ECON 8104

Jakub Pawelczak*

March 2021

Contents

*These notes are intended to summarize the main concepts, definitions and results covered in the first year of micro sequence for the Economics PhD of the University of Minnesota. The material is not my own. Please let me know of any errors that persist in the document. E-mail: pawel042@umn.edu .

1 Herding

Those notes looks into the following issue: why do people tend to converge on similar behavior, in what is known as 'herding'? Why is mass behavior prone to error and fads?

Many of us identify restaurant quality by the fraction of seats occupied; perhaps not coincidentally, restaurants often close off back-room peak-load seating capacity until the main and most visible section becomes quite full. Advertisements report the fractions of doctors or dentists that use certain medications and health products. In the standard herding model, privately informed individuals sequentially see prior actions and then act. An identical action herd eventually starts and public beliefs tend to "cascade sets" where social learning stops. What behaviour is socially efficient when actions ignore informational externalities? To see possible answers we go through following papers:

- Bikchandani, Hirschleifer and Welch(1992) JPE -examples, simple model observable actions and observable signals
- Banerjee (1992) - proposed a remedy for the social learning externality: conceal early actions
- Smith, Sorensen, Tian (2021) forthcoming REStud - main results shrinking cascade sets, contrarianism (the planner skews action choices at the margin)
- Smith, Sorensen (2000) Econometrica

The simplest and most basic cause of convergent behavior is that individuals face similar decision problems, by which we mean that people have similar information, face similar action alternatives, and face similar payoffs. As a result, they make similar choices.

Herding may arise when payoffs are similar even if initial information is not. In this case people **communicate with each other or observe the actions of others or the consequences of these actions**.

We will call influence resulting from rational processing of information gained by observing others **observational learning or social learning**. We focuses mainly on the case where individuals learn by observing the actions of others. Let's formalize our reasonings.

1.1 Bikchandani, Hirschleifer and Welch(1992)

Observable Actions versus Observable Signals

- N individuals starts with some private information,
- obtains some information from predecessors (ordered by time t),
- individuals can observe the actions but not the signals of their predecessors
- and then decides on picking action

Agents decides sequentially whether to adopt or to reject a possible action. The payoff to adopting, V, is either 1 or -1 with equal probability; the payoff to rejecting is 0. In the absence of further information, both alternatives are equally desirable. The order in which individuals decide is given and known to all.

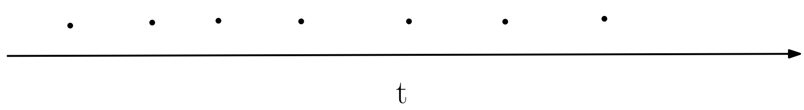

We consider two scenarios. In both, each individual starts with some private information, obtains some information from predecessors, and then decides on a particular action. In the observable actions scenario, individuals can observe the actions but not the signals of their predecessors. We compare this to a benchmark observable signals scenario in which individuals can observe both the actions and signals of predecessors.

$$\Omega = \{L, H\} \quad \text{prior} = (\frac{1}{2}, \frac{1}{2})$$

Each individual's signal is either H or L and H is more likely when adoption is desirable ($V = 1$) than when it is undesirable ($V = -1$). Each observes H with probability $q > 1/2$ if $V = 1$, and with probability $1 - q$ if $V = -1$.

Conditionally IID signals $\sigma \in \{\underline{\sigma}, \bar{\sigma}\}$ are more likely when adoption is desirable $V \in \{L, H\}$. Look at $\mathbb{P}(\sigma|V)$:

$$q = P(\underline{\sigma}|L) > \frac{1}{2} > P(\bar{\sigma}|L)$$

$$q = P(\bar{\sigma}|H) > \frac{1}{2} > P(\underline{\sigma}|H)$$

$$\lambda = \frac{q}{1-q} > 1$$

First player comes after observing only one H, an individual's posterior probability that $V = 1$ is q and the probability that $V = 1$ is only $1 - q$ if he observes L.

What is probability of Low given low signal?

$$P(L|\underline{\sigma}) = P(\underline{\sigma}|L) \cdot \frac{P(L)}{P(\underline{\sigma})} = P(\underline{\sigma}|L) \cdot \frac{P(L)}{P(\underline{\sigma}|L)P(L) + P(\underline{\sigma}|H)P(H)}$$

$$= q \cdot \frac{\frac{1}{2}}{q\frac{1}{2} + (1-q)\frac{1}{2}} = \frac{1}{1 + \frac{1-q}{q}} =$$

$$= \frac{1}{1+\lambda} \geq \frac{1}{2}$$

- So 1st chooses $l \iff$ observe $\underline{\sigma}$

- 1st player plays informatively

- The anaologoue of this is true for h and $\underline{\sigma}$

- Thus, q is the (posterior) probability that the signal is correct.

