Unit 5: Coordination and norms

READINGS AND RESOURCES

- Boyd R, Richerson PJ (2002) Group beneficial norms can spread rapidly in a structured population. *Journal of Theoretical Biology* 215: 287–296 .
- Ehrlich PR, Levin SA (2005) The evolution of norms. *PLOS Biology* 3(6): e194.

The problem of coordination and norms

Cooperation is not the only problem to be solved for social cohesion. Once we have decided to cooperate, we have to figure out how to do so. How do we best cooperate? Who are the best cooperators for *me*? If we don't understand one another's intentions, don't have the same goals, or don't do things the same way, cooperation can be impeded. Doing things the same way is what we will call **coordination**.

Language is a prominent example of coordination—many of the words we use for things are arbitrary, but unless we use the same words for a given situation, we are liable to have a misunderstanding (and of course this can occur regardless). In most of the world, cars and trucks are driven on the right side of the road, though in the UK, Australia, and Japan, cars are driven on the left. It doesn't really matter which side you drive on, only that everyone does things the same way.

The word "norm" is used a few different ways in the social sciences, but I'll use it here to connote behaviors that facilitate coordination. Norms differ from one population to another, and can also appear fairly arbitrary.

Sometimes norms can serve to signal to others a larger suit of norms that are harder to observe. If you go to academic conferences and are male, you are likely to wear a tie if you are an economist, but extremely unlikely to do so if you are an anthropologist. The tie, or lack thereof, doesn't serve any function directly, but it signals to others the sort of person you are likely to be. Likewise, norms often develop among communities or groups of friends that enable rapid communication between their members, but are baffling to outsiders.

Although norms may be arbitrary, that doesn't mean that all norms are always just as good as any other norms. Some, like smoking cigarettes together after vigorous activity, can be directly harmful. Others norms can be prosocial, and improve the collective good. For example:

- Rules against murder and assault encourage civil order, and make us feel relatively secure.

- Property rights allow encourage productive effort and innovation.
- Well-managed taxation provides roads, schools, and other public goods.
- Product standards, building codes, and rules of professional conduct allow more efficient commerce and protect citizens from harm.
- Norms governing the filling of political offices reduce the chances of a civil war over political disputes (hopefully)

In any given domain, there are many possibilities for exactly what norms to put in place. Drive on the right or the left? Speak English or German? Wear a business suit or yoga pants and flip flops? Allow disputes to escalate, or prohibit violence? Some norms may be more in line with basic human psychology and practicality than others (for example, we can probably take driving in swirly spirals off the table), but in other cases it's clear that many options are possible. Human cultural diversity is a testament to this fact.

So, how might norms spread or dominate? In this unit, we'll tackle this question with a series of models. In a sense, this question combines our first two units on the spread of behaviors or opinions with the previous unit, in which the utility of a behavior depended on the behavior of others. We'll again use an evolutionary framework with success-biased transmission. That is, we'll assume that individuals interact with others using particular strategies, and that those interactions lead to some payoffs. After, individuals will have the opportunity to observe others and switch to a strategy that is more successful. Now, we could also use other social learning strategies, such as a conformist strategy in which agents copy the most common norm. This sort of thing is left as an exercise.

To start, let's consider the example in which two norms compete and neither has any intrinsic advantage.

Norms with symmetric payoffs

CODE: coordination_simple.nlogo

When considering competition among strategies, it's simplest to start with only two. Let's imagine a population in which everyone uses one of two norms in their social interactions. In this simple coordination game, each agent uses either norm 1 or norm 2 in their interactions. This is not a cooperative dilemma; it is always better to coordinate than not, and even non-coordination is cooperative and therefore yields a positive payoff. In fact, we can assume that in this scenario, the dilemma of cooperation has been solved, and now the issue is how to maximize the benefit between potential cooperators. Here we consider a symmetric coordination game, in which coordination on norms always yields a benefit δ over non-coordination. Without any loss of generality, we can set the baseline payoff to 1. This yields the payoff matrix below (with the payoffs written for the row player):

	Norm 1	Norm 2
Norm 1	$1+\delta$	1
Norm 2	1	$1+\delta$

Individuals in the model have many interactions with members of their community, during which they accumulate payoffs. In the prisoner's dilemma simulations of the last unit, we allowed individuals to play with a small number of individuals to accumulate payoffs, and then always copy the strategy of their best-performing neighbor. We could do that here as well. However, it's also important to note that this is a somewhat arbitrary modeling assumption about how interactions and observations are structured. For the sake of exploring, we'll try a different set of mechanisms here.

