early 20th century, and most American and European Go terminology is derived from Japanese. With many centuries of paid professional leagues and an abundance of game records (see Fig. 1), Go is an ideal “model system” for the quantitative study of human behavioral change (Gobet, de Voogt, & Retschitzki, 2004).

The rules of the game are easily explained: two players (“Black” and “White”) sit on opposite sides of a large 19×19 coordinate grid. Starting with Black, each player alternatively places a single stone on an empty grid intersection of their choice. At the end of the game, the winner is the player who controls the majority of the board (that is, whose stones surround or occupy the most grid intersections). If the eventual outcome is obvious, one player may resign mid-game; otherwise, once all possible moves have been played, the players count grid intersections to see who is the winner.

Beyond this basic outline of the game, there are two key rules regulating play. First, stones or groups of stones of one player that become completely surrounded on the grid by stones of the other are removed from the board (this is the only time stones are moved, once placed). Second, to prevent the possibility of the game entering an infinite loop, the *ko* rule restricts play so that the same board state cannot appear twice in one game.

Though the rules above are remarkably simple, the game is fantastically complex, and it takes years of dedicated study to become a competitive player. Much of one's skill in Go comes from the ability

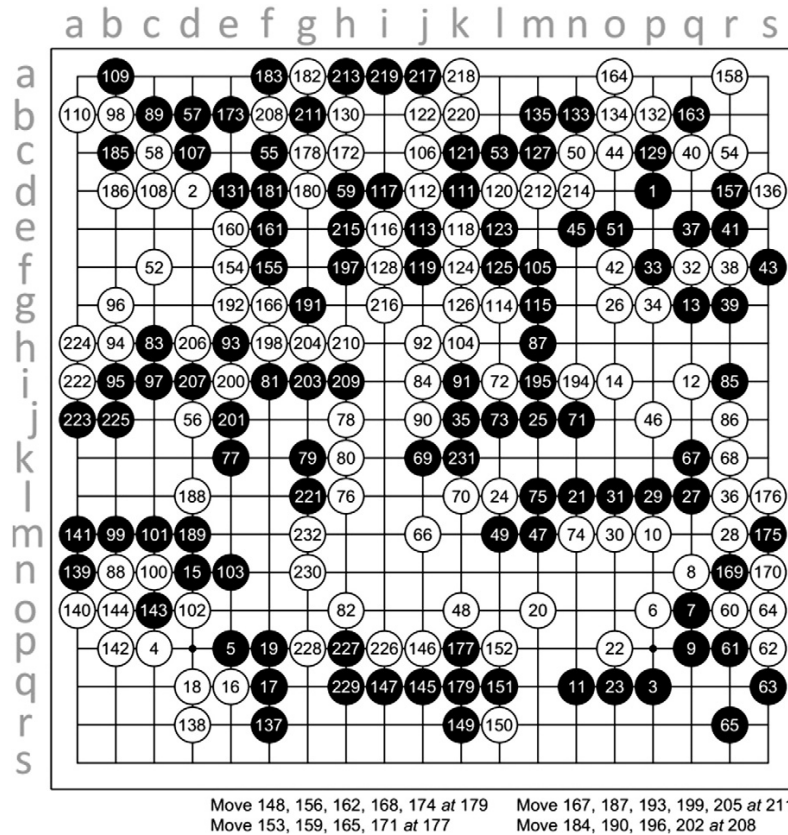


Fig. 1. A complete record, with SGF coordinates, for a 1978 game between Japanese professionals Takemiya Masaki and Kudo Norio. Black (Takemiya) begins with a Fourfour at coordinate *pd* (top right), with White (Kudo) responding at *dd* with another Fourfour. White won by resignation after 232 moves. This game is one of the first successful uses of the *dd* response to the Move 1 Fourfour, and *dd* has subsequently become the dominant response among professionals (see Fig. 3).