In the **observable-signals** scenario, the information signals enter the pool of public information one at a time as individuals arrive. Because all past signals are publicly observed, information keeps accumulating so that individuals, all of whom have the same payoffs from taking the same action, eventually settle on the correct choice and thus behave alike.

- Actions reflect information, it is tempting to infer that if only the actions of predecessors are observable, the public information set will also gradually improve until the true value is revealed almost perfectly.

- However, we now show that a scenario of **observable actions** is actually quite different from a scenario of **observable signals**. In the observable-actions case, individuals often converge on the same wrong action

- Furthermore, behavior is idiosyncratic, in the sense that the error-prone choices of a few early individuals determine the choices of all successors.

1st, adopts if his signal is High and rejects if it is Low. All successors can infer 1st's signal perfectly from his decision: if he adopted then he must have observed High and if he rejected he must have observed Low.

Now consider the choice of the second individual. If 1st adopted, then 2nd should also adopt if her private signal is High; as 2nd sees it, there have now been two high signals, the one she inferred from 1st's actions and the one she observed privately. However, if 2nd's private signal is Low, then as she sees it, there has been a High signal (inferred from 1st's actions) and her own Low signal, and she is exactly indifferent between adopting and rejecting. We assume, for expositional simplicity, that as 2nd is indifferent between the two alternatives, she tosses a coin to decide. (By similar reasoning, if 1st rejected, then 2nd should reject if she observes Low, and toss a coin if her signal is High.)

Let's introduce second agent observing second player

$$P(L|\sigma_1, \sigma_2) \equiv P_1(L|\sigma_2)$$

$$P(L|\sigma_1, \sigma_2) = \frac{P(\sigma_2|L, \sigma_1)P(L|\sigma_1)}{P(\sigma_2|\sigma_1)}$$
$$P_1(L|\sigma_2) = \frac{P_1(\sigma_2|L)P_1(L)}{P_1(\sigma_2}$$

Suppose we observe opposite signals: $\sigma_1 = \underline{\sigma}$ and $\sigma_2 = \bar{\sigma}$

$$P_1(L|\bar{\sigma}) = \frac{P_1(\bar{\sigma}|L)P_1(L)}{P_1(\bar{\sigma})} = \frac{P(\bar{\sigma}|L)P(L|\underline{\sigma})}{P_1(\bar{\sigma}|L)P_1(L) + P_1(\bar{\sigma}|H)P_1(H)} =$$
$$= \frac{q\frac{1-q}{1+\frac{1-q}{q}}}{\frac{1-q}{1+\frac{1-q}{q}} + q(1 - \frac{1-q}{1+\frac{1-q}{q}})} = \frac{q(1-q)}{q(1-q) + q - q^2} = \frac{1}{2}$$

The third individual, faces one of three possible situations: both predecessors adopted, both rejected or one adopted and the other rejected. Let's add third player

1. **both 1st and 2nd adopted, then 3rd also adopt**.

$$(a_1, a_2, \sigma_3) = (L, L, \bar{\sigma})$$

$$P(L|\sigma_1, \sigma_2) \equiv P_1(L|\sigma_2)$$

He knows that 1st observed High and that it is more than likely that 2nd observed High too (although she may have seen Low and flipped a coin). Thus, even if 3rd sees a Low signal, he adopts, because he believes that there is better than an even chance that the value to adoption is 1.[3] Consequently, 3rd's decision to adopt provides no information to his successors about the desirability of adopting. The fourth individual, Donna, finds herself in a similar situation as 3rd and adopts regardless of her signal, as will all her successors. 3rd is said to be in an informational cascade, because his optimal action does not depend on his private information, and **the uninformativeness** of 3rd's action means that no further information accumulates. Everyone after 3rd faces the same decision and also adopts based only on the observed actions of Aaron and Barbara. We therefore call this situation an **Up cascade**.

2. **both 1st and 2nd rejected, then 3rd also rejects**. 3rd and all successors reject even if they all privately observed High signals. This is a **Down cascade**.

3. **one rejected and one adopted**. In the remaining case where 1st adopted and 2nd rejected (or vice versa), 3rd knows that 1st observed High and 2nd observed Low (or vice versa). Thus, 3rd's belief based on the actions of the first two individuals is that the High and Low outcomes are equally likely. He finds himself in a situation similar to that of 1st, so 3rd's decision is based only on his private signal. Then, the decision problem of 4th, the next in line, is the same as Barbara's. Aaron's and Barbara's actions have offset and thus carry no information to the fifth individual, 4th. And if 3rd and Donna both

take the same action-say, adopt - then an Up cascade starts with 4th.

