The Wisdom of the Minority

Steven Callander* Johannes Hörner†

January 2005

Abstract

We consider a simple version of the social learning model, in which agents are either informed or uninformed, and only observe a summary statistic of their predecessors' choices; namely, how many have chosen one or the other alternative. The fraction of informed agents is the only parameter in the model. We study the (uninformed) agents' optimal strategy, and show that, if the fraction of informed agents is small enough, it is optimal to follow the choice made by the minority, provided that this minority is not too large. This occurs independently of the number of agents who have chosen so far. However, such events are unlikely, as we show that the expected fraction of agents taking the correct action tends to one. Always following the choice of the majority is optimal if and only if the fraction of informed agents is at least 7/9.

1 The Model

The model is a simple variant of the canonical model of social learning introduced by BHW. There are two equiprobable states of the world, $\Omega = \{\omega_0, \omega_1\}$. There is a countable set of agents, $i = 1, 2, \ldots$. Each agent must take an action, either 0 or 1. The payoff of an agent is 1 if the state is ω_0 and he takes action 0, or if the state is ω_1 and he takes action 1, and 0 otherwise.

Agents move sequentially. Each agent is either *informed*, with probability p, or *uninformed*. These informational types are privately known and independently distributed. An uninformed agent assigns probability 1/2 to either state. An informed agent learns the realized state. Alternatively, agents can be thought of as homogeneous but receiving signals of different accuracy (either uninformative with probability $(1-p)$ or perfectly informative with probability p).

*Department of Managerial Economics and Decision Sciences, Kellogg School of Management scal@kellogg.northwestern.edu.

†Department of Managerial Economics and Decision Sciences, Kellogg School of Management and HEC School of Management, Paris. j-horner@kellogg.northwestern.edu.

Agent n observes how many agents, among the first $n-1$ agents, have taken action 1, but he does not observe the order of these decisions (Agent 1 observes nothing). Therefore, the information set of Agent $(m_0 + m_1 + 1)$ can be denoted by as an ordered pair (m_0, m_1), where m_0 and m_1 are the number of agents having chosen action 0 and 1, respectively. It is clear that an informed agent always takes the correct action (i.e., 0 if ω_0, 1 otherwise). We therefore focus on the decision of uninformed agents.

The decision problem of an agent is the following: given history (m_0, m_1), which of the two states is more likely? If indifferent (that is, if states are equally likely, given (m_0, m_1)), an agent is assumed to choose either action with probability $\frac{1}{2}$. Let $\alpha(m_0, m_1)$ denote the probability with which an agent takes action 1 after history (m_0, m_1).

2 An Illustration: Minorities of One

The possibility of minority wisdom – that it is more likely the minority rather than the majority has chosen correctly – arises independently of the size of the majority, provided only the minority is small enough. In this section we illustrate the logic behind this finding for minorities of the smallest possible size. Considering histories of the form $(1, m)$, we show that an uninformed agent is often better served following the lone dissenter rather than the overwhelming majority (the case of $(m, 1)$ is analogous).

Establishing that following the minority is optimal for an uninformed agent requires the calculation of all possible decision sequences that can produce a particular history. To make these calculations several observations on equilibrium behavior are required. First, by the tie-breaking assumption, an uninformed agent observing a history (m, m) is indifferent and, again by the tie-breaking assumption, he mixes equally over the two actions. Second, uninformed agents always follow a unanimous choice. That is, if all previous agents have chosen identically, such as for the history $(0, m)$, then an uninformed agent chooses action 1. The uninformed agent can be sure no previous agent was informed that action 0 is optimal, but it may be possible than a previous agent was informed that action 1 is optimal. Therefore, no matter how weak this evidence, an uninformed agent imitates the unanimous selection.

We are now in position to calculate the probability of either state given history $(1, m)$. The earliest possibility for a single agent to compose a minority is $m = 2$. Figures 1a and 1b depicts, for each possible state of the world, the possible paths to the history $(1, 2)$, where the history is represented by the corresponding Cartesian point. In state ω_1 (Figure 1b) there is only one possible path to history $(1, 2)$: agent 1 is uninformed and chooses action 0, agent 2 is informed and chooses action 1, and agent 3 is either informed or uninformed and chooses action 1. The probability of this path conditional on state ω_0 is given by

$$\begin{aligned}\Pr(1,2|\omega_1) &= \left(\frac{1-p}{2}\right) p \left(p + \frac{1-p}{2}\right) \\ &= \left(\frac{1-p}{2}\right) p \left(\frac{1+p}{2}\right).\end{aligned}$$

Observe that if agent 1 chooses action 1 then, conditional on state ω_1, history $(1,2)$ is not possible. In fact, unanimous choice is guaranteed as it can be broken only by an informed agent observing an opposing signal which is not possible in state ω_1. This fact proves critical to the illustration presented here for minorities of size 1, and similar ideas will prove critical to the general results of the following section.