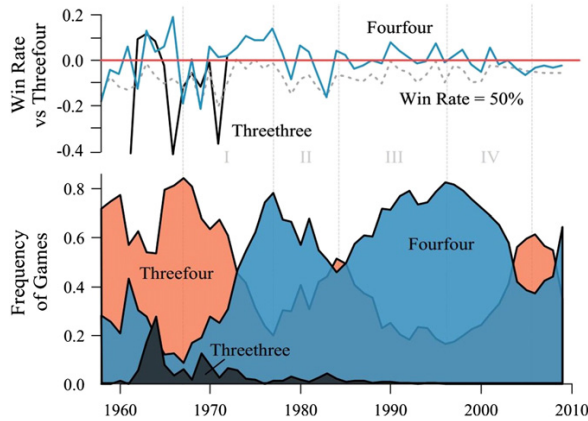


Fig. 2. The ebb and flow of the three major first move variants, with proportion of games won. The Fourfour's rise and fall, here broken into four eras, closely tracks with its relative performance (top) since 1967 ($r = 0.45$). Similarly, the Threethree's brief efflorescence and extinction closely match its relative win rate. Because of its initial dominance, all win rates are expressed relative to the Threefour's, and the dotted gray line indicates the location of the 50% (break-even) point.

to make judgments about the long-term influence of particular moves and formations on the board, a task that currently bedevils attempts at developing strong computer opponents (Rimmel et al., 2010).

The first few moves of game of Go are actually the most predictable. The opening period of the game is often characterized by highly normative, canalized play, where players often follow one of a set of memorized move sequences called *joseki* (or *jungsuk* in Korea), thought to grant neither player a major advantage. Once a particular *joseki* pattern has started, professional players will draw on an encyclopedic knowledge of explored routes to determine their next move.

Despite the conservatism of opening play in professional games, novel variations appear regularly and, if widely copied, become new *joseki* in the canon. Go historian Peter Shotwell writes of the evolution of *joseki*: “Some are hundreds (or thousands) of years old; others are born, die, and then are reborn again as the style of play changes; still others are yet to be born and patiently wait their turn to be discovered in some flash of ingenuity or desperation” (Shotwell, 2003).

As shown in Fig. 2, we can see such dynamics of the very first move (Move 1, hereafter “M1”) over the latter half of 20th century professional play. Once board symmetry is taken into account, there are only two well-known positions Black will play, the Fourfour and the Threefour, each named after their coordinate locations on the corners of the board.

After a brief enthusiasm for the Threethree in the 1960s, the Fourfour rose to dominance, peaking in 1977 (appearing in 78% of games played) and again in 1996 (82% of games played). The primary goal of this analysis is to understand why this pattern exists, and what processes are driving the evolution at M1.

We can trace the rapid diffusion of the Fourfour during this time back to the success of a single innovator, Japanese professional Takemiya Masaki. Beginning his pro career at the age of 14 in the 1960s, Takemiya has used the Fourfour almost exclusively; of his 520 games as Black in the 2009 game database, only 6 do not start with the Fourfour. Takemiya's novel use of this variant in 1968 marks the beginning of its spread within the broader Go community, and his personal performance for the next 5 years (64%) also exceeded the average Black win rate (61.5%).

Subsequent to Takemiya's arrival, the success of the Fourfour appears to have prompted an innovation race in White's Move 2 response. Before 1970, White's response to the Fourfour was almost always a Threefour on an adjacent corner. If we imagine the M1

Fourfour is played in the upper right quadrant (SGF coordinate *pd*), nearly 80% of the time the M2 response in the late 1960s was coordinate *dc* (Fig. 3). By 1977, at the Fourfour's first peak, the *dc* response declined to about 30% of games. In the interim, M2 variants at *dp*, *cq* and *dd* all became established in rapid sequence. As measured by the proportion of games White won, *dp* and *cq* both consistently outperformed *dc* in the 1970s, blunting the Fourfour's effectiveness and anticipating its decline in use after 1977.

2.1. Strategic social learning

From the above, it appears that the waxing and waning success of the Fourfour stems from an ongoing crucible of innovation and counter-innovation. As it rapidly proliferated in the professional community in the 1970s, the Fourfour enjoyed an anomalously large win rate. Conversely, the Fourfour's decline after 1977 occurred when its win rate was consistently below that of its rival variant, the M1 Threefour.

With so many possible opening moves, the few that are observed over the decades of play suggest an important role for cultural transmission from player to player through the generations. But is social learning responsible for the oscillations in move frequencies observed in Figs. 2 and 3?

As strategic actors in a high-stakes competitive environment, professional Go players are under immense pressure to choose effective moves that help defeat their opponents. We should expect that players update their valuation of particular sequences through trial and error in their own games, so that past performance predicts current behavior.

