A Online Appendix

A.1 Further discussion of Section 2 on the Simple Binary Model (SBM)

In §2, §2.7, we discussed that overconfidence can improve social learning in the model of Bernardo and Welch (2001). We now provide algebraic details.

We consider a setting as in the SBM, except that I_2 is overconfident, in the sense that I_2 incorrectly assesses I_2's precision at $p' = p + \epsilon$, where $\epsilon \approx 0$. I_2's true precision of p is known to all except I_2. Observe that if the first two actions are HL or LH, then starting with I_3 the model is isomorphic to the base model with I_3 as the first agent. So the probability of a correct cascade has two terms:

$$p^2 + 2p(1-p)\frac{p(p+1)}{2(1-p+p^2)}.$$

The first term is the probability of a correct cascade occurring starting with I_3,

$$0.5\,\mathbb{P}[a_1a_2 = HH|\theta = H] + 0.5\,\mathbb{P}[a_1a_2 = LL|\theta = L] = p^2.$$

The second term is the probability that the first two actions are opposed, $2p(1-p)$, multiplied by the probability of a correct cascade in the long run in the SBM, as given by equation (1) above. The difference between this expression and the probability that in the long run a correct cascade occurs in the base model (again as given in (1)) is

$$\frac{p(2p-1)(1-p)}{2(1-p+p^2)} > 0,$$

since $p > 0.5$. So overconfidence increases average accuracy and welfare.

A.2 Further discussion of Section 6 on limited communication or observation of past actions

We discuss here imperfectly rational models that make assumptions about the agent's mapping from observed actions and payoffs into the agent's actions.

A.2.1 Heuristic action rules based on observation of averages of past payoffs and of action frequency

If rational agents could observe the average payoffs experienced from possible actions undertaken by many in a large population, they would always choose the right action. However, it is possible that agents fall short of rationality. In the model of Ellison and Fudenberg (1993), there is a continuum of agents, where each period some given fraction of them has the opportunity to update their choices. Agents observe the average payoff of the two actions in the population from the previous period, and their popularities. Agents use rules of thumb for updating their actions, which take the form of specified probability of switching action based on these two variables. The paper explores conditions on the weighting of these considerations that promote or hinder long-run correct decisions.

A.2.2 Heuristic action rules based on sampling of past actions and payoffs

In Ellison and Fudenberg (1995), there is a large number of agents who take simultaneous actions at each discrete period. Each period, some fraction of agents exogenously stick to their current action, and the remainder observes an unbiased sample of the latest actions and payoffs of N and choose action based on this and based on their own latest payoff. Agents follow the following heuristic decision rule. If all the agents in an agent's sample takes the same action, the agent follows that action. If both actions are selected by at least one agent in the sample, the agent chooses whichever action has higher average payoff based on observed reports and the agent's own latest experienced payoff.

When the sample size N is small, learning from the experience of others causes the system to evolve to universal adoption of the correct action. In contrast, when the sample size is large, there is strong pressure toward diversity of behavior.[56] So the system never fixes on the correct action. This analysis is valuable in illustrating how non-obvious interesting conclusions about efficiency derive from reasonably plausible heuristic assumptions. On the other hand, when there is a 'split decision' in an agent's sample, it would be reasonable for an agent to take into account that a preponderance of 99 Adopts to 1 Reject (for example) might suggest almost conclusively that adopt is superior.[57] This would tend to oppose the diversity effect discussed above.

A.2.3 Recent past signals observable

The SBM and the BHW model are based on the premise that agents do not directly communicate their private signals. If the signals of agents are fully communicated to their followers, learning would be efficient, and there would be asymptotic learning.

In many practical contexts, more information is indeed transmitted than just action choices. For example, people are often free to talk about their private information, though owing to time and cognitive constraints, such communication may be imperfect. This raises the question of whether incorrect cascades still occur, and whether there is asymptotic learning.

We consider here the possibility of limited communication of private signals. Every private signal is directly communicated, which in principle could bring about asymptotic learning. However, we consider a case in which each signal is passed on to just one other agent, and provide an example in which incorrect cascades still form and last forever.

Consider the SBM of Section 2, except that in addition to observing all past actions, each agent observes the private signal of the agent's immediate predecessor. In contrast with the SBM, we assume that when indifferent, each agent always follows her own signal. This assumption is convenient but not crucial.

After actions HH or LL, I_3 infers that the first two signals were hh or $\ell\ell$ respectively, and falls into a cascade which could easily be incorrect. This cascade can be broken. To see this, suppose that the first two actions are HH, and that I_3 observes ℓ. So I_3 chooses H. If I_4 observes ℓ, this combines with observation of I_3's ℓ to make I_4 indifferent, so I_4 rejects, breaking the cascade.

[56] This is because if one action is very unpopular, with large N many adherents of the more popular action "hear about" the unpopular one and potentially convert to it.

[57] Ellison and Fudenberg consider a related effect which they call 'popularity weighting.'

Consider next the case in which I_3 observes h as well. Now, I_4 adopts, because I_4 knows that the first 3 signals were all h. At this point the cascade of adoption continues forever. I_5 knows that the first two signals were hh, and that the third and fourth signals were not $\ell\ell$. (I.e., they were either ℓh, hh, or $h\ell$.) So the net evidence in favor of state H is at least slightly stronger than two h signals. So even in the worst case in which I_5 directly observes two ℓ signals (I_5's and I_4's signals), I_5 strictly prefers H. Similar reasoning shows that all subsequent agents adopt.

The general point that inefficient cascades can form and last forever does not require this arbitrary tie-breaking convention.[58] Intuitively, the ability to observe the predecessor's signal is a double-edged blade. The immediate effect is to increase an agent's information, which increases the probability that an agent decides correctly. But to the extent that this is the case, it later makes the action history more informative. That eventually encourages later agents to fall into a cascade, which blocks asymptotic learning.

In particular, the action sequence becomes informative enough that agents fall into cascades despite their access to an extra signal from the predecessor. The actions history eventually overwhelms an agent's information (even inclusive of observation of a predecessor's signal), at which point the accuracy of the social belief stops growing.

Another way to think of this is that being able to observe predecessors' signals is somewhat like being able to observe multiple private signals instead of one. This does not fundamentally change the argument for why incorrect cascades form—that the accumulation of information implicit in past actions must eventually overwhelm the signal(s) of a single agent.

[58] The following modification illustrates the point in a setting with no ties. Suppose that agents differ very slightly in their precisions, and that agents act in inverse order of their precisions, from low to high. Now agents are never indifferent. Agents who are close to indifference strictly prefer to follow their own signals, and the same analysis holds.