** Insert Figure 1a and 1b about here.

For state ω_0, the history $(1,2)$ can be reached by two possible paths (as depicted in Figure 1a), both require the first agent to be uninformed and to choose action 1. The paths diverge for agents 2 and 3 and are as follows: (i) agent 2 is informed and chooses 0, and agent 3 is uninformed and chooses action 1; (ii) agent 2 is uninformed and follows unanimous choice, and agent 3 is informed and chooses 0. The probability of these paths conditional on state ω_0 is given by

$$\begin{aligned}\Pr(1,2|\omega_0) &= \left(\frac{1-p}{2}\right)\left[p\left(\frac{1-p}{2}\right) + (1-p)\,p\right] \\ &= \frac{3}{4}(1-p)^2 p.\end{aligned}$$

An uninformed agent follows the minority if state ω_0 is more likely than state ω_1 given history $(1,2)$; that is, $\Pr(\omega_0|1,2) > \Pr(\omega_1|1,2)$. By Bayes' rule

$$\Pr(\omega_0|1,2) = \frac{\Pr(\omega_0)\Pr(1,2|\omega_0)}{\Pr(\omega_0)\Pr(1,2|\omega_0) + \Pr(\omega_1)\Pr(1,2|\omega_1)}.$$

As the states are equally likely, the condition for minority choice reduces to $\Pr(1,2|\omega_0) > \Pr(1,2|\omega_1)$, which by algebra requires

$$\begin{gathered}\frac{3}{4}(1-p)^2 p > \left(\frac{1-p}{2}\right) p \left(\frac{1+p}{2}\right) \\ \text{or } p < \frac{1}{2}.\end{gathered}$$

Thus, if less than half of all agents are informed, an uninformed agent observing history $(1,2)$ optimally follows the minority and chooses action 0. This result is driven by two factors: the number of possible paths for each state, and the number of agents that are informed versus uninformed. For history $(1,2)$ to be reached in state ω_0, one agent must be informed and two uninformed, and two possible paths reach this history. In contrast, for state ω_1 there exists only one possible path, but it is possible for either one or two agents to be informed. The more likely it is that an agent is informed, therefore, the more likely it is that state ω_1 produced the history. The value $p = \frac{1}{2}$ provides the cut-point on either side of which a different state is more likely to have produced the observed history.

Suppose then that $p < \frac{1}{2}$ and at history $(1,2)$ an uninformed agent follows the minority. The calculations are similar for histories $(1,m)$, where $m \geq 3$, although the bound on minority wisdom, as well as the logic of the result, are slightly different. As depicted in Figures 2a and 2b, only one path to $(1,3)$ exists for each possible state. In state ω_0 (Figure 2a) the final agent must have been informed and deviated from unanimous choice, as if the previous history had been $(1,2)$ then the agent, irrespective of private information, would have chosen with the minority and induced the history $(2,2)$. The reverse path must hold for state ω_1 with the first agent being the dissenter who chose action 0. This path passes through history $(1,2)$ and so the 4th agent must have been informed that state ω_1 is correct. The probabilities for these paths are as follows.

$$\begin{aligned}
\Pr(1,3|\omega_0) &= \left(\frac{1-p}{2}\right)(1-p)^2 p && \text{(1a)} \\
&= \frac{1}{2}(1-p)^3 p. \\
\Pr(1,3|\omega_1) &= \left(\frac{1-p}{2}\right) p \left(\frac{1+p}{2}\right) p && \text{(1b)} \\
&= \frac{1}{4}(1-p)(1+p)p^2.
\end{aligned}$$

Minority choice requires $\Pr(\omega_0|1,3) > \Pr(\omega_1|1,3)$, which by algebra, implies $p < \frac{5-\sqrt{17}}{2} \approx 0.438$.

** Insert Figure 2a and 2b about here.

With an equal number of paths to reach $(1,3)$ for each state, the relative likelihood of the states hinges solely on the numbers of informed and uninformed agents. In state ω_0 there only need be one informed agent to reach history $(1,3)$, although this agent must be in a particular location. In contrast, for history ω_1 there is at least two informed agents, and an additional agent (the third) has the luxury of being informed or uninformed. For a sufficiently small fraction of informed agents, it is more likely that a single agent is informed and, therefore, more likely that the lone minority dissenter is informed rather than uninformed.

The bound on minority wisdom for the history $(1, 3)$ is tighter than for the history $(1, 2)$, and following a minority may not be optimal at $(1, 3)$ despite being optimal at $(1, 2)$. Surprisingly perhaps, this contraction of the bound does not continue for larger populations when the minority is of size one. This can be readily verified for histories $(1, m)$ from the arguments above for history $(1, 3)$. For $m > 3$ there is again only one possible path for each history, and again the dissenter must be the final agent in state ω_0 and the first agent in state ω_1. Note, however, that as minority choice applies at $(1, 2)$, $(1, 3)$, and so on, there must be at least $(m - 1)$ informed agents in state ω_1 but only one informed agent in state ω_0. The probabilities of reaching $(1, m)$, conditional on each state, are therefore generalizations of Equations 1a and 1b and are as follows.

