COOPERATION, CONFORMITY, AND THE COEVOLUTIONARY PROBLEM OF TRAIT ASSOCIATIONS

JEREMY VAN CLEVE^A

^ADEPARTMENT OF BIOLOGY, UNIVERSITY OF KENTUCKY, LEXINGTON, KY 40506 USA

Abstract

In large scale social systems, coordinated or cooperative outcomes become difficult because encounters between kin or repeated encounters between friends are infrequent. Even punishment of noncooperators does not entirely alleviate the dilemma. One important mechanism for achieving cooperative outcomes in such social systems is conformist bias where individuals copy the behavior performed by the majority of their group mates. Conformist bias enhances group competition by both stabilizing behaviors within groups and increasing variance between groups. Due to this group competition effect, conformist bias is thought to have been an important driver of human social complexity and cultural diversity. However, conformist bias only evolves indirectly through associations with other traits, and I show that such associations are more difficult to obtain than previously expected. Specifically, I show that initial measures of population structure must be strong in order for a strong association between conformist bias and cooperative behaviors (cooperation and costly punishment) to evolve and for these traits to reach high frequencies. Additionally, the required initial level of association does not evolve *de novo* in simulations run over long timescales. This suggests that the coevolution of cooperative behaviors and conformist bias alone may not explain the high levels of cooperation within human groups, though conformist bias may still play an important role in combination with other social and demographic forces.

Keywords: social learning, punishment, metapopulation, F_{ST} , maximum entropy

Email address: jvancleve@uky.edu (Jeremy Van Cleve)

1. Introduction

The emergence of large and complex social systems in human and other animals remains a challenging problem of interest to evolutionary biologists and social scientists. Whereas small-scale low cost interactions (e.g., reciprocal exchanges of grooming in baboons; Barrett et al., 1999) or behaviors primarily among kin (e.g., offspring provisioning in many animals) can be explained relatively simply as the outcome of reciprocity (Trivers, 1971; Axelrod and Hamilton, 1981) or kin selection (Hamilton, 1964), interactions in larger groups present special problems. For example, large multi-male groups can be stable in humans (Hill et al., 2011) and chimpanzees (Lehmann and Boesch, 2004) despite the potential for intra-group competition among unrelated males for mates or among males and females for resources. Punishment and sanctioning mechanisms can help explain the stability of cooperation in human social interactions in the laboratory (e.g., Fehr and Gächter, 2000) or cross-culturally (e.g., Henrich et al., 2006). However, since punishment is likely costly, it can be difficult to evolve *de novo* for the same reason as costly cooperation (Sigmund et al., 2001; Boyd et al., 2003; Gardner et al., 2007); this phenomenon is often referred to as the “second-order free-rider problem” (Heckathorn, 1989).

Social behavior in human and primate groups is enriched by the ability of individuals to transmit social information within and between generations, often in the form of “traditions” or “culture” (Whiten, 2005; Laland and Janik, 2006; Whiten and van Schaik, 2007). Generally, social transmission or social learning (Giraldeau et al., 1994) can be favored by natural selection when environmental conditions are uncertain (Boyd and Richerson, 1985; Rogers, 1988; Feldman et al., 1996; Galef and Giraldeau, 2001). The complex information transmitted between individuals by social learning can lead to novel evolutionary dynamics that may help explain the stability of large-scale social interactions in humans (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985).

One of the most studied social learning mechanisms in the context of human cooperation is conformist bias (Asch, 1956; Boyd and Richerson, 1985; Bernheim, 1994; Henrich and Boyd, 1998; Efferson et al., 2008; Claidière and Whiten, 2012). Conformist bias (which can build conformity at the population level) is the tendency of individuals to copy the majority or modal behavior in their social group. It has been documented numerous times in cross-cultural studies of human populations (e.g., Bond and Smith, 1996) and even in some animals, such as stickleback fish (Pike and Laland, 2010) and capuchin monkeys (Dindo et al., 2009). Like other social learning mechanisms, conformist bias facilitates the spread of adaptive behaviors once they become common in a population. This notion underlies mathematical models that show conformist bias can evolve when in competition with unbiased social learning (Boyd and Richerson, 1985; Henrich and Boyd, 1998; Nakahashi et al., 2012) given that the environment varies spatially across different groups (however,

temporal variation tends to select against conformist bias; see refs Nakahashi, 2007; Wakano and Aoki, 2007; Kendal et al., 2009; Nakahashi et al., 2012).

Among social learning mechanisms, conformist bias is particularly good at increasing cultural homogeneity within groups and cultural diversity between groups (Boyd and Richerson, 1985). Multilevel or group selection is strong when trait differences between groups are much stronger than within groups (Price, 1972; Hamilton, 1975; Slatkin and Wade, 1978; Akçay and Van Cleve, 2012), which suggests that conformist bias could facilitate the evolution of group beneficial behaviors (Boyd and Richerson, 1985, 1990; Henrich and Boyd, 2001; Guzmán et al., 2007; Boyd et al., 2011). Using this reasoning, Boyd, Richerson, and Henrich have argued persuasively that conformist bias could be the engine in a cultural group selection process that is responsible both for the rich cultural variation across human groups and for the exceptional levels of cooperation achieved within them (Boyd and Richerson, 1985, 1990; Henrich and Boyd, 1998, 2001). Other researchers have also argued for the importance of group selection and group level processes, such as inter-group conflict (Alexander, 1990; Bowles, 2006) and egalitarian resource distributions within groups (Bowles, 2006; Boehm, 2012), in the evolution of both human cooperation and human culture.

Despite this spirited interest in how group structure can create conditions for cooperative behavior and conformist bias to coevolve, little research has been done to quantitatively explore this coevolutionary process. This lack of research may derive in part from the intuition that conformist bias can stabilize cooperation and punishment of noncooperation (Henrich and Boyd, 2001; Boyd et al., 2011): when potential noncooperators are conformists and cooperators who punish are in the majority in the group, noncooperators will switch to cooperation even at a potential decrease in their payoff from the social interaction. Equally however, conformist bias will stabilize any strategy in the social interaction (e.g., noncooperation) so long as it is in the majority (Henrich and Boyd, 1998). When social interactions occur in a panmictic population, this implies that there is no direct natural selection on conformist bias. Selection on conformist bias must then be indirect through a statistical association between the trait underlying conformity and the trait for the social behavior (Kirkpatrick et al., 2002). Such an association can be due to a population wide correlation between the two traits within an individual or due to population subdivision that increases the chance that a trait sampled from different individuals is of the same type when those individuals are in the same local group or deme; in population genetics, the former association is measured by linkage disequilibrium (LD) and the latter by Wright's F_{ST} (Wright, 1951). These measures of association are general and do not depend on whether traits are cultural or genetic in nature¹. The question of whether conformist bias coevolves with cooperative

¹Though I will use LD to measure between-trait associations and F_{ST} to measure associations between individuals within a group, I will usually refer to them informally as “between-trait” and “within-group” associations, respectively, in the event that references

behaviors (e.g., cooperation and punishment of noncooperation) then hinges on whether and under what conditions positive associations between conformist bias and cooperation can emerge.

Theory for when positive associations can build up between conformist bias and cooperation does not yet exist. One prominent study by Guzmán et al. (Guzmán et al., 2007) finds strong evidence that conformist bias, cooperation, and punishment can evolve in a metapopulation across a range of group sizes but assumes the existence of an initial positive association between conformist bias, cooperation, and punishment. However, an initial positive association may not be required as selection or drift in the metapopulation could build up such an association endogenously. Conversely, an initial positive association may be crucial if the metapopulation process is unlikely to produce such an association on its own. Support for this idea comes from a model of conformist bias where individuals can help group members and groups are randomly formed every generation (Lehmann and Feldman, 2008); random group formation erases initial genetic associations, and conformist bias and helping do not increase together to high frequency in this model.

My goal in this study is to determine whether and when the metapopulation process can build up a strong enough positive association so that conformist bias, cooperation, and punishment can coevolve to high frequency. Using a metapopulation simulation of cooperation, noncooperation, and punishment within groups and conflict between groups (Boyd et al., 2003; Guzmán et al., 2007), I find that conformist bias, cooperation, and punishment coevolve to high frequency only when there is a strong initial positive association between the traits. This result helps reconcile the conflicting results of Guzmán et al. (2007) and Lehmann and Feldman (2008) since the former study assumes there is an initial positive association between the traits whereas the latter study assumes no such association can build up. Without a significant initial positive association, neither conformist bias nor cooperation and punishment evolve to high levels, even after 20,000 generations (400–600k years for a human generation time of 20–30 years (Fenner, 2005; Tremblay and Vézina, 2000)). This suggests that if the metapopulation process can generate a positive association *de novo* between conformist bias and cooperative behaviors, it does so only on very long timescales. Thus, coevolution between conformist bias and cooperative behaviors alone may not be sufficient to explain the high levels of conformity (Gelfand et al., 2011) and cooperative behaviors in current human groups (Henrich et al., 2006).