At the same time, the complexity of the game, dynamic strategy space, and high cost of individual experimentation make objective conclusions about the current value of particular moves very difficult to establish. As such, it is also plausible that players may strategically tap into the knowledge and experience contained in the games of other players, records of which are widely available (with many of the games taking place at the same Go institutes in China, Japan and South Korea).

Theoretical models and experiments have focused on two important forms of such strategic social learning: frequency-dependent and success-biased transmission (Boyd & Richerson, 1985; McElreath et al., 2008). Frequency dependence is one of the simplest kinds of social learning, in which learners use the popularity of different behaviors to decide which to adopt. It can arise through an unbiased or biased sampling scheme, and by copying the behavior of

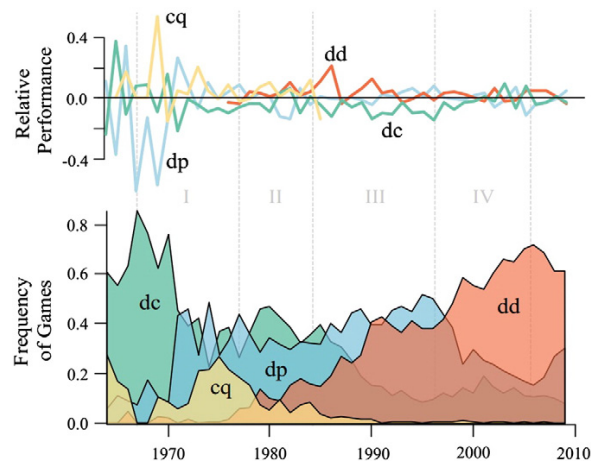


Fig. 3. Frequency and performance for four major M2 responses to the Fourfour. Performance for each is calculated as the win rate subtracting the average win rate of alternatives. As at M1, a variant's frequency closely tracks its relative performance. To compare with Fig. 2, we define four eras (I to IV) corresponding to the Fourfour's rise and fall.

Table 1
M1 Fourfour use logistic regression coefficients (mean and standard error, $n = 31,137$).

	Mean (SE)
Fixed effects	
Personal Fourfour use rate (β)	2.24 (0.18)
× Personal Fourfour win rate	1.49 (0.52)
× Personal win rate	0.18 (0.41)
Personal Fourfour win rate	1.02 (0.16)
Population Fourfour use rate (γ)	2.35 (0.40)
× Population Fourfour win rate	10.29 (2.20)
× Personal win rate	−0.10 (0.71)
Population Fourfour win rate	2.17 (0.48)
Handicap (<i>komi</i>)	0.04 (0.01)
Intercept	0.13 (0.06)
Varying effects	
Player _{<i>j</i>}	0.19 (0.04)
× Personal Fourfour use rate	4.00 (0.68)
× Population Fourfour use rate	9.77 (1.86)
Age _{<i>k</i>} × personal Fourfour use rate	0.08 (0.03)
Age _{<i>k</i>} × population Fourfour use rate	3.22 (1.01)

others, individuals can quickly and cheaply acquire useful information that would otherwise require costly individual learning.

Success-biased social learning exploits the correlation between behaviors and outcomes, so that a move that is associated with greater success among others may be favored simply for that reason, regardless of who is using it, how popular it is, or whether the learner really understands why it works.

Using statistical modeling techniques, we can evaluate how consistent real players are with these idealized models. Each of these three concepts, individual updating, frequency-dependent and success-biased social learning, were formalized as multilevel logistic regression models for the probability any particular game i will start with a Fourfour.

3. Methods

We model the decision to use the Fourfour variant at M1 via a binomial model with a logistic link function of the probability of use, fitted by Hamiltonian MCMC in the STAN software package (Stan Development Team, 2013) and the R programming language. For each game i , we model whether Move 1 was a Fourfour ($y_i = 1$) or not ($y_i = 0$) via a multilevel logistic regression framework. We used two classes of predictors, “Personal Knowledge” and “Population Knowledge.” Personal variables include a player’s recent (last 2 years) use of the Fourfour for Move 1, and how often they won games with it, relative to how often they won games without it. This model corresponds to the hypothesis that players update their views of the move, based on how familiar it is to them and how many games they have won using it. “Population Knowledge” models use various predictors constructed from the recent population use of the Fourfour, and the population win rate of the Fourfour relative to its alternatives. These attempt to capture frequency-dependent or success-biased social learning.

