Bot Got Your Tongue? Social Learning with Timidity and Noise

John Cremin
Aix-Marseille School of Economics
SAET-EWET 2026

10-12 2026

Introduction

- In many situations where we learn by observing others' actions, they first choose whether to make said action public.
  - Having made a purchase do you brag about it?
  - After investing in a particular stock do you let others know?
  - Having formed an opinion in a political debate do you post it online?
- And yet the literature on sequential social learning largely neglects this...
- Whenever agents might face some penalty for being caught adopting the 'wrong' position, and are rewarded for being seen to be in the 'right', this will create endogenous observation networks

Preview I

This project:

- Studies sequential social learning where *social* agents can choose the visibility of their action, and *noise* agents act randomly.
- Finds that the outcome depends on whether the observation network is *sparse* or *dense*.

Preview I

This project:

- Studies sequential social learning where *social* agents can choose the visibility of their action, and *noise* agents act randomly.
- Finds that the outcome depends on whether the observation network is *sparse* or *dense*.

Sparse:

- $\exists\ M \in \mathbb{N}$ max $\#$ neighbours
- Unravelling, exaggerated noise and reduced learning

Dense:

- No such M
- No unravelling
- Can facilitate learning in removing *cascade beliefs*.

Outline

Literature

- SSL models have an infinite sequence of agents $n \in \mathbb{N}$ trying to match action x_n to the state θ, with social and private information.

Literature

- SSL models have an infinite sequence of agents $n \in \mathbb{N}$ trying to match action x_n to the state θ, with social and private information.

Background Literature for SSL in general:

- Bikhchandani et al. 1992, Banerjee 1992, Smith & Sorensen 2002
- Acemoglu et al. 2011, Lobel & Sadler 2015 & 2016
- Smith and Sorensen 2008, Monzon & Rapp 2014

Literature

- SSL models have an infinite sequence of agents $n \in \mathbb{N}$ trying to match action x_n to the state θ, with social and private information.

Background Literature for SSL in general:

- Bikhchandani et al. 1992, Banerjee 1992, Smith & Sorensen 2002
- Acemoglu et al. 2011, Lobel & Sadler 2015 & 2016
- Smith and Sorensen 2008, Monzon & Rapp 2014

This project specifically builds on a small 'visibility' literature.

- Guarino et al. 2011, Herrera & Horner 2015

Model

Model

The basic elements of the model are:

- $\theta \in \{0, 1\}$: $\mathbb{P}(\theta = 0) = \frac{1}{2}$
- Agents have (private) type $\tau_n \in \{S, N\}$
  - Choose action $x_n \in \{0, 1\}$ & visibility $v_n \in \{0, 1\}$
  - $\mathbb{P}(\tau = N) = \rho \Rightarrow x_n$ random and $v_n = 1$.
  - $\tau = S$: Timidity Parameter: $t_n \sim^{i.i.d} \Delta([1, \infty))$
  - Confidence $c_n = \frac{1}{t_n}$

$$u_n(x_n, v_n, \theta) = \begin{cases} 1 + v_n & \text{if } x_n = \theta, \\ -(1 + t_n \cdot v_n) & \text{if } x_n \neq \theta, \end{cases}$$

Impact of Timidity

$$u_n(x_n, v_n, \theta) = \begin{cases} 1 + v_n & \text{if } x_n = \theta, \\ -(1 + t_n \cdot v_n) & \text{if } x_n \neq \theta, \end{cases}$$

The ▸▸ Decision Rules that follow imply they will choose x_n as per usual, and $v_n = 1$ if very confident:

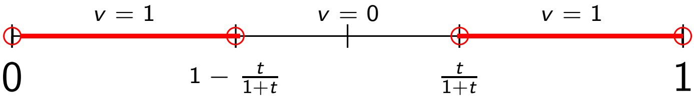

Model II

They have two distinct sources of information:

- A Private Signal: $s_n \in \mathcal{S}_n \sim (\mathbb{F}_0, \mathbb{F}_1) \Rightarrow p_n \sim (\mathbb{G}_0, \mathbb{G}_1)$ (mutually a.c., distinct)
- A Social Signal $\{x_b : b \in B(n)\}$ with $B(n) \sim \mathbb{Q}(n)$ (independent across n)

The Social Signal

There are numerous assumptions about the social signal ('Sample'/'Neighborhood') in the literature:

1. ADLO2011: Non-anonymous & $\mathbb{Q}$ common knowledge
2. SS2008: Anonymous & Unordered Sample & Σ common knowledge
3. MR2014: SS2008 & Uncertainty over own position '$P(i)$'

The Social Signal

There are numerous assumptions about the social signal ('Sample'/'Neighborhood') in the literature:

1. ADLO2011: Non-anonymous & $\mathbb{Q}$ common knowledge
2. SS2008: Anonymous & Unordered Sample & Σ common knowledge
3. MR2014: SS2008 & Uncertainty over own position '$P(i)$'

- This article follows the first approach here, though the unravelling mechanic holds for any of these cases.