Optimal rule Let d be the difference between the number of predecessors who adopted and the number who rejected. Following holds true:

- If $d > 1$, then adopt regardless of private signal.
- If $d = 1$, then adopt if private signal is High and toss a coin if signal is Low.
- If $d = 0$, then follow private signal.
- The decisions for $d = -1$ and $d < -1$ are symmetric.

The fundamental reason the outcome with observable actions is so different from the observable-signals benchmark is that once a **cascade** starts, public information stops accumulating.

Furthermore, the type of cascade depends not just on how many High and Low signals arrive, but the order in which they arrive. For example, if signals arrive in the order HHLL....., then all individuals adopt, because 3rd begins an Up cascade. If, instead, the same set of signals arrive in the order LLHH..... , all individuals reject, because 3rd begins a Down cascade. And if the signals arrive as HLLH .... , then with probability one-half 2nd adopts and 3rd begins an Up cascade. Thus, in the observable-actions scenario, whether individuals on the whole adopt or reject is path-dependent.

Probability of cascades Let the probability that the signal is correct is $p = 0.51$. Then, there is approximately a 75 percent chance that an UP or Down cascade forms after the first two individuals!

- An Up cascade : either 1st and 2nd both High (with p $0.51 \times 0.51 = 0.2601$

) or when 1st High and 2nd Low, flips a coin, and chooses to adopt ($0.51 \times 0.49 \times 0.5 = 0.12495$). Up cascade with 0.38505.

- A Down cascade when 1st and 2nd both Low (with p $0.49 \times 0.49 = 0.2401$) or 1st Low, 2nd High, but flips a coin and reject ($0.49 \times 0.51 \times 0.5 = 0.12495$) . Down cascade with .36505

- Summing up, a cascade occurs with slightly more than 75 percent.

- given that a cascade has occurred, the chance of it being a correct Up cascade rather than an incorrect Down cascade is 51.3 percent ($0.38505/[0.38505 + 0.36505]$).

- individual picking right action .51 based only on the private signal . Gain in accuracy from observing the actions of predecessors is a minimal 0.3 percentage points

- In the observable-signals scenario, publicly observed information signals of predecessors are virtually conclusive as to the right action after many individuals.

- In contrast, when only actions are observed, decisions are little better than when individuals cannot observe predecessors at all.

So far we have argued that cascades are born quickly and idiosyncratically, but they can disappear easily (they ilustrated it in paper). How robust are these conclusions? When some assumptions in the example are relaxed, is the aggregation of information still inefficient or delayed?

1.2 Smith, Sorensen, Tian (2021)

Main results:

- shrinking cascade sets
- contrarianism

Assume that economic theory research fashion is captured by one of two unobserved states, either low-brow theory L or high-brow (intellectual) theory H. A Professor and a Student share a prior belief π on state H.

$$\Omega = \{L, H\} \qquad \pi = P(H)$$

Respectively, they observe conditionally independent draws σ_P, σ_S of a private signal, with cdf's $F^H(\sigma)$ and $F^L(\sigma)$ and densities $f^H(\sigma) = 2\sigma$ and $f^L(\sigma) = 1$, as in Figure 1 .

Two players : Student (S) Professor (P)

Conditionally IID signals with CDFs: $\sigma \in [0, 1]$

$$F_H(\sigma) = \sigma^2. \qquad f_H(\sigma) = 2\sigma$$

$$F_L(\sigma) = \sigma. \qquad f_L(\sigma) = 1$$

$$\frac{f_H}{F_L}(\sigma) = 2\sigma$$

Since the signal likelihood ratio $f^H(\sigma)/f^L(\sigma) = 2\sigma$ in favor of state H increases, higher signals σ lead to higher posterior beliefs p in state H. Also, low $\sigma > 0$ are arbitrarily powerful indicators of state L, but all $\sigma < 1$ have bounded power for H.

After seeing his signal, the Professor either starts a low-brow paper l or high-brow paper h. His Student then learns from his paper choice, and makes his own paper selection. Research pays 1 if the paper and state match, and -1 otherwise (Figure 2). Actions $\{l, h\}$

$$u(\omega) = \begin{cases} 1, & \text{if } \omega = a \\ -1, & \text{otherwise} \end{cases}$$

If the Professor updates to the posterior belief q in state H, his expected payoff is $U(q) \equiv \max(2\pi - 1, 1 - 2\pi)$. We compare two extreme motivations for the Professor: he selfishly only cares about his own expected payoffs, or he is entirely motivated by his Student's expected payoffs.