$$\begin{aligned}
\Pr(1, m|\omega_0) &= \frac{1}{2}(1-p)^m p \\
\Pr(1, m|\omega_1) &= \frac{1}{4}(1-p)(1+p)p^{m-1}
\end{aligned}$$

As $p < \frac{1}{2}$, $\Pr(\omega_0|1, 3) > \Pr(\omega_1|1, 3)$ implies $\Pr(\omega_0|1, m) > \Pr(\omega_1|1, m)$ and minority choice persists. Therefore, if an uninformed agent is prepared to follow a lone dissenter in a population of size four, then it is optimal for an uninformed agent to support a lone dissenter in populations of arbitrary size. In fact, in terms of the ratio of probabilities, $\Pr(1, m|\omega_0) / \Pr(1, m|\omega_1)$, the wisdom in a minority of one increases in m, the size of the opposing majority.

3 Properties of the Optimal Strategy

3.1 Minority Choice

The previous section illustrates that, for any $p < \left(5 - \sqrt{17}\right)/2$, minority choice is optimal independently of the majority size, provided only the minority consists of one agent. The reasoning can be generalized to minority sizes of any order, but the larger the minority size, the smaller the critical bound on p, the fraction of informed agents in the population. If we let p_j denote this bound for minorities of j or smaller, so that $p_1 = \left(5 - \sqrt{17}\right)/2$, this means that p_j is strictly decreasing in j. In fact,

$$\lim_{j\to\infty} p_j = 0.$$

It follows that it is not the case, for some $p > 0$, that minority choice is optimal, independently of the minority (and majority) size. Minority choice 'eventually breaks down' for some configurations. In fact, it does not only break down in the neighborhood of the *diagonal* $\{(m, m) : m \in \mathbb{N}\}$, but also arbitrarily far from this diagonal: for any k, no matter how large, there exists m sufficiently large such that it is optimal to follow the majority choice at least k steps away from the diagonal, despite the existence of a minority: $\alpha(m - k', m) = 1$ for some $k' \geq k$, $k' < m$.

This discussion is summarized in the following Proposition:

Proposition 1 *(i)* $\forall j > 0, \exists p_j > 0, \forall p \in (0, p_j), \forall 0 < m_0 < j, \forall m_1 > m_0$ *such that* $\alpha(m_0, m_1) = 0$.

(ii) $\forall k > 0, \exists m$ *and* $m > k' > k$ *such that* $\alpha(m - k', m) = 1$.

Proof: Part (i), which is a tedious generalization of the argument presented in the illustration, is relegated to an Appendix.

Part (ii): Suppose otherwise. That is, suppose that, for some $k > 0$, minority choice is optimal: $\forall m > k' \geq k$, $\alpha(m - k', m) = 0$. In this case, if the state of nature is ω_0, observe that:

$$\Pr\{(m - k', m) \mid \omega_0\} = \frac{1}{2}(1-p)^m p,$$

for the only path that connects $(0,0)$ to $(m - k', m)$, conditional on state ω_0 consists of a string of m uninformed agents who all chose action 1 (the first of which chose that action at random), immediately followed by one informed agent (the agents that then followed chose action 0 independently of their information).

Suppose now that the state of nature is ω_1. For $0 \leq j \leq m - k'$, consider the path along which the first j agents are uninformed, all of which choosing action 0, immediately followed by one informed agent, and then, after another arbitrary $\max\{j - k', 0\}$ agents (so that the 'band of length k' around the diagonal' is reached), by as many uninformed agents as necessary to obtain that, overall, $m - k'$ agents have chosen action 0 (this cannot require more than another $m + k' - j$ agents). When this occurs, the number of agents having chosen action 1 must necessarily be between $m - 2k'$ and m. Therefore, the probability of reaching $(m - k', m)$ following such a path is at least

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

Since there are $m - k'$ possible choices for the integer j, this means that:

$$\Pr\{(m - k', m) \mid \omega_1\} \geq \frac{1}{2}(m - k')(1-p)^{m+k'} p^{2k'+1}.$$

For fixed k', this number may be chosen to be arbitrarily large relative to $\frac{1}{2}(1-p)^m p$, by picking m large enough. This implies that, for such an m,

$$\Pr\{(m - k', m) \mid \omega_1\} > \Pr\{(m - k', m) \mid \omega_0\},$$

so that the optimal action after such a history is action 1, yielding the desired contradiction.■

It is easy to compute the first terms recursively: $p_1 \simeq 0.44$, $p_2 \simeq 0.29$, $p_3 \simeq 0.17$, $p_4 \simeq 0.08$, ..., but a general formula appears elusive. In the proof of the proposition, it is shown that $p_j > 4^{-j}$.

As a corollary of Proposition 1, we can characterize the optimal strategy for all agents up to any N, as long as p is small enough. Indeed, consider the following symmetric minority choice strategy α^*, defined by, for all agents $n \leq N$:

$$\alpha^*(m_0, m_1) = \begin{cases} 0 \text{ if } m_1 = 0, \\ 1 \text{ if } m_0 > m_1 > 1. \end{cases}$$

Corollary 1: Given $N \in \mathbb{N}$, there exists $\bar{p}$, $\forall p < \bar{p}$, α^* is the unique optimal strategy for all agents $n \leq N$.