2. Methods

I describe the metapopulation simulation in detail in Appendix A and how I control for initial associations in Appendix B. I summarize those sections briefly here.

to LD and F_{ST} would incorrectly imply to the reader that the cooperation and punishment loci must be genetic and not cultural.

2.1. Metapopulation simulation

The metapopulation consists of $n = 128$ groups each with N haploid individuals (N varies between 4 and 256). Each individual has three phenotypic loci, each with two possible types: (1) cooperation or noncooperation in the social interaction, (2) punishment of noncooperators in the group or no punishment, and (3) use of a conformist or payoff-based social learning rule. Following previous work on the evolution of cooperation, punishment, and conformist social learning (Boyd et al., 2003; Guzmán et al., 2007), the metapopulation has the following lifecycle (see Fig. 1; Table 1 summarizes the parameters and their values in the simulation). First, individuals participate in a public goods type social interaction where cooperators pay a cost for cooperation (c_C), punishers pay a cost for punishing (c_R) noncooperators, and noncooperators pay a cost for getting punished (c_P). Cooperators mistakenly do not cooperate with probability e . The payoff individuals accumulate from the social interaction represent an individual resource pool used to produce offspring later in the lifecycle. Groups then experience a period of conflict where a random pair of groups (say, x and y) fights with probability ϵ . Groups are more likely to prevail in the conflict if they accumulate more resources, which occurs if they have a higher frequency of cooperators. The probability group x wins the fight is $\frac{1+p_{x1}-p_{y1}}{2}$, where p_{x1} and p_{y1} are the frequency of cooperators (trait 1) in groups x and y respectively. The winning group eliminates the losing group and produces new individuals to colonize the empty patch. Next, individuals experience a social learning phase where conformists copy the modal cooperation and punishment strategies in their group and payoff-based learners can copy a random individual in their group. A payoff-based learner copies the cooperation and punishment strategy of a random partner with probability $1 - \mu$ times the partner's relative payoff in the pair. With probability μ , the payoff-based learner "innovates" and randomly picks cooperation and punishment strategies. After learning, some individuals die and others will have an opportunity to reproduce and fill the empty patch left by the dying individuals; each individual dies with probability q and is replaced by the offspring of two parents sampled randomly from the group with probability proportional to their relative payoff from the social interaction. The phenotypic loci of the offspring each mutate to a random type with probability ν . Finally, each individual in the group can be replaced by an immigrant from another group with probability m times the relative payoff of the immigrant.

2.2. Distribution of initial population states

Each simulation is initialized with a frequency of approximately $1/n = 1/128$ cooperation, punishment, and conformist types; in other words, one group in the initial metapopulation *could* be constructed of individuals who all cooperate, punish, and conform. My goal is distribute these types across the initial metapopulation in a way that allows us to constrain average initial statistical associations within and between types.

The two initial associations I constrain are F_{ST} , which measures the association *between* individuals within a group with respect to a given phenotype or genotype (Wright, 1951; Weir and Hill, 2002), and within-group normalized linkage disequilibrium D' (Lewontin, 1964; Robinson et al., 1991), which measures the association between two types *within* individuals. More informally, F_{ST} is a measure of population structure since structure can increase the degree of trait association between individuals within a group. In contrast, D' measures the degree of phenotypic or genotypic structure since nonzero values of D' indicate types are non-randomly distributed within individuals. In theoretical models with specific population structures, F_{ST} is determined by parameters such as group size and migration rate. In this study, F_{ST} is measured empirically from the distribution of types in the population at a particular time.

Since I want to isolate the effect of these initial associations, I do not want to assume anything about a biological process that might build these associations. The appropriate distribution to use in this case is a maximum entropy (MaxEnt) distribution, which yields the most random pattern (as measured by information entropy) given a set of constraints (Jaynes and Bretthorst, 2003; Frank, 2009). As calculating the maximum entropy distribution is too difficult in this case, I heuristically approximate it (see Appendix B). Briefly, I find the MaxEnt distribution of allele frequencies within groups given the mean initial F_{ST} at each locus and assuming that allele frequencies can take only one of two values (a simplifying assumption for analytic tractability). I also assume for tractability that the covariance between traits between groups is high, which is favorable to the coevolution of cooperation and conformist bias in this context. Given the MaxEnt allele frequencies calculated from F_{ST} and the three initial pairwise D' values, I calculate group genotype frequencies using standard population genetics equations (Robinson et al., 1991). I then take the heuristic MaxEnt distribution to be the product of multinomial distributions with parameters N and the genotype frequencies for each group; the multinomial is the MaxEnt distribution if the expected genotype frequencies within each group are constrained.

The initial distribution of genotypes in each metapopulation simulation is a sample drawn from this heuristic MaxEnt distribution. Thus, the initial frequencies of cooperation, punishment, and conformist types in the metapopulation simulations are on average $1/n = 1/128$ but vary in each simulation due to random sampling; likewise, the initial F_{ST} and D' associations are on average the parameteric values input into the heuristic MaxEnt distribution with some variation due to random sampling. After initialization, the allele frequencies and associations evolve according to the metapopulation process described above.

2.3. Parameter values

The parameter values used in the simulation are listed in Table 1 and are the same ones used in the previous study of the coevolution of conformist bias and cooperation by Guzmán et al. (Guzmán et al., 2007) (as well as a previous study of cooperation and punishment; Boyd et al., 2003). In short, these parameter values assume that the cost of cooperating and punishing noncooperators are equivalent and that the cost of being punished is relatively high. Mutations at all loci occur relatively slowly compared to the rate of random changes at the cooperation and punishment loci due to mistakes and innovations. Given the number of groups ($n = 128$) and the probability of group conflict ($\epsilon = 0.015$), there is an average of one group conflict every time step. Finally, the mortality rate of $q = 0.1$ implies a generation time of 10 time steps. Recent anthropological surveys measuring maternal age at first reproduction estimate human generation time to be approximately 30 years (Tremblay and Vézina, 2000; Fenner, 2005) and older studies use a value of 20 years (Tremblay and Vézina, 2000), which yields an interval of 2–3 years per time step. Simulations results in Figs. 2 and 4 are for 10^5 and 2×10^5 time steps, respectively, which are equivalent to ranges of 200–300k years and 400–600k years, respectively.

3. Results

Using the heuristic MaxEnt method described in section 2, I simulated the metapopulation process for initial within-group associations (F_{ST} values) of 0.01–0.99 at each of the three phenotypic loci (cooperation, punishment, and conformist bias) and D' values of -0.5 , 0.0 , and 0.5 for each of the three pairwise normalized linkage disequilibria. Fig. 2 uses box plots to display the distribution of the cooperation, punishment, and conformist bias alleles among the last 2000 of 10^5 time steps for 250 replicate simulations; plus symbols denote the maximum and minimum of the distribution. The horizontal axis shows the initial within-group associations for each phenotypic locus. The initial between-trait association is zero. Four different group sizes are displayed, $N = 4, 16, 64$, and 256 . For large group sizes ($N = 64$ and 256), a clear pattern emerges where cooperation, punishment, and conformist bias alleles reach high frequency so long as the initial within-group association is greater than a threshold value of 0.2 . Smaller group sizes show a similar pattern but with a higher threshold value of within-group association and more variation in the final allele frequencies. If there is no initial within-group association at the phenotypic loci, then cooperation, punishment, and conformist bias are lost entirely from the metapopulation (within-group association ≤ 0.2 in Fig. 2). For reference, measures of genetic differentiation, including F_{ST} , in indigenous human groups rarely reach even 0.1 (Langergraber et al., 2011). A realistic comparison between empirical human F_{ST} values and simulation results, however, would require that the simulation be calibrated more closely to human ancestral demography.