Predictor variables for each game used a retrospective 2-year window from the date the game took place, taking the simple average of values for relevant games within that period (with shorter time horizons producing similar results, see Table S2, available on the journal’s Web site at www.ehbonline.org). Each predictor representing a use frequency was centered on 0.5, a natural point of interpretation for whether a behavior is in the majority. Predictors for win rate, in contrast, were centered on the average win rate of alternative variants at that game move.

4. Results

With all predictors centered on means, the intercept estimate in Table 1 shows that an average Go player used the Fourfour about 53.2% of the time during the study period (i.e. $\text{logit}(.532) = 0.13$). A player’s recent use of a particular opening move strongly predicts their future use; for an otherwise-average player who used the Fourfour about 60% of their recent games, the probability they will use the Fourfour in the current game rises to 58.7% (OR: 9.39). There is a smaller (OR: 2.77), but unambiguous, positive effect from past success with the Fourfour, adding about two percentage points to the probability of use for a player who has won 10% more games using the Fourfour than his overall win rate. These two personal effects also interact, so that past success enhances the predictive power of past use (or vice versa).

Most important to the question of social learning and cultural evolution, though, is whether there is a predictive value in the recent behavior of *other* players, controlling for individual behavior. Indeed, when the Fourfour appears in 60% of all recent games, the expected probability of use by an average player rises from 53.2% to 59.0% (OR: 10.475). Two important interaction effects regulate the frequency dependence of a player’s behavior; if the Fourfour appears in 60% of recent professional games, *and* its population performance is 10 percentage points higher than that of its rival, the Threefour, the expected probability a player will use it rises from 59.0% to 61.4% (Fig. 4). A player’s personal performance, on the other hand, has a dampening effect on their sensitivity to population information; as a player’s win rate relative to the Black average increases, they are less sensitive to the popularity of the Fourfour among others.

Because the coefficient for personal use (which we call β) and that for population use (γ) have the same units, comparing their magnitudes tells us the relative predictive value of personal and population knowledge. For an average player, social knowledge is 1.05 times more important than personal experience in predicting the outcome, but this belies the enormous variation from player to player. Partially pooling games of each Black player, we see a varying effects standard deviation of 4.00 for the distribution of player-specific β . For population knowledge, the standard deviation across player

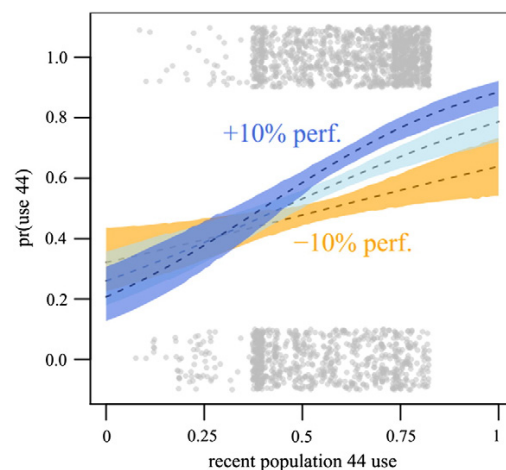


Fig. 4. Predicted probability of using the Fourfour by an average player at different values of population use, with means (black dashed lines) and 95% HPDI (solid regions). Per Table 1, the relative win rate of the Fourfour within the population regulates the predictive importance of population use, −10% versus the Threefour (orange), equal in rate to the Threefour (light blue) and +10% (dark blue). Outcome data are also shown via scatterplot.

intercepts is even larger, 9.78. Plotting the joint posterior densities of β_j and γ_j for each player j (Fig. 5, left) reveals that some professionals are substantially better predicted by social knowledge than by their own recent experience (e.g. Lee Sedol), while other players are almost completely insensitive to population trends (foremost among them, Takemiya Masaki).

Including player age as a varying effect regulating magnitude of β and γ also shows learning heterogeneity over the life course (Fig. 5, right). The games of the youngest professionals, those between the ages of 14 and 18, are least predicted by personal experience and most by population trends. After age 20, the balance steadily shifts toward personal experience, peaking at age 27 well above the average β of 2.24. Players older than 35 appear to return to average levels of social learning; with a few deviations, no trend is apparent after this age.