3.2 Efficiency

If minority choice were always optimal, and informed agents were unlikely ($p < \frac{1}{2}$), information would not be efficiently aggregated over time, as, provided only the first agent made an incorrect choice, the tally would tend to oscillate around the diagonal, with larger departures from ties being caused by strings of informed agents, eventually cancelled out by the larger numbers of uninformed agents that, choosing with the minority, would cause a reversion to a tie. But we already know that minority choice must break down, sooner or later. The question of efficiency becomes therefore nontrivial. What is the limit of the expected fraction of agents who choose correctly, as the number of agents grow large? Given the optimal strategy (obviously unique given our tie-breaking rule), define X_n as the random variable that corresponds to the choice of the nth agent: $X_n = 0$ if he chooses 0, $X_n = 1$ otherwise. Define

$$S_n := \sum_{i=1}^{n} X_n \text{ and } M_n = S_n/n.$$

The optimal strategy is *asymptotically efficient* if $\lim_{n\to\infty} \mathbb{E}\left[M_n \mid \omega_0\right] = 0$. [Obviously, this also implies that $\lim_{n\to\infty} \mathbb{E}\left[M_n \mid \omega_1\right] = 1$], where $\mathbb{E}\left[\cdot\right]$ denotes expectations. A strategy is *efficient* if it minimizes $\mathbb{E}\left[M_n \mid \omega_0\right]$, for all n (or equivalently, $\mathbb{E}\left[S_n \mid \omega_0\right]$). It is true, but not necessarily obvious, that efficient strategies exist. Such a strategy is described in the proof of the following result, based on an application of the welfare improvement principle (see Banerjee and Fudenberg (2004), Smith and Sørensen (1999))

Proposition 2 *The optimal strategy is asymptotically efficient, but not efficient.*

Proof: Consider agent's n probability of making the incorrect decision (before he observes his signal) if he follows the strategy of either following his informative signal, if the case occurs, or mimicking one of his predecessors at random. [The optimal strategy is the one that minimizes this probability.] While this is not a particularly bright strategy, it nevertheless guarantees that this probability does not exceed:

$$q_n = (1-p)\,\frac{\sum_{j=1}^{n-1} q_j}{n-1},$$

where q_j denotes the corresponding probability for agent j, with the understanding that $q_1 = (1-p)/2$. [Of course, q_j is an upper bound on the probability agent $j < n$ takes an incorrect decision, since agent j follows the optimal strategy.] Solving, we get, for all n:

$$q_n = \frac{1-p}{2} \prod_{j=1}^{n-1} \left(1 - \frac{p}{j}\right).$$

Taking logarithms, the convergence of this sequence is equivalent to the divergence of the series $\ln(1 - p/n)$, which follows from the divergence of the series $-p/n$.

Thus, the probability of an incorrect decision tends to 0 along any path. One efficient strategy is: "choose action 0 until at least one agent has chosen action 1 so far" (obviously, there is another one, where the role of 0 and 1 are exchanged). Of course, informed agents follow their information. The argument is rather simple. Consider any agent n, uninformed. If one or more of his predecessors are informed, he is sure to choose the correct action (since either the correct action is 1, and therefore someone has chosen 1 before, so agent n chooses 1; or the correct action is 0, in which case nobody has chosen 1, and agent n chooses 0 as well). If none of his predecessors is informed, then indeed, both actions are equally good, and 0 is optimal. For n agents, the expected number of agents taking the incorrect action under the efficient strategy is:

$$\frac{1-p}{2p}\left(1-(1-p)^n\right),$$

which tends to $(1-p)/(2p)$.

So the total expected number of agents taking an incorrect action is finite under the efficient strategy. The equilibrium strategy is not efficient, because, in some circumstances, while the deciding agent knows that at least one of his predecessors was informed, he is not able to infer what this information was. For instance, if the fourth agent observes two choices for 1 and one choice for 0 -which implies that at least one agent was informed-, it could be that the sequence was $(1,0,1)$, with the second agent being informed of state 0, or it could be that the sequence was $(0,1,1)$, in which case the second agent was informed of state 1. This implies that, among the first three agents, the expected number of agents choosing incorrectly for, say, $p = 1/5$, is $126/125 > 1$, while under the efficient strategy, only $122/125 < 1$ choose incorrectly. (for $p = 1/5$, the equilibrium strategy is minority driven, but suboptimality also obtains under majority choice: for $p = 8/9$, the numbers are $95/1458$ an $91/1458 < 95/1458$ respectively.)$\blacksquare$

While the optimal strategy is not efficient -in fact, for small p, the expected number of agents taking an incorrect action is larger than under an efficient strategy by a factor that is (arbitrarily?) large-, it appears to be the case, from all numerical calculations that we made, that the expected number of incorrect decisions remains bounded, for any $p > 0$. [Clearly, there exists sample paths for which the number of incorrect decisions is arbitrarily large.] This conjecture is easily proven for the case $p > \frac{1}{2}$, but we have not been able to establish it otherwise.[1] As mentioned in the proof of the previous result, this limit is finite for any efficient procedure. However, the proportional sampling scheme used there to prove that $\mathbb{E}[M_n \mid \omega_0]$ tends to zero under the equilibrium strategy does not yield that $\mathbb{E}[S_n \mid \omega_0]$ converges. To the contrary, under this sampling scheme, this number diverges, as is readily verified.