To get more detail on how the initial distribution of types explains the results in Fig. 2, I plot in Fig. 3 the distribution of the initial within-group frequency of the conformist bias allele from the simulations in Fig 2; plots of the initial distribution of the cooperation and punishment types are virtually identical. Each of the five curves represents a different initial within-group association value: starting from dark to light (purple to yellow), $F_{ST} = 0.01, 0.2, 0.5, 0.8$, and 0.99 . Since the initial allele frequency is $1/128$ on average for each of the three phenotypic loci, many groups have few or no conformist alleles; to better resolve the group with an appreciable frequency, Fig. 3 plots the distribution for groups with a conformist allele frequency greater than 0.1 . For all four of the group sizes, low values of the initial within-group association result in a more even distribution of the conformist bias allele across the metapopulation where no groups attain significant frequency of the allele. As the initial within-group association increases, the conformist bias alleles are concentrated in a few groups with higher frequency of the allele. A strong initial association of $F_{ST} = 0.99$ results in one group initially containing all or nearly all the conformist bias alleles. Recall from Fig. 2 that for large group sizes ($N = 64$ and 256) an initial within-group association of 0.2 or larger is required to drive the cooperation, punishment, and conformist bias types to high frequency. The reason this threshold value is 0.2 may be explained in part by observing from Fig. 3 that an initial within-group association of at least 0.2 is required to obtain some groups with an initial frequency 0.5 for the conformist bias, cooperation, and punishment alleles. In such groups, individuals with the conformist bias allele will copy the cooperation and punishment types since those types are likely in the majority in those groups.

In comparison to the simulations presented here, the simulations of Guzmán et al. (Guzmán et al., 2007) begin with a single group in the metapopulation composed entirely of cooperators who punish and conform; the remaining 127 groups contain noncooperators who don't punish or conform. Thus, their simulations assume a high initial within-group association, and their results (Fig. 1 in Guzmán et al., 2007) match mine in that range. However, my results show that once a high initial within-group association is not assumed, cooperation, punishment, and conformist bias do not evolve to high levels regardless of the group size. My simulations cover a range of 10^4 generations or at least 200k years, whereas the simulations of Guzmán et al. run for only 200 generations or 4000 years. Thus, the requirement of a high initial within-group association in order for the coevolution of cooperation and conformist bias holds even over a large timescale. This matches the result of Lehmann and Feldman (Lehmann and Feldman, 2008) whose random group formation model lacks an initial within-group association as measured by F_{ST} and resists the buildup of within-group association due to the lack of limited dispersal between groups.

Comparing the results in Fig. 2 for no average initial between-trait association ($D' = 0$) to results for positive ($D' = 0.5$ in Supplementary Fig. S1) and negative ($D' = -0.5$ in Supplementary Fig. S2) between-trait

associations reveals that initial between-trait association has relatively little effect on whether cooperation, punishment, and conformist bias coevolve relative to the effect of initial F_{ST} . A positive between-trait association marginally decreases the threshold within-group association necessary for the three traits to coevolve, with a slightly larger decrease occurring for low group sizes. Similarly, a negative between-trait association increases the threshold within-group association, though more so at large group sizes. The marginal effect of between-trait associations could be due in part to the skewed initial trait frequencies that make it difficult to achieve extreme between-trait association values (compare Figs. S5 and S6).

In populations with stochastic effects such as genetic drift and mutation, stronger selection typically results in a larger chance that alleles close to a fitness optimum will be of high frequency or will fix. Thus, stronger selection typically reduces the variance in allele frequency when averaging over replicate simulations. This means that the widths of the interquartile intervals in Fig. 2 reflect, in part, the strength of selection on cooperation, punishment and conformist bias. Comparing the width of the interquartile intervals in Fig. 2, there is significantly more variation in the frequency of the conformist bias allele than in the frequency of the cooperation or punishment alleles. This suggests that there is significantly weaker selection on conformist bias. As discussed in the section 1, this is a result of the fact that selection on conformist bias occurs indirectly via an association with cooperation. The strength of such selection is proportional to the product of a selection coefficient times the between-trait association (LD) between the conformist bias and cooperation. Since the between-trait association is less than complete, particularly for smaller populations, indirect selection on conformist bias tends to be weaker than direct selection on cooperation and punishment.

In order to obtain a more refined picture of the importance of initial within-group associations, I vary the within-group association at the conformist bias locus independently from within-group associations at the cooperation and punishment loci. This might occur biologically when prior selection affected these loci differently (resulting in the divergence in F_{ST}), and it is useful conceptually in that it allows us to determine more specifically whether it is the co-occurrence of only cooperation and punishment together in the same group that is required to jump start the coevolution of cooperation, punishment and conformist bias, or whether all three traits must co-occur for such coevolution. Fig. 4 shows the median allele frequency (over the last 2000 of 2×10^5 time steps for 250 replicate simulations) for the three traits where the within-group association at the cooperation and punishment loci varies on the horizontal axis and the within-group association at the conformist bias locus varies on the vertical axis. For large groups ($N = 64$ and 128), within-group association at the conformist bias locus has no effect on whether cooperation and conformist bias coevolve. Rather, within-group association at the cooperation and punishment loci must be larger than 0.2 in order for cooperation and conformist bias to coevolve. The lack of dependence on the initial within-group association

at the conformist bias locus is likely due to the fact that whether conformist bias can stabilize cooperation or not depends primarily on the within-group association at the cooperation and punishment loci.

In contrast, when group size is small ($N = 4$ and 16), the initial within-group association at the conformist bias locus becomes more important. For $N = 16$, the within-group association at the conformist bias locus must be high enough in order for conformist bias to evolve to significant frequencies. For the smallest group size, $N = 4$, a threshold level of conformist bias must be met for both high levels of conformist bias and cooperation to evolve. This dependence at small group sizes on the initial within-group association for conformist bias is likely due to a combination of increased stochasticity in small groups and the reduced effect of intergroup conflict on the frequency of cooperation in the metapopulation. Once intergroup competition is weak enough, the combined initial association between cooperation and conformist bias becomes the important condition.

4. Discussion

Social learning is thought to be one of the key drivers in the development of large and complex social systems. Among such mechanisms, conformist bias can explain both the consistency of behaviors within groups that forms the basis of norms and traditions, as well as the diversity of behaviors between groups that forms cultural diversity. Together, these effects might stabilize cooperative norms within groups and benefit such groups when in competition with groups of noncooperators. However, I show that cooperative behaviors and conformist bias can coevolve only when the initial association between cooperation and punishment is sufficiently strong in terms of within-group association or F_{ST} , which is a form of relatedness (Rousset, 2002). When cooperation, punishment, and conformist bias are initially rare, a high within-group association can ensure that at least one group contains a sufficient number of cooperators and punishers for both types to increase in frequency. As cooperators and punishers spread in the metapopulation, conformist bias can stabilize cooperative groups, those groups can profit in group competition, and the association between cooperation and conformist bias can increase. Over the long run, this can lead to a high median frequency of cooperation, punishment, and conformist bias. If the initial within-group association is not larger than the threshold (≈ 0.2 for larger populations), cooperation, punishment, and conformist bias do not evolve over a timescale of at least 2×10^4 generations or 400–600k years for a human generation time of 20–30 years (Fenner, 2005; Tremblay and Vézina, 2000). Notably, significant genetic and cultural change can occur on timescale of the simulation (e.g., the human–Neanderthal split is ~ 650 k years (Prüfer et al., 2014)). This suggests that it may be unlikely that current levels of cooperation or cultural complexity evolved in ancestral human groups purely due to the coevolution of cooperative behaviors and conformist bias.

The existence of a substantial threshold within-group association for the coevolution of cooperative behaviors and conformist bias over a biologically significant timescale is particularly important because of what it reveals in this model about the interaction between natural selection and genetic drift. Natural selection appears to create positive frequency dependent selection where selection favors metapopulation states with either high or low levels of cooperation, punishment, and conformist bias. On the other hand, genetic drift creates a buildup of within-group association given the group structure, and the association will cross the required threshold where enough cooperators and punishers are in the same group given enough time. This will lead to the metapopulation evolving towards a high level of cooperation, punishment, and conformist bias. Thus, the important question is how long it takes for genetic drift to buildup this level of within-group association given the strength of selection against cooperation and punishment when they are at low frequency.

To get a sense for the relative strengths of drift and selection in this case, I plot in Fig. 5 the time required for the within-group association (F_{ST}) to approach half of its equilibrium value in an infinite island model (Wright, 1931) for different group sizes and migration rates. The solid portions of curves represent combinations of group size and migration rate that result in an equilibrium $F_{ST} > 0.2$, which is the minimum threshold required for the coevolution of cooperative behaviors and conformist bias in Figs. 2 and 4. The plot shows that the timescale required to approach the threshold within-group association (in an island model with the equivalent effective population size and migration rate) is easily an order of magnitude smaller than the timescale of the simulations (even after accounting for a generation time of 10 time steps per generation in the simulation). This suggests that in simulations where the initial within-group association is less than the threshold, selection inhibits the buildup of within-group association to this threshold over the timescale of the simulation. Since the timescale of the simulation is comparable to that thought to be important for the emergence of human social complexity (Marean et al., 2007; Powell et al., 2009), conformist bias alone may not explain the evolution of current levels of cooperation in large human social groups.