Let's assume that each individual has a lot of interactions in which there is an opportunity to coordinate. In this case, we don't necessarily need to simulate all of these interactions. Instead, if we know the distribution of strategies in the population, we can simply calculate the expected payoff under the assumption of random interactions. Recall that agent-based models are great for exploring assumptions of non-random interactions, but it's also important to establish a baseline.

To calculate an agent's payoff, we just need to know the proportion of other agents who do and do not share the agent's norm. Let n_1 be the number of agents who use norm 1, and n_2 be the number of agents who use norm 2, in a population of N (so $n_2 = N - n1$). The payoff to an agent who uses norm 1 is determined by considering the expected payoff over time if interaction partners at random from the population:

$$V_1 = \frac{n_1 - 1}{N-1}(1+\delta) + \frac{n_2}{N-1}(1)$$

Which reduces to:

$$V_1 = 1 + \frac{(n_1 - 1)}{N-1}\delta$$

and similar for V_2. This makes sense and is intuitive if you examine it a bit. It says that the payoff is the baseline of 1 plus the coordination bonus times the proportion of the time the agent's interaction partner uses the same norm as the agent.

Now, it comes time for imitation. Each agent will select one other agent at random, and compare their strategies and payoffs. If we stuck with the assumption we used in the previous unit, of deterministically copying the best, then whichever norm was more popular would be instantly adopted by everyone. This is a valuable thing to realize. In the PD game, focusing only on local interactions influenced and slowed the spread of behaviors. Here, we use the expected payoff, so we ignore the heterogeneity added by local interactions. Instead, we will add heterogeneity another way. Instead of assuming deterministic copying, we will make what is probably a more realistic assumption: that *copying is probabilistic*.

Here's why I think probabilistic copying makes sense. If your payoff is much higher than mine, I should be very likely to copy you. If your payoff is only a little higher than mine, I should be more likely to copy you than not, but also more likely to stick with my current strategy. And if your payoff is worse than mine, I should generally stick with my current strategy, but every now and then I might make an error in judgment or want to explore, so I will still occasionally copy a worse strategy. Luckily, there is a functional form that fits this description nicely: the sigmoid function.

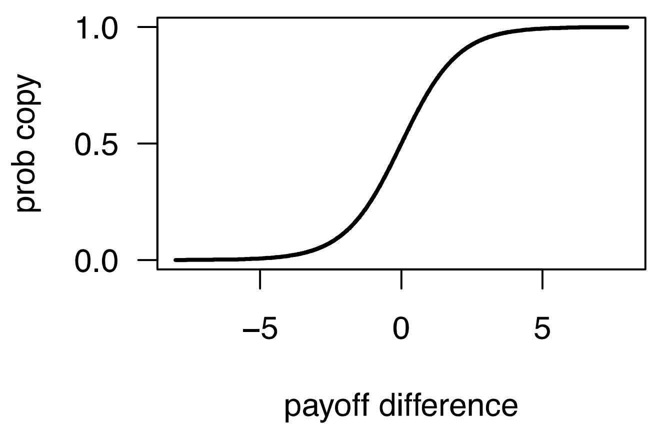

Assume agent *i* observes agent *j*. If agent *i* has payoff x_i and agent *j* has payoff x_j, then the probability of *i* adopting *j*'s strategy is:

$$P(copy) = \frac{1}{1+\exp[-\alpha(x_j - x_i)]}$$

where α is a scaling parameter that controls how quickly small differences become important. For simplicity, we'll use $\alpha = 1$. OK. Now let's go to the code. As we're done before, we'll use colors do designate the differences between agents. Norm 1 agents will be yellow, norm 2 agents will be blue.

SETTING UP THE MODEL

- *init-norm1* slider
- *coordination-benefit* slider
- turtles-own [*norm1?, payoff*]
    - *norm1?* is a Boolean variable that will be true if the agent uses norm 1, false if they use norm 2. We could instead use integers rather than a Boolean switch to allow for more than two strategies in future versions.