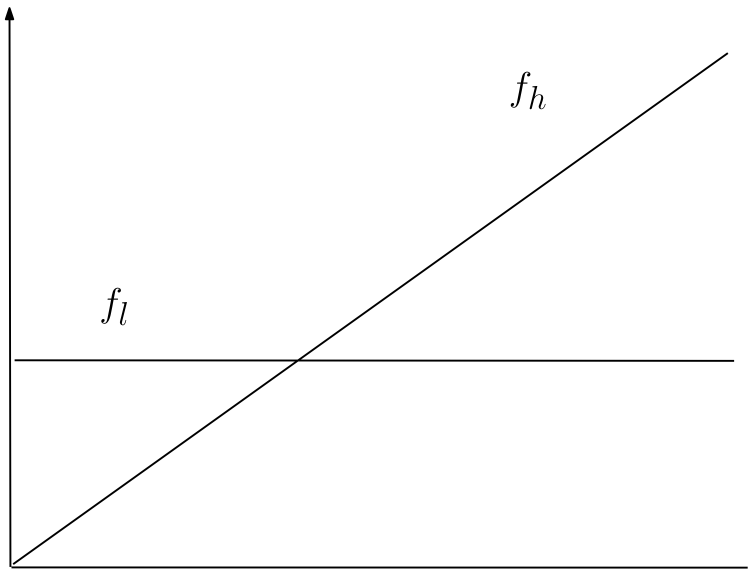

Figure 1: pdfs

Figure 2: payoffs

$$q = P(H|\sigma)$$

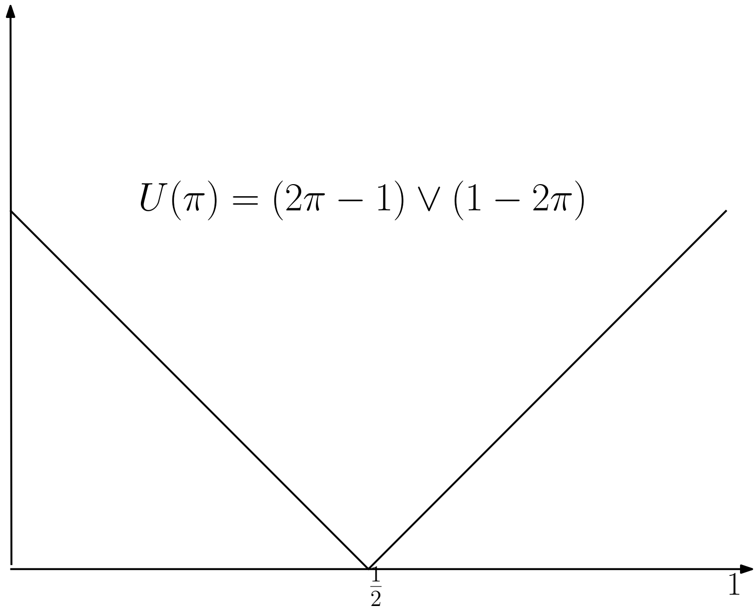

$$\mathbb{E}U_p = U(q)$$

Case 1: The Selfish Professor

Assume the Professor writes paper h when state H is most likely - i.e. for posterior beliefs $q \geq 1/2$. By Bayes rule, this happens when his posterior likelihood ratio of states H to L exceeds one, or $\left[f^H(\sigma)/f^L(\sigma)\right][\pi/(1-\pi)] \geq 1$

Lemma 1. *P writes h* $\iff$ $q \geq \frac{1}{2}$

$$q = P(H|\sigma) \geq \frac{1}{2} \iff \frac{f(\sigma|H)\pi}{F(\sigma|H)\pi + F(\sigma|L)(1-\pi)} \geq \frac{1}{2}$$

$$\frac{1}{1 + \frac{f_l(\sigma)(1-\pi)}{f_h(\sigma)\pi}} \geq \frac{1}{2}$$

$$\frac{f_h(\sigma)\pi}{f_l(\sigma)(1-\pi)} \geq 1$$

This happens for high private signals σ above a selfish threshold signal $\bar{\sigma}(\pi) \equiv (1-\pi)/(2\pi)$. For any prior belief $\pi < 1/3$ in state H, the threshold signal impossibly exceeds one $-$ in this case, the Professor always writes paper l. This event when the prior belief overwhelms all private signals is called a cascade.

This occurs when

$$\sigma \geq \bar{\sigma}(\pi) : \quad \bar{\sigma}(\pi) = \frac{1-\pi}{2\pi}$$

$$\frac{f_h(\sigma)\pi}{f_l(\sigma)(1-\pi)} = 2\sigma\frac{\pi}{1-\pi} \geq 1 \quad \iff \quad \sigma \geq \frac{1-\pi}{2\pi}$$

$$\bar{\sigma}(\frac{1}{3}) = 1$$

$$\forall \quad \pi < \frac{1}{3} \quad \bar{\sigma}(\pi) > 1 \text{ which is unattainable since } \quad \sigma \in [0,1]$$

P writes l in this case $\forall \sigma$.