A number of important assumptions in the metapopulation model shape the likelihood that the metapopulation will evolve a high level of cooperation, punishment, and conformist bias. For example, the rate at which cooperative groups outcompete noncooperative ones is determined in part by the probability or rate of group conflict, ϵ . Increasing the rate of group conflict will lead to cooperative groups proliferating faster in the metapopulation. However, such cooperative groups still rely on an initial within-group association in order to spread, and thus a threshold initial within-group association will still exist, though the threshold could decrease with increasing group conflict. Another important assumption is that social learning only happens once in the lifecycle (see Fig. 1). Alternatively, social learning might occur many times in the lifecycle, which

will accelerate cultural change relative to genetic change. Assuming that cooperation and punishment are cultural loci and the social learning type is genetic, this suggests that cooperation and punishment or non-cooperation might quickly fix in groups relative to the rate of change in the frequency of conformist bias. Even in this scenario however, a threshold initial within-group association is still required for cooperation and conformist bias to coevolve, and its not clear whether faster cultural change relative genetic change will increase or decrease this threshold. Finally, the metapopulation simulation borrows its payoff-based learning strategy from previous work (Boyd et al., 2003; Guzmán et al., 2007) where individuals can copy the successful strategy of only one individual in their group. However, payoff-based learners might be able to observe the strategies of many individuals in their group and thus copy the individual in the group with the highest payoff. This kind of payoff-based learning might be more resistant to the invasion of conformist bias than the payoff-based learning used here.

Although the coevolution of conformist bias and cooperative behaviors alone may not produce high levels of cooperation without a substantial initial within-group association, conformist bias may still be an important mechanism stabilizing social norms for cooperative behavior. For example, it could be the case the cooperative tendencies evolve first, likely alongside a stabilizing mechanism such as punishment (Axelrod, 1986; Sethi, 1996; Boyd et al., 2003, 2010), reciprocity (Axelrod and Hamilton, 1981; Boyd and Richerson, 1988, 1989; Nowak and Sigmund, 1998; Leimar and Hammerstein, 2001), or behavioral responsiveness (André and Day, 2007; Akçay et al., 2009; Akçay and Van Cleve, 2012), and then conformist bias evolves to increase the level of cooperation. However, selection on conformist bias in such cases will still be indirect and will likely occur more slowly than direct selection due to the reliance on an association with cooperation and punishment.

Alternatively, there might be direct selection on conformist bias, which could evolve independently of cooperation or any other behavior. One possible source of such selection could come from spatially variable environments where conformist bias, as a form of social learning, helps spread adaptive behaviors (Boyd and Richerson, 1985; Henrich and Boyd, 1998; Nakahashi et al., 2012). Notably, this advantage of conformist bias disappears when environments are temporally variable (Eriksson et al., 2007; Nakahashi et al., 2012) and knowledge of all possible cultural variants is limited (Eriksson et al., 2007). Another possibility is that coordinating on tasks, such as hunting, predator defense, or flocking, requires quick decision making, and there is a cost to coordination failure regardless of which behavior the group eventually coordinates on. This kind of payoff structure defines coordination (Binmore, 2007) or stag-hunt games (Skyrms, 2001), and creates a kind of fitness valley (Van Cleve and Lehmann, 2013). If conformism reduces the risk of coordination failure, it could provide a direct benefit to individuals when performing coordination tasks by ensuring that

they end up at one of the fitness peaks. In a coordination game where the two strategies are cooperation and noncooperation, cooperation is the higher fitness peak, and populations evolve to that peak so long as the “risk” of cooperation is low or relatedness is high (Van Cleve and Lehmann, 2013). Regardless of its source, any direct selection for conformist bias still poses a problem for the origination of cooperation since conformity can inhibit the increase of initially rare cooperative tendencies (Lehmann and Feldman, 2008). Thus, even in this case, a complete understanding of high levels of cooperation and conformist bias still requires an explicit model that accounts the role of initial statistical associations and how they evolve in response to both deterministic and stochastic forces.

Finally, the dependence of the evolution of conformist bias on indirect selection through a statistical association with another trait is a general feature of many complex coevolutionary problems. For example, genetic recombination breaks down associations between alleles at different loci regardless of whether such alleles confer a selective advantage or not. When there is negative between-trait association between beneficial alleles, increased recombination can be favored by selection since it increases the likelihood that beneficial alleles will be brought together (Otto and Lenormand, 2002; Barton, 2010); LD is consistently negative when there is negative epistasis (Barton, 1995) or genetic drift (Hill and Robertson, 1966; Barton, 2010). In addition to conformist bias and recombination, traits that change associations at other loci, without depending on the specific alleles at those loci, include migration and assortative and disassortative mating. When these traits lack direct selection, they evolve through indirect selection on their association with other loci, and the same coevolutionary issue arises of determining the conditions under which such associations are created or broken down. Thus, properly capturing the complexity of these coevolutionary problems requires a detailed analysis of the interplay between direct and indirect selection, recombination, and often genetic drift. Such analyses are particularly important for studying complex social behaviors in humans and other animals since these behaviors often interact in ways that make associations between them pivotal to understanding how they evolve.

Acknowledgements

JVC was supported by a Santa Fe Institute Omidyar Fellowship and the National Evolutionary Synthesis Center (NESCent), NSF #EF-0423641. Special thanks to Erol Akçay, Laura Fortunato, Anne Kandler, and Laurent Lehmann for insightful comments on the manuscript and to two reviewers whose comments improved the manuscript.

References

- 372 Akçay, E., Van Cleve, J., 2012. Behavioral responses in structured populations pave the way to group opti-
 374 mality. *Am. Nat.* 179 (2), 257–269.
- 374 Akçay, E., Van Cleve, J., Feldman, M. W., Roughgarden, J., 2009. A theory for the evolution of other-regard
 integrating proximate and ultimate perspectives. *Proc. Natl. Acad. Sci. U. S. A.* 106 (45), 19061–19066.
- 376 Alexander, R. D., 1990. *How Did Humans Evolve?* The University of Michigan Museum of Zoology, Ann
 Arbor, MI.
- 378 André, J.-B., Day, T., 2007. Perfect reciprocity is the only evolutionarily stable strategy in the continuous
 iterated prisoner's dilemma. *J. Theor. Biol.* 247 (1), 11–22.
- 380 Asch, S. E., 1956. Studies of independence and conformity: I. A minority of one against a unanimous majority.
Psychol. Monogr. 70 (9), 1–70.
- 382 Axelrod, R., 1986. An evolutionary approach to norms. *Amer. Polit. Sci. Rev.* 80 (4), pp. 1095–1111.
- Axelrod, R., Hamilton, W. D., 1981. The evolution of cooperation. *Science* 211 (4489), 1390–1396.
- 384 Barrett, L., Henzi, S. P., Weingrill, T., Lycett, J. E., Hill, R. A., 1999. Market forces predict grooming reciprocity
 in female baboons. *Proc. R. Soc. B* 266 (1420), 665–670.
- 386 Barton, N. H., 1995. A general model for the evolution of recombination. *Genet. Res.* 65 (2), 123–45.
- Barton, N. H., 2010. Genetic linkage and natural selection. *Philos. Trans. R. Soc. B* 365 (1552), 2559–69.
- 388 Bernheim, B. D., 1994. A theory of conformity. *J. Polit. Economy* 102 (5), 841–877.
- Binmore, K. G., 2007. *Playing for Real: A Text on Game Theory*. Oxford University Press, Oxford.
- 390 Boehm, C., 2012. Ancestral hierarchy and conflict. *Science* 336 (6083), 844–847.
- Bond, R., Smith, P. B., 1996. Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956)
 392 line judgment task. *Psychol. Bull.* 119 (1), 111–137.
- Bowles, S., 2006. Group competition, reproductive leveling, and the evolution of human altruism. *Science*
 394 314 (5805), 1569–1572.
- Boyd, R., Gintis, H., Bowles, S., 2010. Coordinated punishment of defectors sustains cooperation and can
 396 proliferate when rare. *Science* 328 (5978), 617–20.