5. Discussion

Though only a preliminary treatment of a very rich data set, the above results demonstrate how a socially transmitted cultural pattern of behavior—the first move variants in the game of Go—undergo evolutionary change. In deciding which opening move to use, it appears Go professionals attend to both their own past experiences and those of other players, placing substantial weight on the recent popularity and performance of different variants.

As each player's decision and subsequent game outcome adds new information to the body of population knowledge, this process proceeds recursively, driving the observed cyclic trends in move frequencies. We can also see how the breakaway success of the Fourfour variant at M1 was followed by a period of innovation and diversity in more effective M2 responses, consistent with an ongoing evolutionary arms race conducted through strategic social learning.

These results, calculated using a 2-year retrospective window from the date of each game, are compared with identical analyses from windows set at 6, 12 and 18 months retrospective in Table S2, available on the journal's Web site at www.ehboneline.org. One notable difference across these analyses is the interaction between recent population use of the Fourfour and its population win rate, which strengthens as more of the recent past is taken into account. This time-increasing interaction could represent the greater evidence attested by 24 months of trends rather than only 6, or possibly the time needed for such information to fully diffuse across the professional community.

According to player-specific posterior estimates, the players who most strongly track population trends are South Korean professional Lee Sedol and Cho Namchul, while Chinese player Peng Quan and Japanese player Takemiya Masaki lie at the other extreme. This

segregation of nationalities may not be a coincidence; as Table 2 shows, the average of the player-specific estimates for individual and population knowledge differs across country, with South Korean players tending toward greater-than-average sensitivity to social information (positive mean γ_j). The average Japanese player, in contrast, is better predicted by their individual knowledge (positive mean β_j), with the 76 Chinese players nearest to the population intercepts β and γ , on the average. It is important to emphasize that these distributions show substantial overlap, but this result raises the tantalizing possibility that normative or institutional differences across East Asia may encourage different forms of social learning within each national league.

Although we can do much to reconstruct and explore the historical record, several important questions remain outstanding. First, we do not know the nature of the observed success bias. We can see that relative performance of the Fourfour versus the Threefour consistently improves our predictions of which one a particular player will use. This seems to indicate players focus on how well a move does, but it is equally possible, however, that successful *players* rather than successful variants are what are emulated.

This is not captured by crude measures like population use frequency, and is indicative of a larger problem: establishing exactly who players are learning from. The statistical models above treat the general population as an equally connected, homogenous pool. Real social influence in this situation undoubtedly involves a variety of complex social relationships between unique personalities. The nature of these influences may be better captured by network-based approaches that allow for heterogeneous interconnections between players. It may be possible to use information about players' nationalities or recent matches to develop measures of social proximity, allowing the detection of theoretical effects like conformist learning or prestige.

Another problem is the nature of innovation in Fourfour play. Even if Takemiya's innovative opening style (widely known as "Cosmic Go") was the catalyst for the Fourfour's rise in the 1970s, as the evidence above seems to indicate, we do not know if other adopters were using it in the same way. It could be, for example, that Takemiya's early successes in 1968–1970 encouraged other professionals to revisit the Fourfour and develop their own innovative applications, rather than emulate his.

Finally, even if players really are using a combination of individual and population knowledge, we should consider the possibility that this offers no real strategic advantage. A well-understood problem with success bias is that it aims to exploit an association between use and performance, but can lead to the proliferation of neutral or even deleterious practices if that association is not causal. Rather than