Here, the Professor's (prior expected) value - or highest expected payoff - is $V(\pi) = U(\pi)$ when $\pi < 1/3$. P's value:

$$V(\pi) = U(\pi) \quad \pi \leq \frac{1}{3}$$

$$\text{if} \quad \pi > \frac{1}{3}$$

$$V(\pi) = \underbrace{\pi}_{H}[\underbrace{(1 - F_H(\bar{\sigma}(\pi)))}_{\sigma > \bar{\sigma}} \underbrace{\cdot(+1)}_{h} + \underbrace{F_H(\bar{\sigma}(\pi))}_{\sigma < \bar{\sigma}} \underbrace{\cdot(-1)}_{l}] + (1-\pi)[(1-F_L(\bar{\sigma}(\pi)))\cdot(-1) + F_L(\bar{\sigma}(\pi))\cdot(+1)]$$

$$= \pi[1 - 2F_H] + (1-\pi)[2F_L - 1] = \pi[1 - 2(\frac{1-\pi}{2\pi})^2] + (1-\pi)[2\frac{1-\pi}{2\pi} - 1] = \cdots =$$

$$= \frac{1}{2}[5\pi - 4 + \frac{1}{\pi}]$$

which is strictly convex on $(\frac{1}{3}, 1)$

S updates beliefs to

$$p = p_l(\pi) \quad \text{or} \quad p_h(\pi)$$

If P sometimes writes l and. h

$$p_l(\pi) < \pi < p_h(\pi)$$

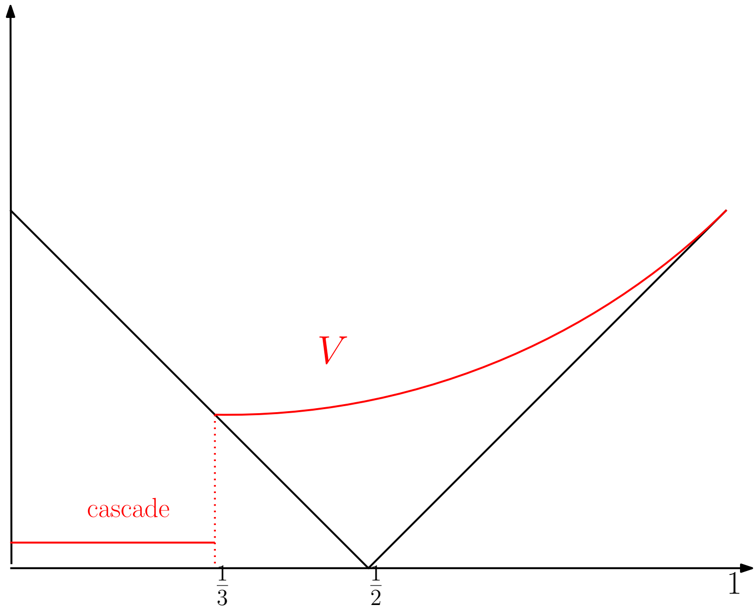

S writes l iff for S's signal τ

$$\tau < \bar{\sigma}(p) = \frac{1-p}{2p}$$

$$\text{S's payoff} \quad = V(p)$$

IF $\pi \leq \frac{1}{3}$ then

$$p = p_l(\pi) \leq \pi$$

P writes l $\Rightarrow$ S writes l too. S copies P!

Case 2: The Altruistic Professor

Consider next that the Professor chooses his paper genre to maximize his Student's expected value. Since $\bar{\sigma}(1/3) = 1$, with a prior belief at or just below 1/3, the selfish Professor always chooses the low brow paper, which sends the student a useless signal.

Suppose P chooses l or h to maximize S's utility

$$\bar{\sigma}(\frac{1}{3}) = 1 \quad \Rightarrow \quad \pi \leq \frac{1}{3}$$

it would lead P to choose l, useless for S.