INITIALIZATION

- Each patch of the grid will sprout a turtle. With probability *init-norm1*, the agent uses norm 1, otherwise norm 2. Norm 1 agents are yellow, norm 2 agents are blue.

DYNAMICS

- Stop if one norm completely dominates.
- Calculate payoffs for norm 1 and norm 2 agents, and assign those payoffs accordingly.
- ASK TURTLES:
    - Choose another agent to observe at random.
    - Copy the other agent's norm with a probability derived from a sigmoid function based on the difference between payoffs.

PLOTTING

- The agents' colors will represent their norms, and we can see the change happen on that level. We will also plot the frequency of agents using norm 1 in the population over time.

RESULTS

- If there is no benefit to coordination ($\delta=0$), both norms can persist indefinitely, as the frequency of each norm over time is essentially a random walk. In the long run, one norm will probably go to fixation just by chance – this is called *neutral drift* – but it can take a very long time.
- As soon as there is any benefit to coordination ($\delta>0$) the more common norm will almost always spread. There is some uncertainty when they start out at similar numbers due to stochasticity, but otherwise the more common one is favored.
- So when there is a benefit to coordinating, the more popular norm should spread. If uncommon norms persist, there must be mechanisms beyond coordination, or the assumption of random interactions may not be met. For our purposes, we have established a baseline model that illustrates dynamics of norms when those norms are completely arbitrary and neither is intrinsically better. Now, let's relax that assumption, and consider the case where one norm is clearly superior to another. Will it spread?

Group-beneficial norms

CODE: coordination_asymmetric.nlogo

Things get more interesting if we consider two competing norms in which one norm is prosocial – that is, better for the group – but carries a cost when rare, since others will fail to coordinate. Consider the following payoff matrix for the asymmetric coordination game.

	Norm 1	Norm 2
Norm 1	$1+\delta+g$	$1-h$
Norm 2	$1+g$	1

Norm 1 here is the prosocial norm. Those employing norm 1 confer a benefit *g* on everyone, regardless of their norm. There is also an additional benefit δ of coordinating on norm 1. Coordinating on norm 1 is obviously preferable to coordinating on norm 2, and it is clear that if norm 1 is common, anyone using norm 2 should switch (to receive the additional benefit δ). However, to make things more interesting, norm 1 also carries a cost when rare. Norm 1 may be costly to implement, it may be burdensome to execute alone, or it may be actively punished by non-users. All these are real possibilities for uncommon behaviors. We'll model this by imposing a cost *h* to employ norm 1 when one's partner uses norm 2.

It's easy to see that the prosocial norm should persist when very common. In the language of evolutionary game theory, norm 1 should resist invasion by rare norm 2 agents. Our question, then, will be how costly prosocial norms can spread when rare? To tackle this, let's update our model.

SETTING UP THE MODEL

- *norm1-group-benefit* slider (*g*)
- *norm1-self-benefit* slider (δ)
- *norm2-deviance-cost* slider (*h*)

INITIALIZATION

- Each patch of the grid will sprout a turtle. With prob *init-norm1*, the agent uses norm 1, otherwise norm 2. Norm 1 agents are yellow, norm 2 agents are blue.

DYNAMICS

- Stop if one norm completely dominates.
- Calculate payoffs for norm 1 and norm 2 agents, and assign those payoffs accordingly. **Use updated payoff matrix**.
- ASK TURTLES:

- Choose another agent to observe at random.
- Copy the other agent's norm with a probability derived from a sigmoid function based on the difference between payoffs.