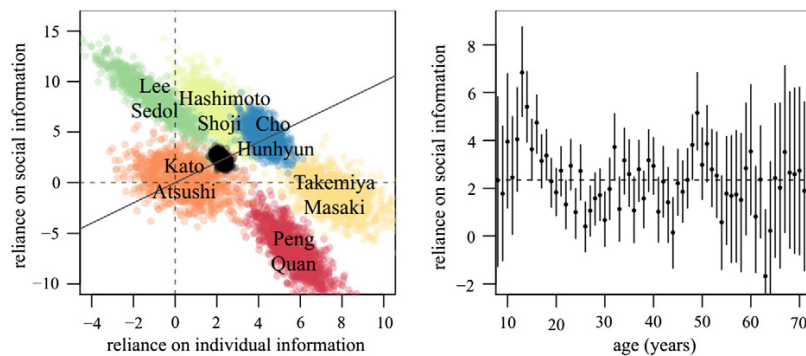


Fig. 5. (Left) Six player-specific joint posterior densities of the importance of individual (β_j) and social (γ_j) information on predicting move outcome, along with the main effect density (black region). Players who lie above the line of equality (gray) are more predicted by recent population trends than individual experience, while the reverse is true for players below it. (Right) Estimates for the importance of social knowledge for each age k , showing means and 95% HPDI intervals. The youngest professional players tend to have values of γ_k above the main effect mean γ (dotted), but this declines during their late teens and 20s.

Table 2

National averages of player-specific estimates of reliance on individual (β_j) and social (γ_j) information.

Nationality	<i>n</i>	β_j (SE)	γ_j (SE)
Chinese	75	0.05 (0.19)	0.05 (0.27)
Japanese	64	0.65 (0.20)	−0.76 (0.27)
South Korean	58	−0.69 (0.17)	0.72 (0.31)
Taiwanese	7	−0.85 (0.35)	0.55 (0.62)

produce adaptation, social learning may simply be diffusing the (arbitrary) behaviors of a high-performing elite.

There are many more possible avenues of exploration in this very rich data set. Venturing beyond Move 1, we can study the diffusion of particular *joseki* patterns, the importance of particular high-status individuals or, indirectly, the diffusion of general strategies within professional play. The scope of study can also be expanded; large archives of professional games go back at least several hundred years, and tens of thousands of amateur games are available online, to say nothing of similar archives available for other games like chess.

In any case, the statistical and theoretical tools employed above can be used to expand our understanding of cultural dynamics. The above results hopefully demonstrate the value of adopting a quantitative, Darwinian point of view toward cultural change. As in evolutionary ecology, knowledge gained from such high-resolution data sets can carry over to human historical dynamics for which few data remain and variants are harder to define.

Indeed, we feel that the theoretical and empirical evidence is so overwhelming that the key question is no longer *whether* culture evolves, but rather *which* evolutionary processes operating in day-to-day life brought about the particular patterns of cultural diversity we see in human history.

Appendix 1. The GoGoD database

Games were pulled from the 2009 Games of Go on Disk (GoGoD) library, published by T Mark Hall and John Fairbairn. Data are drawn from all database games within the years 1954 to 2009, but because predictors use a 2-year retrospective window from the day and month the game was played, the first games analyzed took place in 1956. From this year range, we excluded all games that took place with handicap stones (in which White plays first, with Black stones already on the board). We also excluded exhibition games with unusual rules, games with known corruptions described in the comments, and games in which one of the players was amateur player by rank, rather than a paid professional (the one exception to this rule was Kikuchi Yasuro, b. 1929, who played evenly with top professionals).

Since we are using player-specific predictors and grouping games by player in our varying effects, we excluded games whose Black players did not have at least 50 games on record between 1956 and 2009. After these exclusions, a total of 31,137 games remained within the database, with 207 unique players taking Black. In addition to the data contained within the game record and the metadata, we also collected player nationality and date of birth was drawn from a variety of sources, chiefly player biography pages from the Go wiki “Sensei’s Library” (<http://senseis.xmp.net/>). Although the GoGoD database is itself commercially licensed by Fairbairn and Hall, the regression table used here is available by request to the corresponding author.

Appendix 2. Full regression model

Our basic question is whether or not a player will use the Fourfour opening move in any particular game, modeled as a binomial random variable with a logistic probability model. This regression model has three important features: terms relating to personal information, terms relating to population information, and stochastically modeled

player-specific intercepts and slopes. Specifically, if the probability of using the Fourfour in game *i* of Black player *j* who is aged *k* is p_{ijk} , then the log-odds that the Fourfour opening move is used is modeled by

$$\log\left(\frac{p_{ijk}}{1-p_{ijk}}\right) = \alpha + \alpha_j + \beta_{ijk}x_i + \gamma_{ijk}y_i + \kappa q_i + \eta r_i + \text{controls}$$