To help the Student, by informing him of high signals, the altruistic Professor therefore leans against the prevailing prior belief, by choosing a lower altruistic threshold signal $\hat{\sigma} < \bar{\sigma}(\pi)$. In other words, he writes the high brow paper more often, yielding respectively lower continuation beliefs:

To help S, P chooses

$$\hat{\sigma} < \bar{\sigma}(\pi)$$

P writes h more often

$$\hat{p}_l(\pi, \hat{\sigma}) < p_l(\pi|\bar{\sigma}(\pi))$$

$$\hat{p}_h(\pi, \hat{\sigma}) < p_h(\pi|\bar{\sigma}(\pi))$$

$$\hat{p}_l(\pi, \hat{\sigma}) = \frac{P(l|H)P(H)}{P(l)} = \frac{\pi\hat{\sigma}^2}{\pi\hat{\sigma}^2 + (1-\pi)\hat{\sigma}} < \pi < \hat{p}_h(\pi, \hat{\sigma}) = \frac{\pi(1-\hat{\sigma}^2)}{\pi(1-\hat{\sigma}^2) + (1-\pi)(1-\hat{\sigma})}$$

P chooses $\hat{\sigma}$ to max S's value.

If $\hat{p}_l \geq \frac{1}{3}$

$$\mathbb{E}V(p) = \mathbb{E}[\frac{1}{2}(5p - 4 + \frac{1}{p})] = 5\pi - \frac{4}{2} + \frac{1}{2}\mathbb{E}[\frac{1}{p}]$$
$$\mathbb{E}[\frac{1}{p}] = \frac{(\pi\hat{\sigma}^2 + (1-\pi)\hat{\sigma})^2}{\pi\hat{\sigma}^2} + \frac{(\pi(1-\hat{\sigma}^2) + (1-\pi)(1-\hat{\sigma}))^2}{\pi(1-\hat{\sigma}^2)}$$
$$= \cdots = \frac{1}{\pi} + \frac{(1-\pi)^2}{\pi}\frac{1-\hat{\sigma}}{1+\hat{\sigma}}$$

and this function is decreasing in $\hat{\sigma}$ and it is maximized at $\hat{\sigma} = 0$

we can rule out the case where the continuation beliefs obey $\hat{p}_l(\pi) > 1/3$ for any prior belief $\pi > 1/3$

$$\hat{p}_l(\pi) < \frac{1}{3} \quad \forall \pi > \frac{1}{3}$$

Rather, the Professor optimally endows the Student with the two continuation beliefs $\hat{p}_l(\pi) \le 1/3 < \hat{p}_h(\pi)$. Since the Student's value is $V(p) = 1 - 2p$ on $[0, 1/3]$, and $V(p) = 5p - 4 + 1/p$ on $[1/3, 1]$, its expectation is:

$$E[V(P) \mid \pi, \hat{\sigma}] = \left[\pi\hat{\sigma}^2 + (1-\pi)\hat{\sigma}\right](1 - 2\hat{p}_l(\pi)) + \left(\pi\left[1 - \hat{\sigma}^2\right] + (1-\pi)[1-\hat{\sigma}]\right)(5\hat{p}_h(\pi) - 4 + 1/\hat{p}_h(\pi))$$

$$\hat{p}_l(\pi) > \frac{1}{3} \quad \hat{\sigma} > 0 \quad \text{low enough}$$

then

$$\hat{p}_l(\pi) = \frac{1}{1 + \frac{1-\pi}{\pi\hat{\sigma}}} < \frac{1}{3}$$

Unless $p_h(\pi) \ge \frac{1}{3}$, S is in a cascade so gets U

$$V(p) = \begin{cases} 1 - 2p, & \text{on } [0, \frac{1}{3}] \\ 5p - 4 + \frac{1}{p}, & \text{on } [\frac{1}{3}, 1] \end{cases}$$

Optimal wrt to $\hat{\sigma}$

If we substitute the continuation beliefs this expression reduces to

$$E[V(P) \mid \pi, \hat{\sigma}] = V(\pi) + \frac{(3\pi - 1)\left[1 - \pi\hat{\sigma}^2 - \pi\hat{\sigma} - \pi\right]\hat{\sigma}}{\pi(1+\hat{\sigma})} + \frac{\pi(1-\pi)\hat{\sigma}^2(1-\hat{\sigma})}{\pi(1+\hat{\sigma})}$$

Taking logs of the second two terms, the first order condition in $\hat{\sigma}$ then simplifies to

$$\pi^2\hat{\sigma}^3 + 2\pi(5\pi - 1)\hat{\sigma}^2 + 4\pi(2\pi - 1)\hat{\sigma} + 3\pi^2 - 4\pi + 1 = 0$$

Let us first see what this says about the cascade set. At the highest cascade prior belief $\bar{\pi}$, the Professor optimally always chooses the low brow paper. In other words, the altruistic threshold signal is $\hat{\sigma} = 1$. This implies

$$0 = 4\bar{\pi}^2 + 2\bar{\pi}(5\bar{\pi} - 1) + 4\bar{\pi}(2\bar{\pi} - 1) + 3\bar{\pi}^2 - 4\bar{\pi} + 1 = (5\bar{\pi} - 1)^2$$