- Boyd, R., Gintis, H., Bowles, S., Richerson, P. J., 2003. The evolution of altruistic punishment. *Proc. Natl. Acad. Sci. U. S. A.* 100 (6), 3531–3535.
- Boyd, R., Richerson, P. J., 1985. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Boyd, R., Richerson, P. J., 1988. The evolution of reciprocity in sizable groups. *J. Theor. Biol.* 132 (3), 337–56.
- Boyd, R., Richerson, P. J., 1989. The evolution of indirect reciprocity. *Social Networks* 11 (3), 213–236.
- Boyd, R., Richerson, P. J., 1990. Group selection among alternative evolutionarily stable strategies. *J. Theor. Biol.* 145 (3), 331–342.
- Boyd, R., Richerson, P. J., Henrich, J., 2011. Rapid cultural adaptation can facilitate the evolution of large-scale cooperation. *Behav. Ecol. Sociobiol.* 65 (3), 431–444.
- Cavalli-Sforza, L. L., Feldman, M. W., 1981. *Cultural Transmission and Evolution: A Quantitative Approach*. Vol. 16. Princeton University Press, Princeton, N.J.
- Claidière, N., Whiten, A., 2012. Integrating the study of conformity and culture in humans and nonhuman animals. *Psychol. Bull.* 138 (1), 126–145.
- Dindo, M., Whiten, A., de Waal, F. B. M., 2009. In-group conformity sustains different foraging traditions in capuchin monkeys (*Cebus apella*). *PLoS ONE* 4 (11), e7858.
- Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., Lubell, M., 2008. Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evol. Hum. Behav.* 29 (1), 56 – 64.
- Eriksson, K., Enquist, M., Ghirlanda, S., 2007. Critical points in current theory of conformist social learning. *J Evol Psychol* 5 (1), 67–87.
- Fehr, E., Gächter, S., 2000. Cooperation and punishment in public goods experiments. *Am. Econ. Rev.* 90 (4), 980–994.
- Feldman, M. W., Aoki, K., Kumm, J., 1996. Individual versus social learning: Evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* 104 (3), 209–231.
- Fenner, J. N., 2005. Cross-cultural estimation of the human generation interval for use in genetics-based population divergence studies. *Am. J. Phys. Anthropol.* 128 (2), 415–423.
- Frank, S. A., 2009. The common patterns of nature. *J. Evol. Biol.* 22 (8), 1563–85.

Galef, Jr, B. G., Giraldeau, L.-A., 2001. Social influences on foraging in vertebrates: causal mechanisms and
adaptive functions. *Anim. Behav.* 61 (1), 3–15.

Gardner, A., West, S. A., Barton, N. H., 2007. The relation between multilocus population genetics and social
evolution theory. *Am. Nat.* 169 (2), 207–226.

Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., Duan, L., Almaliach, A., Ang, S., Ar-
nadottir, J., Aycan, Z., Boehnke, K., Boski, P., Cabecinhas, R., Chan, D., Chhokar, J., D’Amato, A., Ferrer,
M., Fischlmayr, I. C., Fischer, R., Fülöp, M., Georgas, J., Kashima, E. S., Kashima, Y., Kim, K., Lempereur,
A., Marquez, P., Othman, R., Overlaet, B., Panagiotopoulou, P., Peltzer, K., Perez-Florizno, L. R., Pono-
marenko, L., Realo, A., Schei, V., Schmitt, M., Smith, P. B., Soomro, N., Szabo, E., Taveesin, N., Toyama,
M., Van de Vliert, E., Vohra, N., Ward, C., Yamaguchi, S., 2011. Differences between tight and loose cul-
tures: a 33-nation study. *Science* 332 (6033), 1100–4.

Giraldeau, L.-A., Caraco, T., Valone, T. J., 1994. Social foraging: individual learning and cultural transmission
of innovations. *Behav. Ecol.* 5 (1), 35–43.

Guzmán, R. A., Rodríguez-Sickert, C., Rowthorn, R., 2007. When in Rome, do as the Romans do: the coevo-
lution of altruistic punishment, conformist learning, and cooperation. *Evol. Hum. Behav.* 28 (2), 112–117.

Hamilton, W. D., 1964. The genetical evolution of social behaviour. I. *J. Theor. Biol.* 7 (1), 1–16.

Hamilton, W. D., 1975. Innate social aptitudes of man: an approach from evolutionary genetics. In: Fox, R.
(Ed.), *Biosocial Anthropology*. Malaby Press, London, pp. 133–155.

Heckathorn, D. D., 1989. Collective action and the second-order free-rider problem. *Ration. Soc.* 1 (1), 78–
100.

Henrich, J., Boyd, R., 1998. The evolution of conformist transmission and the emergence of between-group
differences. *Evol. Hum. Behav.* 19 (4), 215 – 241.

Henrich, J., Boyd, R., 2001. Why people punish defectors. weak conformist transmission can stabilize costly
enforcement of norms in cooperative dilemmas. *J. Theor. Biol.* 208 (1), 79–89.

Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M.,
Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., Ziker, J., 2006. Costly punishment across
human societies. *Science* 312 (5781), 1767–70.

450 Hill, K. R., Walker, R. S., Bozicević, M., Eder, J., Headland, T., Hewlett, B., Hurtado, A. M., Marlowe, F.,
 Wiessner, P., Wood, B., 2011. Co-residence patterns in hunter-gatherer societies show unique human social
 452 structure. *Science* 331 (6022), 1286–9.

Hill, W. G., Robertson, A., 1966. The effect of linkage on limits to artificial selection. *Genet. Res.* 8 (03),
 454 269–294.

Jaynes, E. T., Bretthorst, G. L., 2003. *Probability Theory: The Logic of Science*. Cambridge University Press,
 456 Cambridge, UK.

Kendal, J., Giraldeau, L.-A., Laland, K., 2009. The evolution of social learning rules: payoff-biased and
 458 frequency-dependent biased transmission. *J. Theor. Biol.* 260 (2), 210–9.

Kirkpatrick, M., Johnson, T., Barton, N. H., 2002. General models of multilocus evolution. *Genetics* 161 (4),
 460 1727–1750.

Laland, K. N., Janik, V. M., 2006. The animal cultures debate. *Trends Ecol. Evol.* 21 (10), 542–547.

462 Langergraber, K., Schubert, G., Rowney, C., Wrangham, R., Zommers, Z., Vigilant, L., 2011. Genetic dif-
 ferentiation and the evolution of cooperation in chimpanzees and humans. *Proc. R. Soc. B* 278 (1717),
 464 2546–52.

Lehmann, J., Boesch, C., 2004. To fission or to fusion: effects of community size on wild chimpanzee (*Pan*
 466 *trogodytes verus*) social organisation. *Behav. Ecol. Sociobiol.* 56 (3), 207–216.

Lehmann, L., Feldman, M. W., 2008. The co-evolution of culturally inherited altruistic helping and cultural
 468 transmission under random group formation. *Theor. Popul. Biol.* 73 (4), 506–516.

Leimar, O., Hammerstein, P., 2001. Evolution of cooperation through indirect reciprocity. *Proc. R. Soc. B*
 470 268 (1468), 745–753.

Lewontin, R. C., 1964. The interaction of selection and linkage. I. General considerations; heterotic models.
 472 *Genetics* 49 (1), 49–67.

Marean, C. W., Bar-Matthews, M., Bernatchez, J., Fisher, E., Goldberg, P., Herries, A. I. R., Jacobs, Z., Jer-
 474 ardino, A., Karkanas, P., Minichillo, T., Nilssen, P. J., Thompson, E., Watts, I., Williams, H. M., 2007.
 Early human use of marine resources and pigment in South Africa during the Middle Pleistocene. *Nature*
 476 449 (7164), 905–908.

Nakahashi, W., 2007. The evolution of conformist transmission in social learning when the environment
478 changes periodically. *Theor. Popul. Biol.* 72 (1), 52–66.

Nakahashi, W., Wakano, J. Y., Henrich, J., 2012. Adaptive social learning strategies in temporally and spatially
480 varying environments. *Hum. Nat.* 23 (4), 386–418.

Nowak, M. A., Sigmund, K., 1998. The dynamics of indirect reciprocity. *J. Theor. Biol.* 194 (4), 561–574.

482 Otto, S. P., Lenormand, T., 2002. Resolving the paradox of sex and recombination. *Nat. Rev. Genet.* 3 (4),
252–61.

484 Pike, T. W., Laland, K. N., 2010. Conformist learning in nine-spined sticklebacks' foraging decisions. *Biol.*
Lett. 6 (4), 466–468.