Predictor x_i is the proportion of Black player *j*’s games that use the M1 Fourfour recently (within the last 2 years), y_i is the frequency of the M1 Fourfour among recent games within the general population, q_i is that player’s recent win rate specifically among their Fourfour games, and r_i is the win rate among all recent Fourfour games in the population. Of particular interest are parameters β_{ijk} and γ_{ijk} , which represent the importance of personal and population use frequencies in the model. These two parameters are themselves modeled by

$$\beta_{ijk} = \beta + \beta_j + \beta_k + \beta_S s_i + \beta_Q q_i$$

$$\gamma_{ijk} = \gamma + \gamma_j + \gamma_k + \gamma_S s_i + \gamma_R r_i$$

In addition to interactions for q_i and r_i , which allow the Fourfour’s recent win rate for the player and the population to regulate the importance of its frequency, predictor s_i is the Black player’s overall win rate among games played in the last 2 years. Within the model, intercept α_j and slopes β_j and γ_j for player *j* treated as varying effects following a joint normal distribution with mean 0, as are those of players of age *k*.

The controls for each game *i* are made up of six predictors: White’s win rate using the M1 Fourfour in past games (where they played as Black), White’s loss rate against the M1 Fourfour, White’s frequency of use of the Fourfour, the number of past games Black has seen the Fourfour used, whether or not Black lost the last time someone used an M1 Fourfour on him, and the handicap of game *i*.

This model was first fit using the lme4 package in the R language and by Hamiltonian Monte Carlo via the STAN software package called by the glmer2stan command (McElreath, 2013), running 5000 iterations with a 1000-iteration burn-in. Convergence of chains was determined by the R-hat Gelman & Rubin diagnostic (Stan Development Team, 2013) and by visual inspection of the chains. Estimates for each parameter of the full model are presented in Table S1, available on the journal’s Web site at www.ehbonline.org of the ESM and compared against those of two other models, which exclude personal and population knowledge predictors, respectively.

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.evolhumbehav.2014.04.001>.

References

- Barkow, J., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford University Press.
- Baum, W., Richerson, P. J., Efferson, C., & Paciotti, B. (2004). Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evolution and Human Behavior*, 25(5), 305–326.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Caldwell, C. A., & Millen, A. E. (2008). Studying cumulative cultural evolution in the laboratory. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 363(1509), 3529–3539.
- Durham, W. H. (1992). *Coevolution: Genes, culture, and human diversity*. Stanford University Press.
- Efferson, C., Takezawa, M., & McElreath, R. (2007). New methods in quantitative ethnography: Economic experiments and variation in the price of equality. *Current Anthropology*, 48, 912–919.
- Gobet, F., de Voogt, A., & Retschitzki, J. (2004). *Moves in mind: The psychology of board games*. Psychology Press.
- Grant, P. R., & Grant, B. R. (2002). Unpredictable evolution in a 30-year study of Darwin’s finches. *Science*, 296(5568), 707–711.
- Henrich, J., & Boyd, R. (2002). On modeling cognition and culture: Why cultural evolution does not require replication of representations. *Journal of Cognition and Culture*, 2(2), 87–112.
- Laland, K. N., & Brown, G. R. (2002). *Sense and nonsense: Evolutionary perspectives on human behaviour*. Oxford University Press.
- McElreath, R. (2013). *glmer2stan: Rstan models defined by glmer formulas, R package version 0.99*.

- McElreath, R., et al. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and payoff-biased social learning strategies. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 363, 3515–3528.
- Mesoudi, A. (2011). An experimental comparison of human social learning strategies: Payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, 32(5), 334–342.
- Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29, 329–383.
- Michel, J. B., et al. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014), 176–182.
- Ozgul, A., et al. (2009). The dynamics of phenotypic change and the shrinking sheep of St. Kilda. *Science*, 325(5939), 464–467.
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. University of Chicago Press.
- Rimmel, A., et al. (2010). Current frontiers in computer go. *Computational Intelligence and AI in Games. IEEE Transactions*, 2(4), 229–238.
- Shennan, S. (Ed.). (2009). *Pattern and process in cultural evolution*. University of California Press.
- Shotwell, P. (2003). *Go! More than a game*. Tuttle Publishing.
- Stan Development Team (2013). Stan: A C++ Library for Probability and Sampling, Version 20. <http://mc-stan.org/>.