[1] If $p > 1/2$, observe that, by following the majority decision, the nth agent can secure a probability of error not exceeding $(1-p)\, q_{\lceil n/2 \rceil}$, where q_j is the jth agent's probability of error under that method, and $\lceil x \rceil$ is the smallest integer no smaller than x. This implies that the sum of probabilities of errors does not exceed $\sum_{j=1}^{\infty} \frac{1-p}{2} 2^j (1-p)^j = (1-p)^2 / (2p-1)$, which is indeed less than $(1-p)/2p$.

Conjecture 1 $\lim_{n\to\infty}\mathbb{E}[S_n \mid \omega_0] < \infty$.

3.3 Majority Choice

We view the central finding of this paper to be that minority choice is optimal under rather natural circumstances. Nevertheless, it requires, in its simplest expression, the fraction of informed agents not to exceed $\left(5-\sqrt{17}\right)/2$. What about the optimal strategy when p exceeds $\left(5-\sqrt{17}\right)/2$?

It turns out that this does not necessarily imply that majority choice is 'always' optimal. Define the majority choice strategy as the strategy that necessarily mimics the choice of the majority: $\alpha(m_0, m_1) = 0$ if and only if $m_0 > m_1$. The following rather striking conclusion emerges.

Proposition 3 *The majority choice strategy is optimal if and only if $p \geq 7/9$.*

Proof: Under the majority choice strategy, it is possible to describe quite explicitly the conditional probabilities of the various events. Namely,

$$\Pr\{(m_0, m_1) \mid \omega_0\} = \sum_{r=1}^{m_0+1} 2^{-r}\frac{m_1 - m_0 + r - 1}{m_1 + m_0 - r + 1}\binom{m_1 + m_0 - r + 1}{m_0 - r + 1} p^{m_0}(1-p)^{m_1}, \; m_1 > m_0,$$

$$\Pr\{(m, m) \mid \omega_0\} = \Pr\{(m-1, m) \mid \omega_0\},$$

$$\Pr\{(m_0, m_1) \mid \omega_0\} = \frac{1+p}{2}\Pr\{(m_1, m_1) \mid \omega_0\}, \; m_1 < m_0,$$

The first formula follows from the formula for the number $N(m_0, m_1)$ of paths going from $(0,0)$ to (m_0, m_1) touching exactly r times the horizontal axis without ever crossing it (see Mohanty (1979), formula (4.9)):

$$N(m_0, m_1) = \frac{m_1 - m_0 + r - 1}{m_1 + m_0 - r + 1}\binom{m_1 + m_0 - r + 1}{m_0 - r + 1}.$$

The other two formulas are obvious. Trite computations show that, for $n \in 2\mathbb{N}$:

$$\frac{\Pr\{(m, m) \mid \omega_0\}}{p^m(1-p)^m} = \frac{C(2; 2m)}{2^{2m}} := \sum_{r=0}^{2m-1}(-1)^r 2^{-r} C_{2m-1-r} + (-1)^{2m} 2^{-2m},$$

where $C_n = \binom{2n}{n}/(n+1)$ is the Catalan number ($C(2;n)$ is known as a generalized Catalan number).

Let $p(m_0, m_1) = \Pr\{(m_0, m_1) \mid \omega_0\}$. For $p \geq 3/4$, the probability $p(m_0, m_1)$ is non-increasing in m_1, for $m_1 > m_0$. To see this, observe that $p(m_0, m_1 + 1)$, $p(m_0, m_1)$ are given by the aforementioned summations, and consider the ratio of the corresponding summands:

$$\begin{aligned}
\frac{N(m_0, m_1 + 1)}{N(m_0, m_1)}(1 - p) &\leq \frac{1}{4}\frac{N(m_0, m_1 + 1)}{N(m_0, m_1)} \\
&= \frac{1}{4}\frac{m_1 + 1 - (m_0 + 1 - r)}{m_1 - (m_0 + 1 - r)}\frac{m_1 + (m_0 + 1 - r)}{m_1 + 1} \\
&\leq \frac{1}{4}\frac{r + 1}{r}\frac{2m_0 + 2 - r}{m_0 + 2} \leq \frac{1}{4}4 = 1,
\end{aligned}$$

where the second inequality follows by observing that the expression is decreasing in m_1, and the last one from the fact that the expression is decreasing in $r \leq m_0 + 1$.