486 Powell, A., Shennan, S., Thomas, M. G., 2009. Late Pleistocene demography and the appearance of modern
human behavior. *Science* 324 (5932), 1298–301.

488 Price, G. R., 1972. Extension of covariance selection mathematics. *Ann. Hum. Genet.* 35 (4), 485–490.

Prüfer, K., Racimo, F., Patterson, N., Jay, F., Sankararaman, S., Sawyer, S., Heinze, A., Renaud, G., Sudmant,
490 P. H., de Filippo, C., Li, H., Mallick, S., Dannemann, M., Fu, Q., Kircher, M., Kuhlwilm, M., Lachmann,
M., Meyer, M., Ongyerth, M., Siebauer, M., Theunert, C., Tandon, A., Moorjani, P., Pickrell, J., Mullikin,
492 J. C., Vohr, S. H., Green, R. E., Hellmann, I., Johnson, P. L. F., Blanche, H., Cann, H., Kitzman, J. O.,
Shendure, J., Eichler, E. E., Lein, E. S., Bakken, T. E., Golovanova, L. V., Doronichev, V. B., Shunkov, M. V.,
494 Derevianko, A. P., Viola, B., Slatkin, M., Reich, D., Kelso, J., Pääbo, S., 2014. The complete genome sequence
of a Neanderthal from the Altai Mountains. *Nature* 505 (7481), 43–49.

496 Robinson, W. P., Asmussen, M. A., Thomson, G., 1991. Three-locus systems impose additional constraints on
pairwise disequilibria. *Genetics* 129 (3), 925–930.

498 Rogers, A. R., 1988. Does biology constrain culture? *Am. Anthropol.* 90 (4), 819–831.

Rousset, F., 2002. Inbreeding and relatedness coefficients: what do they measure? *Heredity* 88 (5), 371–380.

500 Sethi, R., 1996. Evolutionary stability and social norms. *J. Econ. Behav. Organ.* 29 (1), 113–140.

Sigmund, K., Hauert, C., Nowak, M. A., 2001. Reward and punishment. *Proc. Natl. Acad. Sci. U. S. A.* 98 (19),
502 10757–62.

Skyrms, B., 2001. The stag hunt. *Proc. Adresses Am. Philos. Assoc.* 75 (2), pp. 31–41.

- 504 Slatkin, M., Wade, M. J., 1978. Group selection on a quantitative character. *Proc. Natl. Acad. Sci. U. S. A.*
75 (7), 3531–3534.
- 506 Tremblay, M., Vézina, H., 2000. New estimates of intergenerational time intervals for the calculation of age
and origins of mutations. *Am. J. Hum. Genet.* 66 (2), 651 – 658.
- 508 Trivers, R. L., 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46 (1), 35–57.
- Van Cleve, J., Lehmann, L., 2013. Stochastic stability and the evolution of coordination in spatially structured
510 populations. *Theor. Popul. Biol.* 89, 75–87.
- Wakano, J. Y., Aoki, K., 2007. Do social learning and conformist bias coevolve? Henrich and Boyd revisited.
512 *Theor. Popul. Biol.* 72 (4), 504–12.
- Weir, B. S., Hill, W. G., 2002. Estimating F-statistics. *Annu. Rev. Genet.* 36, 721–750.
- 514 Whiten, A., 2005. The second inheritance system of chimpanzees and humans. *Nature* 437 (7055), 52–55.
- Whiten, A., van Schaik, C. P., 2007. The evolution of animal ‘cultures’ and social intelligence. *Philos. Trans.*
516 *R. Soc. B* 362 (1480), 603–620.
- Wright, S., 1931. Evolution in Mendelian populations. *Genetics* 16 (2), 97–159.
- 518 Wright, S., 1951. The genetical structure of populations. *Ann. Eugen.* 15 (1), 323–354.

Parameter	Value	Definition
N	varies	population size in each group (or deme)
n	128	number of groups
c_C	0.2	cost of cooperation
c_R	0.2	cost of punishing
c_P	0.8	cost of being punished
ε	0.015	probability of group conflict
e	0.02	probability a cooperator makes a mistake and does not cooperate
μ	0.01	probability that imitators randomly innovate
q	0.1	mortality rate
ν	0.001	mutation rate
m	0.01	immigration rate

Table 1: Parameter definitions and values used in the simulations. These values are also used in refs (Guzmán et al., 2007) and (Boyd et al., 2003).

Figure 1: Each time step in the simulation involves one iteration of the above lifecycle, which begins with the social interaction stage and ends with the migration stage. The same lifecycle is used in refs (Guzmán et al., 2007) and (Boyd et al., 2003).

Figure 2: Box plots for cooperation, punishment, and conformist bias allele frequencies given the initial within-group association (F_{ST}) for all three traits on the horizontal axis and an initial within-group between-trait association of zero ($D' = 0.0$). Boxes extend from lower to upper quartile of the distribution of simulation results with a line at the median. Simulations were run for 10^5 time steps and replicated 250 times; the box plots represent the last 2000 time steps of all the replicates. Plus symbols denote the maximum and minimum frequency. Group size is given above each column of plots.

Figure 3: Distributions of the initial within-group frequency of the conformist bias allele in simulations from Fig. 2 plotted as Gaussian kernel density estimates. The curves going from dark to light (purple to yellow) represent initial within-group associations (F_{ST}) of 0.01, 0.2, 0.5, 0.8, and 0.99. Since the mean initial conformist allele frequency is approximately $1/128$, many groups have an initial frequency of zero. Thus, this peak and others at allele frequencies less than 0.1 are excluded from the above plots to better resolve the peaks with allele frequencies greater than 0.1. Plots of the distributions of the initial within-group frequency of the cooperation and punishment alleles are virtually identical.

Figure 4: Median allele frequency for cooperation, punishment, and conformist bias given an initial within-group association (F_{ST}) for cooperation and punishment on the horizontal axis, an initial within-group association for conformist bias on the vertical axis, and an initial between-trait association of zero ($D' = 0.0$) within groups. Group size is given above each column of plots.

Figure 5: Number of generations for the within-group association F_{ST} to approach half of its equilibrium value in an infinite island model of population structure. In such a model, equilibrium $F_{ST} = (1 - m)^2 / (1 + m(2 - m)(N - 1))$ where N is the haploid group size and m is the migration rate. The four curves represent from top to bottom $N = 256, 64, 16$, and 4 . The solid portion of the curves indicate an equilibrium $F_{ST} > 0.2$ and the dotted portions $F_{ST} < 0.2$

Appendix A. Metapopulation simulation

The metapopulation consists of n groups (or demes) each composed of N haploid individuals each with three biallelic loci. The first two loci are cultural in that the traits at those loci can be transmitted vertically from parent to offspring and horizontally between different individuals in the same generation. Traits at the third locus are transmitted only vertically and thus this locus can be either genetic or cultural. I track the state of each individual using the variable p_{ijk} , which is 1 when individual j in group i has allele A_{k1} at locus k and is 0 when that individual has allele A_{k0} . More specifically, I am interested in how functions of p_{ijk} changes over time, such as the mean allele frequency at locus k within each population i , $p_{ik} = \sum_j p_{ijk} / N$, the mean in the whole metapopulation, $p_k = \sum_i p_{ik} / n$, or the diversity within populations at locus k , $H_{1k} = \sum_i 2p_{ik}(1 - p_{ik}) / n$, relative the diversity in the whole metapopulation, $H_{0k} = 2p_k(1 - p_k)$, usually quantified in population genetics as Wright's F_{ST} , $F_{STk} = (H_{0k} - H_{1k}) / H_{0k}$ (Wright, 1951). I am also interested in the correlation between different loci, k and l , which is measured in population genetics by linkage disequilibrium (LD), $D_{kl} = \frac{1}{nN} \sum_{i,j} p_{ijl}p_{ijk} - p_l p_k$. Since it is often useful to normalize LD so that it varies between -1 and 1 , and I use Lewontin's normalization (Lewontin, 1964): $D'_{kl} = D_{kl} / \min(p_l(1 - p_k), p_k(1 - p_l))$ for $D_{kl} > 0$ and $D'_{kl} = D_{kl} / \min(p_l p_k, (1 - p_k)(1 - p_l))$ for $D_{kl} < 0$. For the three-locus system, a similar normalization exists but contains extra constraints (see Robinson et al., 1991).

The simulation uses the following lifecycle given in Guzmán et al. (Guzmán et al., 2007), which is a modification of Boyd et al. (2003).