Unravelling

Unravelling

- Suppose we have a sparse network topology with common bound M. Any neighbourhood could be entirely noise agents $\Rightarrow [1 - \overline{SB}, \overline{SB}]$

Unravelling

- Suppose we have a sparse network topology with common bound M. Any neighbourhood could be entirely noise agents $\Rightarrow [1-\overline{SB}, \overline{SB}]$
- For a given timidity level t, we can have $v_n = 1$ only if:

$$\mathbb{P}(\theta = 1 | s_n) > \frac{t[1-\overline{SB}]}{t[1-\overline{SB}] + \overline{SB}}$$

$$\text{or } \mathbb{P}(\theta = 1 | s_n) < 1 - \frac{t[1-\overline{SB}]}{t[1-\overline{SB}] + \overline{SB}}$$

Unravelling

- Suppose we have a sparse network topology with common bound M. Any neighbourhood could be entirely noise agents $\Rightarrow [1-\overline{SB}, \overline{SB}]$
- For a given timidity level t, we can have $v_n = 1$ only if:

$$\mathbb{P}(\theta = 1|s_n) > \frac{t[1-\overline{SB}]}{t[1-\overline{SB}] + \overline{SB}}$$

$$\text{or } \mathbb{P}(\theta = 1|s_n) < 1 - \frac{t[1-\overline{SB}]}{t[1-\overline{SB}] + \overline{SB}}$$

Thus, with $t > t^* := \frac{(1-\frac{\rho}{2})^M}{(\frac{\rho}{2})^M}$, we have $v_n = 0$ for all:

$$\mathbb{P}(\theta = 1|s_n) \in [\underbrace{1 - \frac{t(\frac{\rho}{2})^M}{t(\frac{\rho}{2})^M + (1-\frac{\rho}{2})^M}}_{L_0}, \underbrace{\frac{t(\frac{\rho}{2})^M}{t(\frac{\rho}{2})^M + (1-\frac{\rho}{2})^M}}_{U_0}]$$

Timidity & Visibility II

Therefore any social agent chooses $v_n = 0$ with at least probability:

$$\lambda_0 := \min\{\mathbb{G}_0(U_0) - \mathbb{G}_0(L_0), \mathbb{G}_1(U_0) - \mathbb{G}_1(L_0)\}$$

Timidity & Visibility II

Therefore any social agent chooses $v_n = 0$ with at least probability:

$$\lambda_0 := \min\{\mathbb{G}_0(U_0) - \mathbb{G}_0(L_0), \mathbb{G}_1(U_0) - \mathbb{G}_1(L_0)\}$$

Neighbors are drawn only from the $v_n = 1$ pool, thus the relevant probabilities are not $\rho, 1 - \rho$, since at most $(1 - \lambda_0)$ of the social agents choose visible actions.

- $\rho_1^* := \frac{\rho}{\left(\rho+(1-\lambda_0)(1-\rho)\right)}$ is a lower bound for the noise probability.

Timidity & Visibility III

We can then iterate this reasoning ad infinitum:

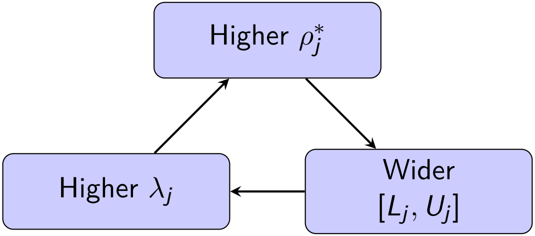

- $\underline{\rho}_\infty = \lim_{j\to\infty} \rho_j^*$ (inc. & bounded above)