It follows that majority choice is optimal provided that:

$$\Pr\{(m + 1, m) \mid \omega_0\} > \Pr\{(m, m + 1) \mid \omega_0\}.$$

We will show that this inequality holds for all m if and only if $p \geq 7/9$. The previous inequality is equivalent to:

$$\frac{C(2; n + 1) - C(2; n)}{2C(2; n)} \leq \frac{p}{1 - p},$$

for $n := 2m$. As can be readily verified, $C(2; n + 1)/C(2; n)$ is increasing in n and bounded, and converges therefore to some limit l. It is clear, from the definition of $C(2; n)$, that:

$$\begin{aligned}
C(2; n + 1) + C(2; n) &= 2^{n+1}C_n, \text{ and thus also} \\
C(2; n + 2) - C(2; n) &= 2^{n+2}C_{n+1} - 2^{n+1}C_n.
\end{aligned}$$

From the identity,

$$\frac{C(2; n + 2) - C(2; n + 1)}{C(2; n + 1)} = \frac{C(2; n + 2) - C(2; n)}{C(2; n + 1)} - \frac{C(2; n + 1) - C(2; n)}{C(2; n)}\frac{C(2; n)}{C(2; n + 1)},$$

it follows that:

$$\begin{aligned}
l &= \lim_n \frac{C(2; n + 2) - C(2; n)}{C(2; n + 1) + C(2; n)} = \lim_n \frac{2^{n+2}C_{n+1} - 2^{n+1}C_n}{2^{n+1}C_n} \\
&= 2\lim_n \frac{C_{n+1}}{C_n} - 1 = 2\lim_n \left(4 - \frac{6}{n + 2}\right) - 1 = 7.
\end{aligned}$$

Hence,

$$\frac{C(2; n + 1) - C(2; n)}{2C(2; n)} \leq \frac{p}{1 - p}\forall n \Leftrightarrow p \geq \bar{p} := 7/9.$$

It is clear that the optimal action given some event (m_0, m_1) only depends on the optimal actions at all events (m_0', m_1') for $m_1' < m_1$, and $m_0' \leq m_0$. It follows that there are thresholds p_i such that the majority rule is optimal as long as the minority does not exceed i, provided $p \geq p_i$. The threshold p_i is increasing in i and tends to $\bar{p}$ as $n \to \infty$.■

What happens if p falls short of $7/9$? The optimal strategy is then rather complicated, prescribing majority choice in some circumstances, minority choice in others, and no general principle or comparative statics seems to emerge. The following example illustrates the intricacies of the optimal strategy.

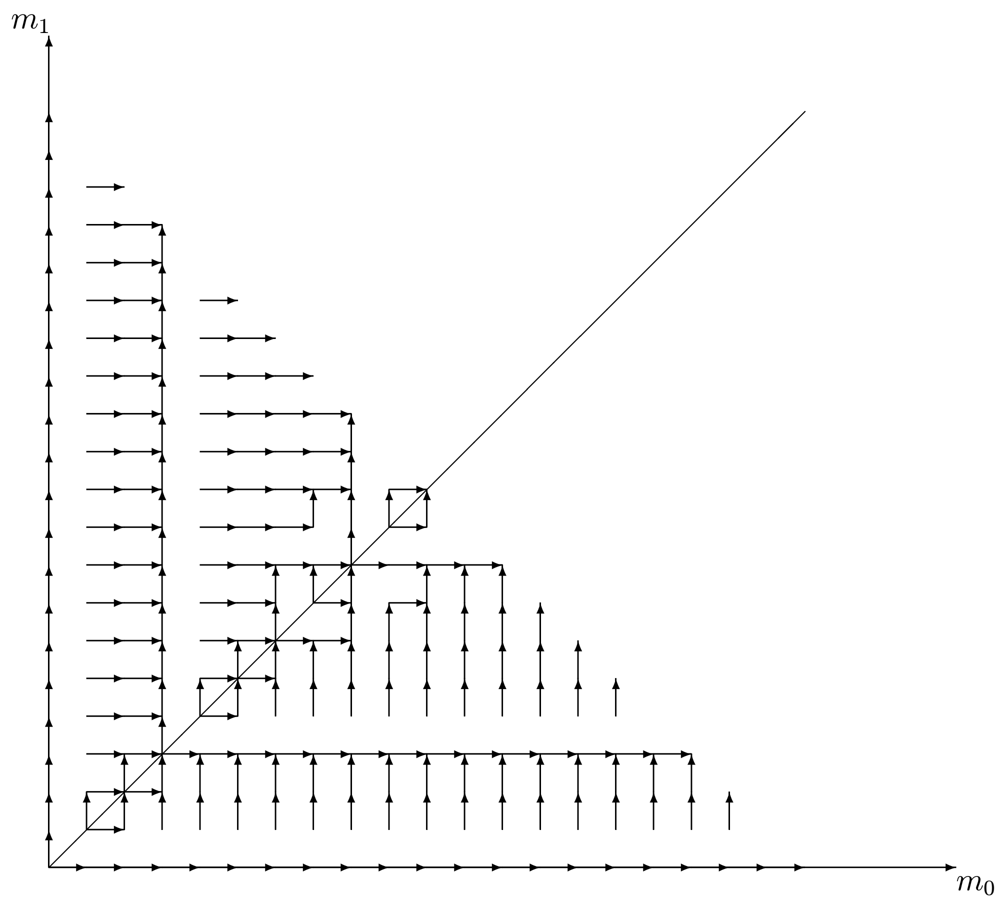

Optimal strategy for uninformed agents ($p = 1/4$)

References:

Banerjee, A., "A Simple Model of Herd Behavior", *Quarterly Journal of Economics*, 107 (1992), 3 (August), pp. 797-817.