For an alternative insight here, starting just above the prior belief $\bar{\pi} = 1/5$, if the Professor almost always chooses the low brow paper, the high brow paper signals σ_P very close to one, with likelihood ratio near 2 . Hence, this endows the Student with the continuation belief just above 1/3. This in turn leaves the Student just outside his cascade set. In summary, the altruism shrinks the Professor's set of cascade beliefs from $[0, 1/3]$ to $[0, 1/5]$. Next, one might ask how altruism impacts the Professor's actions at the margin, as he slowly grows more confident in state H. Consider the Professor's posterior odds $2y = 2\pi\hat{\sigma}/(1-\pi) \geq 1/2$ for state H, when he is at the knife-edge. Substituting this expression into (5), the optimal posterior odds $2y(\pi)$ on the domain $1/5 \leq \pi \leq 1$ obey

$$4(1-\pi)^2 y(\pi)^3 + 2(5\pi - 1)(1-\pi)y(\pi)^2 + 4\pi(2\pi - 1)y(\pi) - (3\pi - 1)\pi = 0$$

We can fortunately factor this cubic to get

$$[2(1-\pi)y(\pi) + 3\pi - 1]\left[2(1-\pi)y(\pi)^2 + 2\pi y(\pi) - \pi\right] = 0$$

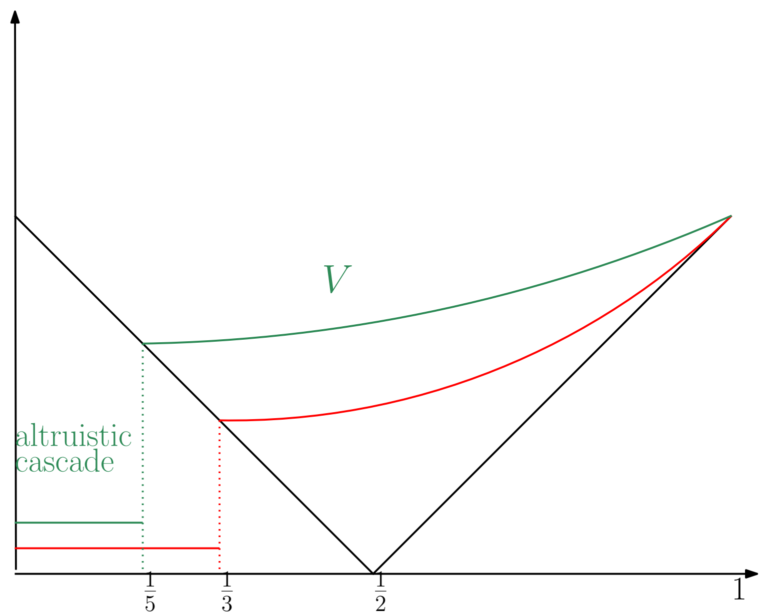

With a cubic derivative with positive lead coefficient, the SOC requires the first or third solution of the FOC. But the first solution is negative. We thus need the positive solution of the quadratic:

$$y(\pi) = \frac{\sqrt{\pi^2 + 2\pi(1-\pi)} - \pi}{2(1-\pi)}$$

This yields the optimal altruistic threshold signal $\hat{\sigma}(\pi) = (1-\pi)y(\pi)/\pi$ falling from 1 to 0 as π increase from 0.2 to 1 . The corresponding threshold posterior belief increases. So the Professor leans more against the high action the higher is π, as his indifference posterior odds $y(\pi)$ are higher. We call this property of an optimal solution **contrarianism**.

1.3 Smith and Sorensen (2000)

Players $n = 1, 2 \ldots$

Prior $= (\frac{1}{2}, \frac{1}{2})$, $\omega \in \{L, H\}$,

$a \in \{1, 2, \ldots, A\}$

State dependent utility functions

$$u_\omega(a) \quad \text{s.t.} u_L(1) > u_L(a) \forall a \neq 1, \quad u_H(A) > u_H(1) \quad \forall a \neq A$$

$$u_H(1) - u_L(1) < u_H(2) - u_L(2) < \cdots < u_H(A) - u_L(A)$$

Action a has myopic payoff :

$$\bar{u}(a, r) = (1-r)u_L(a) + r u_H(a)$$

σ_n n-th private signal IID with

$$\sigma_n = P(H|\sigma_n) \sim F^w$$

F_L and F_H are mutually absolute and signal is the interim belief

$$\sigma = P(H|\sigma) \Rightarrow \quad \frac{dF_H}{dF_L} = \frac{\sigma}{1-\sigma} \quad \Rightarrow \quad F_H(\sigma) \leq F_L(\sigma)$$

means that sampling an individual with signal σ should be just as informative as observing signal σ.