RESULTS

- Let's start out with some arbitrary values. Let $\delta = g = 1$, and *h* = 0.5. If we let the initial frequency of norm 1 be *init-norm1* = .5, we see that the prosocial norm spreads every time. Even if we lower its initial frequency to 0.4, it spreads every time! Huzzah! The prosocial norm spreads even when rare! …Or does it? Keep lowering its initial frequency. You'll find that right around an initial frequency of 0.33, norm 2 starts to occasionally dominate. If init-norm1 ≤ 0.3, the prosocial norm almost never spreads. If you play around with the payoffs, you'll find that this threshold moves around, but there is usually some threshold below which norm 2 spreads. Let's take a look at why.
- In general, I've avoided mathematical derivations for this course and stuck to insights that can be gleamed from simulations. However, in this case it is pretty easy to calculate just where the threshold is, so we'll do it. This will be an approximation, for which we will assume a large enough population where the difference between *N* and *N* – 1 is negligible. Let the proportion of agents using norm 1 be *p*, so the proportion using norm 2 is 1 – *p*. The expected payoff to an agent using norm 1 is the proportion of times it receives a coordination benefit plus the proportion of times it pays a non-coordination cost, $V_1 = p(1+\delta+g)+(1-p)(1-h)$. The payoff to a norm 2 agent is similarly based on our payoff matrix: $V_2 = p(1+g)+(1-p)(1)$. The threshold separating when each norm will spread will be the point at which these two payoffs are equal, so that below the threshold norm 2 will be dominant and above it norm 1 will be. Stretch out your algebra muscles and try it! Setting the two payoffs equal and solving for p yields a threshold of $p^{¿} = h/(\delta+h)$.
- We can add a monitor that automatically calculates and displays the threshold based on our model parameters. For the parameters we used in our initial exploration, we see that the threshold is *p** = 0.3333, exactly what we observed.
- I ran 100 simulations under these parameter values for different initial frequencies of norm 1, plotted the proportion of them in which norm 1 dominated. You can see quite clearly the dividing line.

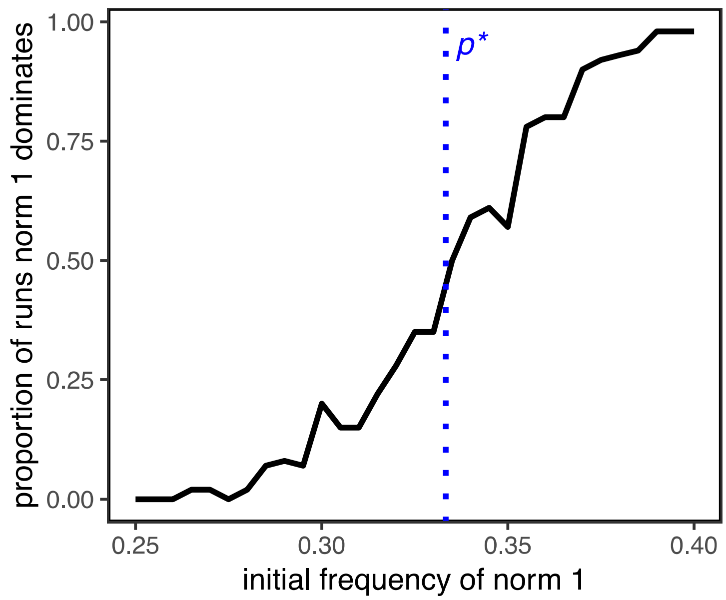

- A few things are illuminated by the calculation of *p**. First, the threshold for norm 1 to spread is lower when the benefit of coordinating on norm 1 is stronger. Second, that threshold is higher when the cost of using norm 1 with a norm 2 agent is greater. These results make sense. Even rare interactions may be attractive if they have a high payoff, but these benefits must outweigh the costs of non-coordinating interactions.
- The third thing we learn is that the prosocial benefit generated by norm 1, *g,* is completely irrelevant to its spread. It can be very large, or it can be zero, and it will no influence on the spread of norms in this model. The reason is apparent when we compare the equations for the payoffs. Norm 1 agents provide the benefit *g* to everyone in the population, so *g* plays no role in differentiating the payoffs of the two norms. This indicates that to some extent we have modeled a norm that is better during coordination but costly when rare. Its prosocial nature is an afterthought. Two comments on this point. First, we are often interested in how norms that are better for the population but costly when rare can spread, even if they are not explicitly altruistic. Second, the model *does* teach us something about the spread of prosocial norms, which is that their spread probably depends primarily on the costs and benefits of coordination and non-coordination rather than the group-level benefits they confer.

A major thing we've learned from this model is that if a prosocial (or better) norm is not sufficiently common at the outset, it will not spread, but instead will disappear. This is a little depressing. What if we introduce a superior norm that, if adopted, will make us all better off? The model suggests that, unless a sufficiently large faction can be coerced to adopt the new norm to demonstrate its clear superiority, it will fizzle and die. Is there any other hope for the spread of new beneficial norms?