Banerjee, A., "The Economics of Rumours", *Review of Economic Studies*, 60 (1993), 2 (April), pp. 309-327.

Banerjee, A. and D. Fudenberg, "Word-of-mouth learning", *Games and Economic Behavior,* 46 (2004), pp. 1–22.

Bikhchandani, S., Hirshleifer, D. and I. Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades", *Journal of Political Economy*, 100 (1992), 5 (October) pp. 992-1026.

Mohanty, G. (1979), *Lattice Path Counting and Applications*, Academic Press, London, U.K.

Smith, L. and P. Sørensen (1999), "Rational Social Learning by Random Sampling", working paper, University of Michigan.

Appendix:

Proof of Proposition 1, part (i): The following remarks will prove useful in the argument.

Remark 1: The problem is obviously symmetric: if one action is optimal when m_1 out of n agents have taken action 1 so far, then the other action is optimal when $n - m_1$ out of n agents have taken that action so far. Formally:

$$\alpha(m_0, m_1) = 1 - \alpha(m_1, m_0).$$

Remark 2: In case of unanimity -that is, if none or all n agents so far have chosen action 1- it is optimal (for an uninformed agent) to take the same action as they have: indeed, there is a positive probability that some of these agents were informed, so that this is the right decision. In case of tie -that is, if exactly half of the n agents have chosen action 1- an uninformed agent is indifferent between both action, and so, according to the tie-breaking rule, takes each action with probability $\frac{1}{2}$.

Remark 3: Let $\Pr\{\omega_i \mid (m_0, m_1)\}$ denote the probability that the state is $\omega_i \in \Omega$, given that m_i agents have taken action $i = 0, 1$, and let $\Pr\{(m_0, m_1) \mid \omega_i\}$ denote the probability that m_i agents take action i, given state $\omega_i \in \Omega$ (conditional, of course, on $m_0 + m_1$ agents exactly having chosen). It follows from Bayes rule that:

$$\begin{aligned}
\Pr\{\omega_i \mid (m_0, m_1)\} &= \frac{\Pr\{(m_0, m_1) \mid \omega_i\} \Pr\{\omega_i\}}{\Pr\{(m_0, m_1) \mid \omega_0\} \Pr\{\omega_0\} + \Pr\{(m_0, m_1) \mid \omega_1\} \Pr\{\omega_1\}} \\
&= \frac{\Pr\{(m_0, m_1) \mid \omega_i\}}{\Pr\{(m_0, m_1) \mid \omega_0\} + \Pr\{(m_0, m_1) \mid \omega_1\}}.
\end{aligned}$$

Therefore,

$$\Pr\{\omega_0 \mid (m_0, m_1)\} > \Pr\{\omega_1 \mid (m_0, m_1)\} \Leftrightarrow \Pr\{(m_0, m_1) \mid \omega_0\} > \Pr\{(m_0, m_1) \mid \omega_1\}.$$

The advantage of working with the probabilities $\Pr\{(m_0, m_1) \mid \omega_i\}$ is that they obey simple recursions, namely:

$$\begin{aligned}
\Pr\{(m_0, m_1) \mid \omega_0\} &= \Pr\{(m_0, m_1 - 1) \mid \omega_0\} \cdot (1 - p)\,\alpha(m_0, m_1 - 1) \\
&\quad + \Pr\{(m_0 - 1, m_1) \mid \omega_0\} \cdot (p + (1 - p)(1 - \alpha(m_0 - 1, m_1))), \\
\Pr\{(m_0, m_1) \mid \omega_1\} &= \Pr\{(m_0, m_1 - 1) \mid \omega_1\} \cdot (p + (1 - p)\,\alpha(m_0, m_1 - 1)) \\
&\quad + \Pr\{(m_0 - 1, m_1) \mid \omega_1\} \cdot (1 - p)(1 - \alpha(m_0 - 1, m_1)),
\end{aligned}$$

with boundary condition $\Pr\{(0, 0) \mid \omega_i\} = 1 \; \forall \omega_i \in \Omega$.