1. **Social interaction.** Individuals have the opportunity to invest in some public good and can either cooperate and contribute to the good or can not cooperate and not contribute. I assume that cooperation is controlled by a single (cultural) locus with two alleles (A_{11} or cooperate and A_{10} or do not cooperate). Non-contributing members are then punished by punishers. Punishment is controlled by another (cultural) locus with two alleles (A_{21} or punisher and A_{20} or non-punisher). The payoff to individual j in group i is given by (Boyd et al., 2003; Guzmán et al., 2007)

$$\pi_{ij} = z - c_P p_{i2} (1 - p_{ij1}) - ((1 - e)c_C + e c_P p_{i2}) p_{ij1} - c_R (1 - (1 - e)p_{i1}) p_{ij2}, \quad (\text{A.1})$$

where c_C is the cost of cooperation, c_P is the cost of being punished, c_R is the cost of punishing, e is the chance a cooperator makes a mistake and does not cooperate instead, and z is a normalizer to make sure that payoffs are nonnegative.

2. **Group Conflict.** Each group pairs with another group at random and each pair has a conflict with probability ε . Group x wins a conflict with group y with probability

$$\frac{1 + p_{x1} - p_{y1}}{2} . \quad (\text{A.2})$$

The winning group obtains the resources of the losing group and replaces it completely. This is accomplished by allowing the winning group to produce a copy of itself, and the $2N$ individuals from the winning group and its copy are each randomly assigned to two new groups that replace groups x and y .

3. **Social learning.** At a third locus (genetic or cultural), individuals are one of two types: conformists or imitators. Conformists copy the modal (most common) types at the cooperate and punishment loci. With probability $1 - \mu$, imitators pick a random individual in their group and copy the types at the cooperate and punishment loci of that individual with a probability equal to their payoff relative to the total payoff of the pair; i.e., they are more likely to copy the “successful” types. With probability μ , imitators “innovate” and pick random types at the cooperate and punishment loci.

4. **Reproduction.** The population follows an overlapping generations model of reproduction (i.e., Cannings type model; see refs. Ewens, 2004; Etheridge, 2011) where each individual at each time step has probability q of dying. If a focal individual dies, two individuals are chosen from the group (with replacement and including the focal individual) with probability equal to their relative payoff in the group. These two individuals produce a single offspring sexually with recombination rates r_{kl} ($r_{kl} = 0$ in these simulations in order to match (Guzmán et al., 2007)) using a standard model for sexual production (see equation 34 in Hastings, 1986). Each locus in the offspring has a probability ν of mutating to a random allele. This offspring replaces the focal individual.

5. **Immigration.** Each time step each individual can be replaced by an immigrant from another group. With probability m , the focal individual meets an immigrant from another group and that immigrant replaces the focal with a copy of itself with a probability equal to the payoff of the immigrant relative to the combined payoff of the focal and the immigrant.

The metapopulation begins evolving with a state sampled from the approximate MaxEnt distribution described in Appendix B given fixed values of mean allele frequency, F_{ST} , and D' at the three loci. In the simulations presented in the main text, there are $n = 128$ groups and the initial frequency of each allele is $1/n = 1/128$; the values of N , F_{ST} , and D' are varied. Each simulation is run for 10^5 time steps and each set

of parameter values (population size N and initial F_{ST} and D') is replicated 250 times. Parameters values for
the life cycle are given in Table 1.

Appendix B. Distribution of initial population states

Here, I develop a heuristic implementation of a broader approach to sensitivity analysis in complex metapopulation models as function of initial statistical associations among state variables in the model. Most metapopulation models are a Markov process on the state space spanned by all possible genotypes in the metapopulation (a larger state space is needed for multiple demographic or environmental states or for frequency dependent selection). If there is even a small amount of bidirectional mutation (or other random process for generating variation), then the Markov process will be ergodic and will have a unique attracting stationary distribution of genotypes that the metapopulation will eventually approach regardless of its initial state (Karlin and Taylor, 1975). This stationary distribution is often the simplest and most complete description of the long-term outcome of the model. For example, a preponderance of the probability mass on one or a few genotypes would suggest that these genotypes are most likely to evolve in the long run.

Obtaining the stationary distribution of a Markov process is straightforward when it is easy to obtain the transition matrix of the process, which describes the probability of transitions from each state to every other state. In this case, the stationary distribution is the dominant left eigenvector of the transition matrix (Karlin and Taylor, 1975). However, for complex Markov models, such as the metapopulation model that I study in the main manuscript, the state space is large enough that it may be difficult to construct the transition matrix, much less calculate the stationary distribution from it. In these circumstances, the only feasible computational route may be to simulate the metapopulation for a large number of time steps. When there is a single state in the state space that is strongly attracting, as might be the case in a simple population with a single local fitness maximum on the fitness landscape (a “potential” landscape in more general models; Graham, 1980; Ao, 2004), the metapopulation will evolve towards this state and any excursions away from it due to random events (mutations, etc) are expected to be short. This means that once the metapopulation has appeared to reach an equilibrium in a simulation run, one can be reasonably sure that equilibrium is the state with the most probability density in the stationary distribution. However, if there are two local fitness maxima, the metapopulation may spend a considerable amount of time at each local fitness maximum, even if it spends a disproportionate amount of time at one maximum relative to the other under the stationary distribution (Gardiner, 2009). In such cases, especially when the population size is large, the transitions between these local fitness maxima occur extremely infrequently (Gardiner, 2009). Computationally, this means that it may be very difficult to tell if a candidate equilibrium observed in a simulation is really the most likely state

under the stationary distribution or if its some other less likely local fitness maximum. On average then, the metapopulation will be attracted to these local maxima when it is close to them in the state space, and it will eventually leave for other higher maxima on a timescale that is likely to be larger than the timescale of most simulation runs. Importantly, this timescale also may be larger than the biological timescale over which the actual phenotype evolved.

In order to obtain some sense of the relative stability of these locally attracting states, I need a way of sampling initial states for the Markov process. Since I don't know anything about the basins of attraction for potential attracting states, the distribution of initial states should be agnostic to the structure of the Markov process. However, I may know which statistics are important to calculate in order to describe the state of the metapopulation. Thus, I may wish to observe the evolution of the Markov process given constant initial values of these statistics. E.T. Jaynes (Jaynes, 1957, 1982; Jaynes and Bretthorst, 2003) and many others have advocated the use of a maximum entropy (MaxEnt) distribution in these cases since MaxEnt yields a probability distribution that conforms to the set of constraints but does not embody any other assumptions about the process. If $\mathbf{x}$ is some vector of state variables and there are m constraints of the form $E[f_k(\mathbf{x})] = \eta_k$, then the MaxEnt probability density is given by

$$q(\mathbf{x}) = \frac{\exp(\sum_k \lambda_k f_k(\mathbf{x}))}{Z(\lambda_1, \dots, \lambda_m)} \quad (\text{B.1})$$

where Z normalizes the probability density and λ_i are Lagrange multipliers that solve

$$\frac{\partial \log(Z)}{\partial \lambda_k} = -\eta_k \quad (\text{B.2})$$

(see Harte, 2011, p. 125). My goal will be to heuristically approximate the MaxEnt distribution of initial states in the metapopulation given a set of constraints on the genotype frequencies in the metapopulation.

Ideally, I would calculate the MaxEnt distribution of initial genotypes in the metapopulation given fixed values of mean allele frequency p and F_{ST} for three loci and the three two-locus linkage disequilibria D' . The full MaxEnt distribution is too difficult at present to obtain as it would require integrating F_{ST} and D' over all possible population states, but I will be able to derive a useful approximation. The main idea behind my approximation is to fix the mean allele frequencies across the metapopulation and then build a MaxEnt distribution with the three F_{ST} constraints that yields sample allele frequencies in each group. I can do this separately for each locus since, given no other constraints, F_{ST} at each locus depends only on the alleles at that locus, which means the MaxEnt distribution in (B.1) factors into separate MaxEnt distributions for each F_{ST} constraint.