Yes. In fact, there are several possible mechanisms for the spread of rare norms, but we will just tackle one. We've assumed so far, in all our models actually, that we are dealing with a single population of interacting agents. However, humans are structured into groups, who often interact primarily with other group members even if they occasionally interact with non-group members. Might group structure play a role in the spread of norms?

Group-beneficial norms in a structured population

CODE: coordination_2groups.nlogo

Let's consider not one but two groups. Within each group, things will work as before. Individuals will interact with group members, accumulate payoffs, and then observe another agent for success-biased social learning. If the groups never interact in any way, we will simply have two exact copies of our previous model. However, individuals do sometimes observe those from other groups. This can happen through a variety of mechanisms, including travel, trade, and intermarriage. If a member of another group uses a different norm and is doing demonstrably better than you, might you consider adopting that norm? This is the situation we consider here. The mechanism of success-biased copying in a group-structured selection provides the basis for a type of **cultural group selection**, in which group structure facilitates the spread of group-beneficial norms (other mechanisms include migration and military conquest; see readings in further directions).

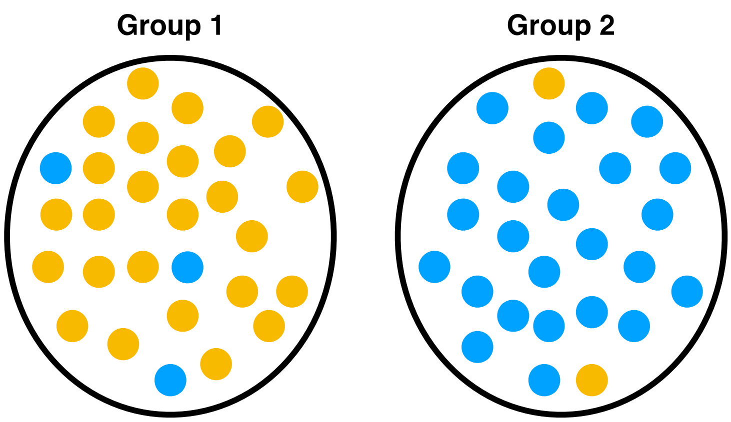

Given that norms may arise in a variety of ways, it is not unreasonable to assume that different groups will initially converge on different norms for fairly arbitrary reasons. We have seen that, once a dominant norm is established, it can be difficult for a new norm, even a superior one, to spread within a population. However, there may also be variation between groups, and if that variation is observable, a norm may be able to cross group boundaries.

To set up this model, we will extend our grid and divide it into two, so that group on the left will be separate from the group on the right, and we will assign each agent a group ID in addition to a norm. In principle, the model can be extended for an arbitrary number of groups.

SETTING UP THE MODEL

- *init-norm1-group1* slider
- *init-norm1-group2* slider
- *prob-outgroup-observation* slider
- turtles-own [*groupID*] (0 or 1)

INITIALIZATION

- Double the width of the grid, and set up patches into two territories.
- Each patch of the grid will sprout a turtle. With prob init-norm1, the agent uses norm 1, otherwise norm 2. Assign groupID based on which territory the agent is in. Norm 1 agents are yellow, norm 2 agents are blue.

DYNAMICS

- Stop if one norm completely dominates.
- Calculate payoffs for norm 1 and norm 2 agents, and assign those payoffs accordingly. Use updated payoff matrix.
- ASK TURTLES:
    - Choose another agent to observe at random. With probability prob-outgroup-observation, observe an agent from the out-group. Otherwise observe an agent from the in-group.
    - Copy the other agent’s norm with a probability derived from a sigmoid function based on the difference between payoffs.