The general argument is by induction, assuming that it holds for $m_1 \leq M - 1$, and proving it holds then for $m_1 = M$ as well, for $p \leq 4^{-M}$. For $m_1 = 1$, it is straightforward to check, using the recursion from Remark 3, that :

$$\begin{aligned}
\Pr\{(m_0, 1) \mid \omega_0\} &= \frac{1}{4} p^{m_0 - 1}(1 - p)(1 + p), \; \forall m_0 \geq 2, \\
\Pr\{(2, 1) \mid \omega_1\} &= 3p(1 - p)^2/4, \; \Pr\{(m_0, 1) \mid \omega_1\} = \frac{1}{2} p (1 - p)^{m_0}, \; \forall m_0 \geq 3,
\end{aligned}$$

so that the conclusion obtains provided $p \leq \left(5 - \sqrt{17}\right)/2$. [Observe that we only need verify the claim for $m_0 \geq 2$.] Conditional on state ω_0, if 1 out of n agents have taken action 1, no more than 2 agents among the n agents could have been uninformed.

Suppose now that the proposition holds for all $m \leq M - 1$, as well as the following claim, valid for $m = 1$: conditional on state ω_0, if $m \leq M - 1$ out of n agents have taken action 1 (where m and n satisfy the assumptions of Proposition 1), then no more than $2m$ agents, out of these n agents, could have been uninformed. We establish these two claims for $m = M$ by induction on the number of agents, n. Suppose that $n = 2M + 1$ (the smallest number of agents that must be considered given the assumptions of Proposition 1). For the second claim, we must show that no more than $2M$ were uninformed, that is, at least one agent was informed. This, however, is obvious, since unanimity would obtain if all agents were uninformed (see Remark 2). For the first claim, observe that:

$$\Pr\left\{(M+1, M) \mid \omega_1\right\} = \frac{1-p}{2} \Pr\left\{(M, M) \mid \omega_1\right\} + \Pr\left\{(M+1, M-1) \mid \omega_1\right\}.$$

To see this, observe that either M out of the first $2M$ agents had taken action 1, in which case, conditional on state 1, only an uninformed agent could take action 0 (furthermore, only with probability $\frac{1}{2}$), or $M-1$ out of the first $2M$ agents had done so, in which case, even an uninformed would have taken action 1 (by the induction hypothesis on M). Similarly:

$$\Pr\left\{(M+1, M) \mid \omega_0\right\} = \left(\frac{1-p}{2} + p\right) \Pr\left\{(M, M) \mid \omega_0\right\} + (1-p) \Pr\left\{(M+1, M-1) \mid \omega_0\right\}.$$

Therefore:

$$\Pr\left\{(M+1, M) \mid \omega_1\right\} > \Pr\left\{(M+1, M) \mid \omega_0\right\}$$
$$\Longleftrightarrow$$
$$\Pr\left\{(M+1, M-1) \mid \omega_1\right\} > p \Pr\left\{(M, M) \mid \omega_0\right\} + (1-p) \Pr\left\{(M+1, M-1) \mid \omega_0\right\},$$

since $\Pr\left\{(M, M) \mid \omega_i\right\}$ is independent of i. Observe that, if M out of $2M$ agents have taken action 1, at least one of them must have been informed. And if $M - 1$ out of $2M$ agents have taken action 1, then, as no more than $2(M-1)$ could have been uninformed, at least two of them must have been informed. Therefore:

$$\begin{aligned} & p \Pr\left\{(M, M) \mid \omega_0\right\} + (1-p) \Pr\left\{(M+1, M-1) \mid \omega_0\right\} \\ \leq\ & p \sum_{i \geq 1} \binom{2M-1}{i} p^i (1-p)^{2M-i} + (1-p) \sum_{i \geq 2} \binom{2M-1}{i} p^i (1-p)^{2M-i} \\ <\ & 2p^2 (1-p)^{2M-1}, \end{aligned}$$

as long as $p < 1/2$. [The fact that the first agent to act must have been uninformed has been used in the binomial coefficient.] In addition:

$$\Pr\left\{(M+1, M-1) \mid \omega_1\right\} = p (1-p)^{M-1}.$$

Therefore, the desired inequality holds for $p \leq 4^{-M}$.

Suppose now that both claims hold for some $n \geq 2M+1$, and suppose that M out of $n+1$ agents have taken action 1. As for the second claim, either $M-1$ out of the first n agents had taken action 1, in which case no more than $2(M-1)+1 < 2M$ out of the $n+1$ agents can be uninformed, or M out of the first n agents had taken action 1. In that case, however, the $(n+1)^{th}$ agent must have informed, as, by the induction hypothesis, an uninformed agent would have taken action 1. Therefore, in this case as well, no more than $2M$ out of the first $n+1$ agents can be uninformed. Next, observe that:

$$\Pr\{(n+1-M, M) \mid \omega_1\} = \frac{1}{2} p (1-p)^M$$

is independent of n. As for $\Pr\{(n+1-M, M) \mid \omega_0\}$, since no more than $2M$ agents were uninformed, it is bounded above by:

$$\sum_{i \leq 2M} \binom{n}{i} (1-p)^i p^{n+1-i}.$$

This upper bound is decreasing in n, for $p \leq 4^{-M}$ [Observe that $\binom{n+1}{i} p^{n+2-i} / \binom{n}{i} p^{n+1-i} = (n+2)p/(n-i+2)$]. The result follows (this inequality is satisfied for $n = 2M+2$ if $p \leq 4^{-M}$).■