Individuals observe the history of actions but not private signals. Before choosing action a_n, the n'th agent observes the history of $n1$ predecessors' actions and deduce the updated public belief:

$$\pi_n = P(H|a_1, \dots a_n)$$

Combining his private signal σ and public belief π gives the posterior belief

$$r = R(\pi, \sigma) = \frac{\pi\sigma}{\pi\sigma + (1-\pi)(1-\sigma)}$$

This paper explores welfare properties of this model: Abstractly, the planner may modify how agents map private signals into actions, after any given history

A choice rule ξ prescribes action for every signal.

$$a = \psi(\sigma)$$

A strategy s_n for the n'th agent assigns a choice rule to action history of length $n-1$.

$$s_n : (a_1, \dots a_n) \to \psi$$

$$s = (s_1, \dots s_n) \in S$$

The planner's preference depends on a discount factor $\delta \in [0, 1)$ that trades off payoffs earned by present and future agents.

The planner chooses the strategy profile $s = (s_1, s_2, \dots) \in \mathcal{S}$ to maximize the expected present value of the utility stream $u_\omega (a_n)$,:

$$V_\delta(\pi) = \sup_{s \in S} \mathbb{E}[(1-\delta)\sum_{n \geq 1} \delta^{n-1} u_w(s_n)]$$

BHW: deals with $\delta = 0$

DP solution of planner pb A stationary, or Markovian, policy assigns a choice rule ξ for every public belief π, our state variable. With rule ξ, action a happens with probability $\psi(a, \omega, \xi) = \int_{\xi^{-1}(a)} dF^\omega$ in state ω, and unconditionally with probability $\psi(a, \pi, \xi) = \pi\psi(a, H, \xi) + (1-\pi)\psi(a, L, \xi)$ (slightly abusing notation). Action a results in continuation public belief $p(a, \pi, \xi) = \pi\psi(a, H, \xi)/\psi(a, \pi, \xi)$ when $\psi(a, \pi, \xi) > 0$. We call action a, and its continuation belief, active if $p(a, \pi, \xi) > 0$. For any policy, starting at belief π, the continuation value of (10) is a function $v_\delta(\pi)$. By dynamic programming, the optimal (average present) value function v_δ solves the Bellman equation:

$$v(\pi) = \sup_{\xi \in \Xi} (T_\xi v)(\pi)$$

where the policy operator T_ξ maps any continuation value v into the current value, namely:

$$(T_\xi v)(\pi) = \sum_{a=1}^{A} \psi(a, \pi, \xi)[(1-\delta)\bar{u}(a, p(a, \pi, \xi)) + \delta v(p(a, \pi, \xi))]$$

Since the upper envelope of affine functions is convex, the value function v_δ solving is convex, and therefore everywhere has a left and right derivative.

They implement pivot mechanism -namely, one that rewards agents for their marginal contribution to social welfare — here, by changing the public belief. They pay state ω contingent transfers equal successor's incremental values. Characterization? Look into paper but Value Function Tangents determine socially optimal transfers.

The subtangent $\tau(p, r)$ to the value function at public belief p yields the present value for someone with any posterior belief r (thm 1). Thus, higher posterior beliefs raise this value iff the value function slopes up. At extreme posteriors $r = 0, 1$, the tangents at public beliefs $p = p_a$ or $p = \pi$ yield the state-contingent transfers $t_L(a \mid \pi)$ and $t_H(a \mid \pi)$ for action a, directly proportional to ($\propto$) the respective indicated axis gaps.

Corollary 1. *V_S is bounded, continuous and convex in public beleifs π with extremem slopes:*

$$V_\delta'(0+) \geq u_H(1) - u_L(1)$$

$$V_\delta'(1-) \leq u_H(A) - u_L(A)$$

Gittins(1979- old literature on indexes

Theorem 1 (Optimal Behaviour). ***For public belief π, an agent with posterior belief r takes the action a with maximal welfare index***

$$w(a, \pi, r) = (1 - \delta)\bar{u}(a, r) + \delta\tau(p(a, \pi, \xi), r)$$

for some supporting subtangent $\tau(p, r)$ to v at public belief p, when evaluated at posterior r.

Recall that the optimal strategy in the selfish herding model of SS was a simple interval rule: Choose action a if one's posterior belief r lies in an interval I_a, where $\{I_a\}$partitions [0,1] .

Actions with empty intervals are not taken. The rule may randomize at the threshold (boundary) θ_a between adjacent intervals I_a and I_{a+1}.

Since the welfare index $w(a, \pi, r)$ is affine in r, interval rules remain socially optimal.

Theorem 2 (Interval Rules)**. .** ***An interval rule*** $\{I_a\}$ ***is optimal at any public belief*** π***.***