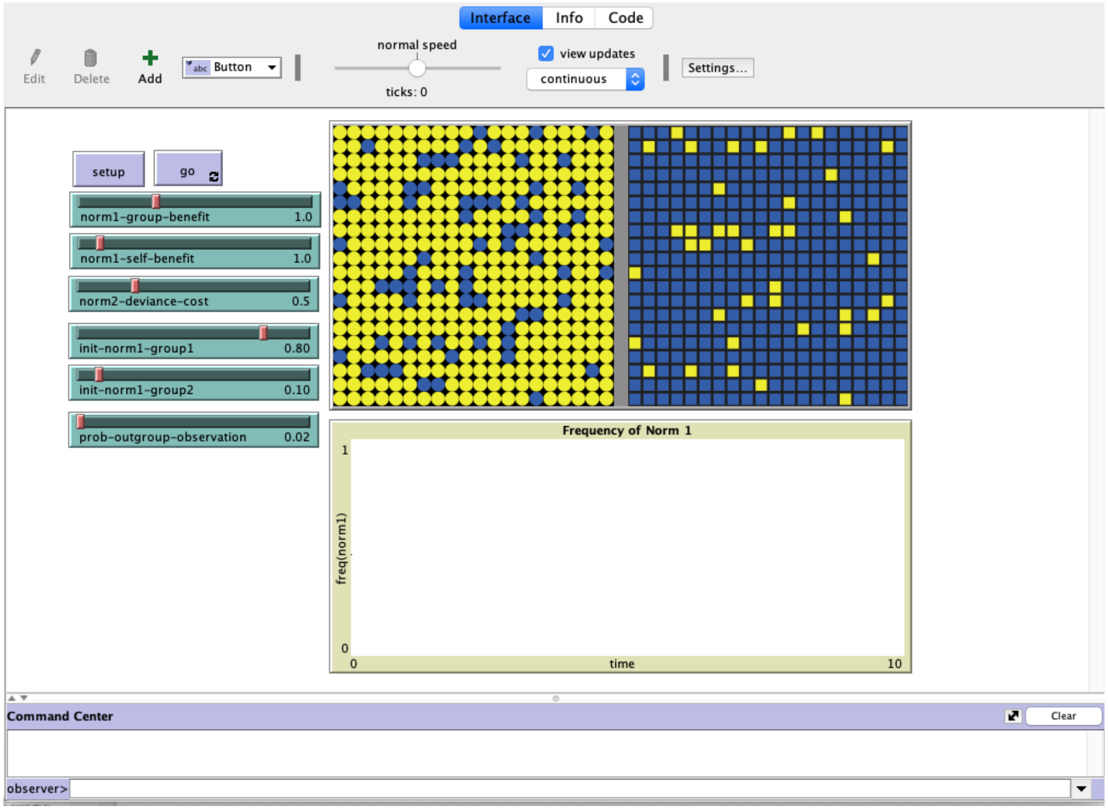

RESULTS

- Consider our parameters from our previous analysis. Let $\delta = g = 1$, and *h* = 0.5. Set the initial frequency of norm 1 to 60% in group 0 and 10% in group 1. If *prob-outgroup-observation* is 0, then what happens is predictable. We know that the threshold for norm 1 to spread is p* = .33. So norm 1 will take over in group 0 and die out in group 1.
- Now, increase prob-outgroup-observation to .01. We see that individuals in group 1 do occasionally adopt it, but because most observations are still within-group, and norm 1 agents do poorly in group 1, the norm doesn't spread. If we increase *prob-outgroup-observation* to .02, however, we see larger oscillations, and then… bam! Norm 1 spreads through. So even a small amount of outgroup interactions can facilitate the rapid spread of group-beneficial norms from one group to another. Unsurprisingly, it can also be shown that the higher the value of p*, the more outgroup observation is required for the norm to spread from a group where it is common to one where it is rare.

If individuals sometimes observe others from other groups and are willing to switch to any more successful strategy, then the group beneficial norm can spread with even a small amount of outgroup observation. Note here that all *interactions* still occur within one's group. The assumption here is that individuals are willing to consider another's success as a reason to copy them, even if they are successful in a different social context. Based on the psychological evidence from humans, this seems like a reasonable assumption.

Another strong assumption of this model is that both norms and their associated payoffs are easily observable. This may not always be the case. As many who have entered a new cultural environment know, the right way to behave isn't always obvious, and it's certainly not necessarily clear precisely which norms lead to a group's success when many candidates are present. For a striking example, consider the cargo cults of the Pacific in the aftermath of World War 2, who constructed wooden air traffic control towers and performed pseudo-military drills, emulating US soldiers in hopes of attaining similar levels of material wealth.