Timidity & Visibility IV

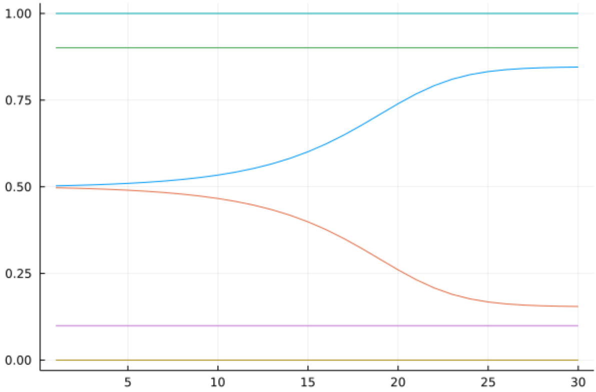

- $f_\theta(\cdot) = g_\theta(\cdot),\ M = 1,\ \rho = 0.2,\ t \approx 9,\ \underline{\rho}_\infty = 0.754...$

Unravelling with a t-Distribution

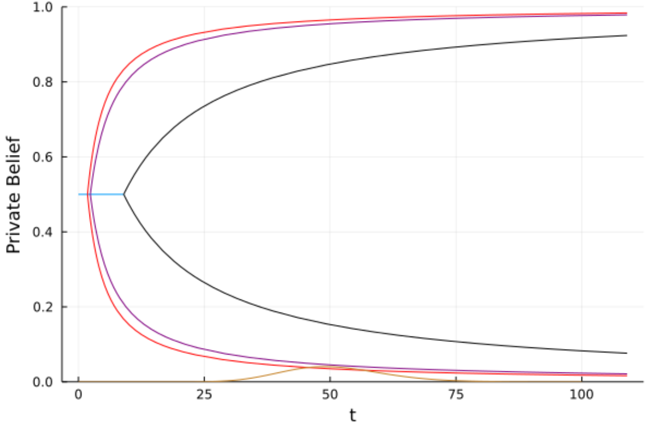

Figure: The parameters for this example are $f_1(x) = 2x$, $f_0(x) = 2 - 2x$, $M = 1$, $\rho = 0.2$; the distribution of $t_n - 1$ is a Chi-squared with 50 degrees of freedom.

Theorem

$\underline{\rho}_\infty = \frac{\rho}{\rho+(1-\lambda_\infty)(1-\rho)}$ is a lower bound on the probability that the nth visible-acting agent is a noise type for any $n \in \mathbb{N}$, and the asymptotic fraction of agents that are noise agents is almost surely greater than $\underline{\rho}_\infty$.

Sparse Networks

Sparse Networks

Timidity and Noise agents have two countervailing effects in general:

- It increases the probability any neighbour is a noise agent (bad)
- It means those neighbours who are not noise agents are acting on the basis of stronger beliefs (good)

Which of these effects dominates asymptotically depends on the network topology, the exact shape of the timidity distribution and the private signal distributions.

Poisson Beta Example

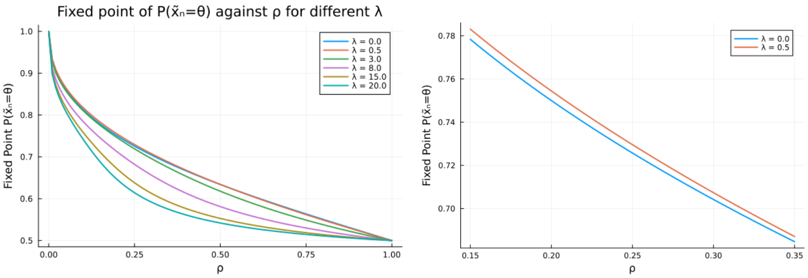

- Agents form an immediate predecessor network, and have private signal and belief density functions $\{2(1-s), 2s\}$.

Normal vs Beta Signals Example

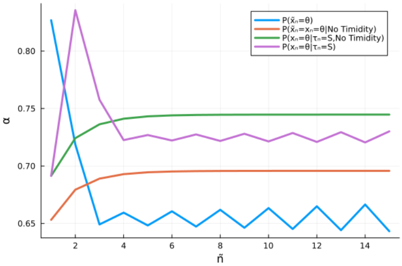

(a) $N(\theta, 1)$

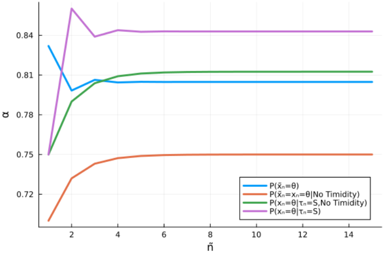

(b) $\{2(1-s), 2s\}$

- Agents form an immediate predecessor network , have timidity that is either 1 or 82 with 0.5 probability each, and $\rho = 0.2$.