Integrating F_{ST} over all n group allele frequencies is difficult, so I instead assume for each locus k that there are two types of groups, one with frequency p_{Ak} and the other with frequency p_{Bk} , with mean metapopulation allele frequency p_k . For biallelic loci, Wright's F_{ST} can be generally expressed as

$$F_{STk} = \frac{\frac{1}{n} \sum_i p_{ik}^2 - p_k^2}{p_k(1 - p_k)} \quad (B.3)$$

for locus k , which simplifies to

$$F_{STk} = \frac{(p_{Ak} - p_k)(p_{Bk} - p_k)}{p_k(1 - p_k)}. \quad (B.4)$$

The MaxEnt distribution is obtained by solving equation (B.2) for λ_k where η_k is the fixed value of F_{STk} and Z is calculated by integrating over $0 < p_{Ak} < p_k < p_{Bk} < 1$. After integrating Z , I find that

$$-\frac{\partial \log(Z)}{\partial \lambda_k} = \frac{\text{Shi}(\lambda_k) - \text{Chi}(\lambda_k) + e^{-\lambda_k} + \log(\lambda_k) + \gamma - 1}{\lambda_k (\text{Shi}(\lambda_k) - \text{Chi}(\lambda_k) + \log(\lambda_k) + \gamma)}, \quad (B.5)$$

where Shi and Chi are hyperbolic sine and cosine integrals, respectively, and γ is Euler's constant. Once allele frequencies have been sampled for each group from these MaxEnt distributions with constrained F_{ST} for each locus, I can then sample genotypes frequencies in each group.

Given the MaxEnt distribution for group allele frequencies for each locus, the resulting joint distribution (the product of the individual distributions) makes no specific assumption about metapopulation LD, which would otherwise result in relatively low covariance between different loci when each locus is sampled from a different group. This covariance is a kind of interdemic LD or two-locus F_{ST} . Finding the joint MaxEnt distribution of this LD alongside the one locus F_{ST} values is also computationally difficult (specifically, calculating the normalizer Z), so I instead assume that this interdemic LD is *high* by grouping the high allele frequencies for each locus in the same groups. Then, I calculate D values for each group from fixed D' values. Given fixed values of p_k and D_{kl} for each group, the MaxEnt distribution of chromosomes is a multinomial distribution with parameters N and genotypes frequencies given in equation (5) of Robinson et al. (1991). To obtain D from D' , I have to take in account the three-locus constraints given in equation (10) of Robinson et al. (1991). I do this heuristically by first calculating D from Lewontin's definition of D' (see equation 9 in Robinson et al., 1991) for each of the three pairs of loci. If these three D values do not lie within the convex hull generated by the three-locus constraints, I find the closest set of D values on the surface of the convex hull. Finding the shortest (Euclidean) distance between a point and a polygon is a equivalent to solving a quadratic optimization problem (Boyd and Vandenberghe, 2004), which I solve using the CVXOPT package (Andersen et al., 2012a,b).

I test the above heuristic approximation of the MaxEnt distribution by plotting F_{ST} values estimated from samples from the approximate distribution. In Supplementary Fig. S3, I show box plots of the Weir and Cockerham F_{ST} estimator (Weir and Cockerham, 1984; Weir and Hill, 2002) applied to the initial values used in the simulations in the main text. For all population sizes, the median estimated F_{ST} value tracks closely the initial fixed value given on the x -axis. As population size increases, the interquartile distance becomes narrower around the median. There is a slight bump at $F_{ST} = 0.4$ in all the panels; this represents a discontinuity induced by the simplification that each population can only take one of two initial allele frequencies, p_{Ak} or p_{Bk} .

The performance of the MaxEnt heuristic is good for other estimates of F_{ST} . For example, in Supplementary Fig. S4, I calculate F_{ST} by simply evaluating

$$F_{STk} = \mathbb{E}(H_{0k} - H_{1k}) / H_{0k} \quad (\text{B.6})$$

using sampling without replacement, which should correct the bias in small samples exhibited when (B.6) is calculated with replacement. Without replacement, I find that

$$F_{STk} = \frac{Q_{wk} - Q_{bk}}{1 - Q_{bk}} \quad (\text{B.7})$$

where

$$Q_{wk} = \frac{1}{n} \sum_i \frac{N(p_{ik}^2 + (1 - p_{ik})^2) - 1}{N - 1} \quad (\text{B.8})$$

and

$$Q_{bk} = \frac{n(p_k^2 + (1 - p_k)^2) - \frac{(N-1)Q_{wk}+1}{N}}{n - 1}. \quad (\text{B.9})$$

Supplementary Figs. S3 and S4 are almost indistinguishable, which indicates that the approximate MaxEnt distribution works well with both the Weir and Cockerham and the sampling without replacement estimators. Also, the sampling without replacement estimator appears to perform similarly to the Weir and Cockerham estimator, which is known to be unbiased (Weir and Hill, 2002).

I also test the sample D' values from the heuristic MaxEnt distribution. In Supplementary Fig. S5, I plot D' (using equation (14) from Robinson et al., 1991) from 250 samples of the heuristic distribution with $n = 128$ groups and $p_k = 0.5$ and $F_{STk} = 0.1$ for all loci k ; group size N is given in each panel. For even moderate group size, $N = 16$ and higher, the samples of the heuristic distribution recover the parametric value of D' in the distribution. However, due to sampling effects in small groups, it is hard to recover very high or low D' in the $N = 4$ case. This is exacerbated by very low or high allele frequencies in each group, which occur when

F_{ST} is high or mean allele frequencies are low or high. The initial values used in the simulations in the main text use an initial allele frequency of $p_k = 1/128$, which flattens the sampled values of D' as can be seen in Supplementary Fig. S6 for $F_{STk} = 0.1$ especially for low group size. High $F_{ST} = 0.9$ exacerbates this pattern further in Supplementary Fig. S7.

Appendix References

Andersen, M. S., Dahl, J., Liu, Z., Vandenberghe, L., 2012a. Interior-point methods for large-scale cone programming. In: Sra, S., Nowozin, S., Wright, S. J. (Eds.), Optimization for Machine Learning. Neural information processing series. MIT Press, Cambridge, MA, pp. 55–83.

Andersen, M. S., Dahl, J., Vandenberghe, L., 2012b. CVXOPT: A python package for convex optimization, version 1.1.5. <http://abel.ee.ucla.edu/cvxopt>.

Ao, P., 2004. Potential in stochastic differential equations: novel construction. J. Phys. A 37 (3), L25.

Boyd, R., Gintis, H., Bowles, S., Richerson, P. J., 2003. The evolution of altruistic punishment. Proc. Natl. Acad. Sci. U. S. A. 100 (6), 3531–3535.

Boyd, S. P., Vandenberghe, L., 2004. Convex Optimization. Cambridge University Press, Cambridge, UK.

Etheridge, A., 2011. Some Mathematical Models from Population Genetics. Vol. 2012 of Lecture notes in mathematics. Springer, Berlin Heidelberg.

Ewens, W. J., 2004. Mathematical Population Genetics. Springer, New York.

Gardiner, C. W., 2009. Stochastic Methods: A Handbook for the Natural and Social Sciences, 4th Edition. Springer series in synergetics. Springer, Berlin.

Graham, R., 1980. Solution of Fokker Planck equations with and without manifest detailed balance. Z. Phys. B 40 (1), 149–155.

Guzmán, R. A., Rodríguez-Sickert, C., Rowthorn, R., 2007. When in Rome, do as the Romans do: the coevolution of altruistic punishment, conformist learning, and cooperation. Evol. Hum. Behav. 28 (2), 112–117.

Harte, J., 2011. Maximum Entropy and Ecology: A Theory of Abundance, Distribution, and Energetics. Oxford series in ecology and evolution. Oxford University Press, Oxford.

Hastings, A., 1986. Multilocus population genetics with weak epistasis. II. Equilibrium properties of multilocus models: what is the unit of selection? Genetics 112 (1), 157–71.

Jaynes, E., 1982. On the rationale of maximum-entropy methods. *Proc. IEEE* 70 (9), 939 – 952.

710 Jaynes, E. T., 1957. Information theory and statistical mechanics. *Phys. Rev.* 106 (4), 620–630.

Jaynes, E. T., Bretthorst, G. L., 2003. *Probability Theory: The Logic of Science*. Cambridge University Press,
712 Cambridge, UK.

Karlin, S., Taylor, H. M., 1975. *A First Course in Stochastic Processes*, 2nd Edition. Academic Press, New
714 York.

Lewontin, R. C., 1964. The interaction of selection and linkage. I. General considerations; heterotic models.
716 *Genetics* 49 (1), 49–67.

Robinson, W. P., Asmussen, M. A., Thomson, G., 1991. Three-locus systems impose additional constraints on
718 pairwise disequilibria. *Genetics* 129 (3), 925–930.

Weir, B. S., Cockerham, C. C., 1984. Estimating F-statistics for the analysis of population structure. *Evolution*
720 38 (6), 1358–1370.

Weir, B. S., Hill, W. G., 2002. Estimating F-statistics. *Annu. Rev. Genet.* 36, 721–750.

722 Wright, S., 1951. The genetical structure of populations. *Ann. Eugen.* 15 (1), 323–354.