Further directions

- Cultural group selection. Although selection at the level of the group is controversial when considering the evolution of genetic traits, group selection is well established for cultural traits due to the normative nature of human psychology and culture. The model considered in this unit is just one approach to understanding how group structure and competition facilitates the evolution of cultural practices. For reviews and other approaches, see the following papers.
    - Henrich J (2004) Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization* 53: 3–35.
    - Boyd R, Richerson PJ (2010) Transmission coupling mechanisms: Cultural group selection. *Philosophical Transactions of the Royal Society B* 365: 3787–3795.
    - Richerson PJ et al. (2016) Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences* 39: e30.
- Signals and markers. Many norms of behavior are opaque until one is already engaged in interaction. If a person could express signals about their likely normative portfolio, those signals would enable more effective assortment on norms and more efficient coordination. An early treatment of this idea was provided by the economist Michael Spence (1973) in the context of how employers evaluate potential employees. McElreath et al. (2003) later showed that if individuals preferred to interact with those who shared arbitrary markers, markers can emerge to be associated with particular norms. Signaling also can facilitate the persistence of multiple norms in a population, with the norm-marker association strongest near the boundaries between populations who hold differing norms.

- Spence M (1973) Job market signaling. *The Quarterly Journal of Economics* 87: 355–374.
- McElreath R, Boyd R, Richerson PJ (2003) Shared norms and the evolution of ethnic markers. *Current Anthropology* 44: 122–130.

- Noisy signals and intragroup variation. Within populations, there may still be variation, but avoiding those who don't share all your norms isn't always feasible. Within a society, we sometimes have to interact with those who don't share our norms, and we may want to avoid signals that highlight our differences. My colleagues and I considered two types of signals – those sent overtly and covertly – and showed that covert signals should be dominant in cases where assortment is imperfect and uncertainty about shared norms between strangers is both high and consequential. Related work by Hoffman et al. has explored conditions for intentionally noisy signals in which the act of obscuring is itself ascribed value.
  - Smaldino PE, Flamson TJ, McElreath R (2018) The evolution of covert signaling. *Scientific Reports* 8: 4905.
  - Hoffman M, HIlbe C, Nowak MA (2018) The signal-burying game can explain why we obscure positive traits and good deeds. *Nature Human Behaviour* 2: 397–404.
- Language and conventions. Much of communication involves an issue of coordination – the need for linguistic signals to represent similar things to both sender and receiver. Although language is in reality much more than a set of signaling conventions — its most important feature perhaps being the flexibility to convey almost any idea—its conventional aspect is nevertheless important to understand. Models have shown how repeated coordination games can give rise not only to shared signals, but shared categories, as with the conventions for color names or safe vs. dangerous foods.
  - Cangelosi A, Parisi D (1998) The emergence of a 'language' in an evolving population of neural networks. *Connection Science* 10: 83–97.
  - Puglisi A, Baronchelli A, Loreto V (2008) Cultural route to the emergence of linguistic categories. *Proceedings of the National Academy of Sciences* 105: 7936–7940.
  - Contreras Kallens PA, Dale R, Smaldino PE (2018) Cultural evolution of categorization. *Cognitive Systems Research* 52: 765–774.

Exercises

- *When in Rome*. Adapt the first model of symmetric coordination so that, instead of success-biased social learning, an agent observes several other agents and adopts whichever norm is more common among those observed. Does this assumption of conformity change any of your conclusions about either which norm will spread or how quickly it will do so? Consider how quickly norms spread as a function of the number of agents being observed at a time.

- *Changing norms is hard*. Consider the emergence of a more group beneficial norm *b* in a population where most people use norm *a*. This second norm confers a lower benefit than would norm *b* if it were common, but all norms do poorly when they are rare. Perform simulations in which a proportion *p* of the population begin to employ norm *b*. How large does *p* have to be for norm *b* to spread? Do your results track our calculation of p^*?
- *Group, there it is.* Now consider our group-structured model, in which one group uses inferior norm *a* and the other group uses superior norm *b*. Let each individual observe a member of the other group with probability *m*. Show that sufficient out-group contact can facilitate the spread of group-beneficial norms, and that the amount of contact needed varies with the costs and benefits of each norm.