Enough Timidity Kills Learning

Proposition

Suppose we have a learning game with $M < \infty$, and consider a sequence of timidity distributions $\Delta^k(t)$ for $k \in \mathbb{N}$, in which each distribution in this sequence is simply the previous one shifted to the right by some fixed value $\delta > 0$. $\underline{\rho}^k_\infty \to 1$, and from this the limiting visible accuracy $\tilde{\alpha}^* \to 0.5$ and accuracy $\alpha^* \to \frac{1}{2}\mathbb{G}_0(0.5) + \frac{1}{2}(1 - \mathbb{G}_1(0.5))$.

Enough Timidity Kills Learning

Proposition

Suppose we have a learning game with $M < \infty$, and consider a sequence of timidity distributions $\Delta^k(t)$ for $k \in \mathbb{N}$, in which each distribution in this sequence is simply the previous one shifted to the right by some fixed value $\delta > 0$. $\underline{\rho}^k_\infty \to 1$, and from this the limiting visible accuracy $\tilde{\alpha}^* \to 0.5$ and accuracy $\alpha^* \to \frac{1}{2}\mathbb{G}_0(0.5) + \frac{1}{2}(1 - \mathbb{G}_1(0.5))$.

- The positive effect of timidity can dominate if we introduce a little, but the bad must eventually dominate.

Enough Timidity Kills Learning II

- Also, the lower is M, the sooner this first effect must dominate in terms of the lower bound $\underline{\rho}_\infty$.
- This is true for the two remaining points in the $\underline{\rho}_\infty$ comparative statics lemma below:

Lemma (Comparative Statics)

$\underline{\rho}_\infty$ is:

- Increasing in ρ.
- Increasing with First order stochastic shifts in $\Delta(t)$.
- Decreasing in M.
- Decreasing with any mean-preserving spread in $\mathbb{G}_0$ or $\mathbb{G}_1$.

Enough Timidity Kills Learning II

- Also, the lower is M, the sooner this first effect must dominate in terms of the lower bound $\underline{\rho}_{\infty}$.
- This is true for the two remaining points in the $\underline{\rho}_{\infty}$ comparative statics lemma below:

Lemma (Comparative Statics)

$\underline{\rho}_{\infty}$ is:

- Increasing in ρ.
- Increasing with First order stochastic shifts in $\Delta(t)$.
- Decreasing in M.
- Decreasing with any mean-preserving spread in $\mathbb{G}_0$ or $\mathbb{G}_1$.

The Improvement Reasoning on which many of these models depend also interacts curiously with timidity.

Dense Networks

Dense Networks

- In dense networks, there is no M, and therefore no unravelling.
- Taking the complete network as the canonical dense network, we can show how learning can benefit from timidity and noise:

Learning in the Complete Network

The following two statements characterise learning in the complete network:

1. With unbounded beliefs, there is complete learning, and in fact the social belief converges to certainty on the true state almost surely.
2. With bounded beliefs supported on $[\underline{p}, \overline{p}] \subset (0, 1)$, and a confidence distribution $\Delta(c)$ such that for some $\epsilon > 0$, $(0, c^* + \epsilon) \cup (1 - \epsilon, 1) \subseteq \text{supp}(\Delta(c))$, where $c^* := \frac{\underline{p}(1-\overline{p})}{\overline{p}(1-\underline{p})}$, the social belief converges to certainty on the true state almost surely.

Timidity and Noise as a Cure for Cascade Beliefs

- Timidity and Noise are both necessary.
  - Without noise, timidity would do if agents could observe neighbours choosing to act invisibly, but here they cannot.
  - Without this additional information, in cascade regions in either state of the world they would just observe the same action forever (and learn nothing).
- With noise agents, there is at least probability $\rho/2$ that any agent will take either action for any social belief, and (with appropriate timidity) the probability action 1 is chosen is *strictly* higher if $\theta = 1$ than $\theta = 0$, hence learning continues at any interior social belief.

Timidity and Noise as a Cure for Cascade Beliefs II

To see this, consider a complete network. The social belief at a given visible history is λ, and we are interested in the probability with which the next visible action is a 1, $\tilde{\psi}(\lambda, \theta)$:

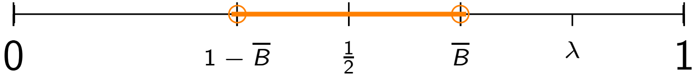

Timidity and Noise as a Cure for Cascade Beliefs II

To see this, consider a complete network. The social belief at a given visible history is λ, and we are interested in the probability with which the next visible action is a 1, $\tilde{\psi}(\lambda, \theta)$:

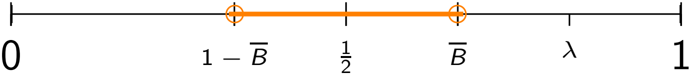

Noise but no Timidity:

- $\tilde{\psi}(\lambda, 1) = 1 - \rho/2$
- $\tilde{\psi}(\lambda, 0) = 1 - \rho/2$
- Learning stops

Timidity and Noise as a Cure for Cascade Beliefs II

To see this, consider a complete network. The social belief at a given visible history is λ, and we are interested in the probability with which the next visible action is a 1, $\tilde{\psi}(\lambda, \theta)$:

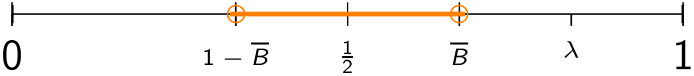

Noise but no Timidity:

- $\tilde{\psi}(\lambda, 1) = 1 - \rho/2$
- $\tilde{\psi}(\lambda, 0) = 1 - \rho/2$
- Learning stops

Timidity but no Noise:

- $\tilde{\psi}(\lambda, 1) = 1$
- $\tilde{\psi}(\lambda, 0) = 1$
- Learning stops

Timidity and Noise as a Cure for Cascade Beliefs III

To see this, consider a complete network. The social belief at a given visible history is λ, and we are interested in the probability with which the next visible action is a 1, $\tilde{\psi}(\lambda, \theta)$:

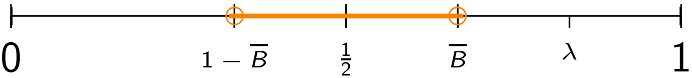

But with *both* noise and timidity satisfying Part 2 of our theorem:

- $\tilde{\psi}(\lambda, 1) = 1 - \mathbb{P}(\tilde{\tau}_n = N | \theta = 1)/2 < 1 - \rho/2$
- $\tilde{\psi}(\lambda, 0) = 1 - \mathbb{P}(\tilde{\tau}_n = N | \theta = 0)/2 < \tilde{\psi}(\lambda, 1)$
- Learning continues...

Benchmarks

Without Noise

In the absence of noise, the impact of timidity is more all-or-nothing:

Proposition

For any bounded private signals, there is some $\overline{c}$ such that if $supp(\Delta(c)) \subseteq [0, \overline{c}]$, every agent's accuracy is the same as $n = 1$, and only noise agents (if there are any) comment visibly.

Without Noise

In the absence of noise, the impact of timidity is more all-or-nothing:

Proposition

For any bounded private signals, there is some $\overline{c}$ such that if $supp(\Delta(c)) \subseteq [0, \overline{c}]$, every agent's accuracy is the same as $n = 1$, and only noise agents (if there are any) comment visibly.

Theorem

With $\rho = 0$, unbounded signals and expanding observations are sufficient for complete learning in every equilibrium.

- Improvement principles still function in this case, and in fact each 'improvement' is greater.

Without Timidity

- Without timidity, noise still hurts:

Proposition

If there is some integer M such that $|B(n)| \leq M$ for any $n \in \mathbb{N}$, and non-zero measure of noise agents ρ, ex-ante accuracy α_n is bounded above by:

$$\alpha_n \leq \frac{1}{2}\mathbb{G}_0\left(\frac{(1-\frac{\rho}{2})^M}{(1-\frac{\rho}{2})^M + (\frac{\rho}{2})^M}\right) + \frac{1}{2}\left(1 - \mathbb{G}_1\left(1 - \frac{(1-\frac{\rho}{2})^M}{(1-\frac{\rho}{2})^M + (\frac{\rho}{2})^M}\right)\right)$$

Without Timidity II

Corollary

With timidity this bound becomes:

$$\frac{1}{2}\mathbb{G}_0\left(\frac{(1-\frac{\rho_\infty}{2})^M}{(1-\frac{\rho_\infty}{2})^M+(\frac{\rho_\infty}{2})^M}\right)+\frac{1}{2}\left(1-\mathbb{G}_1\left(1-\frac{(1-\frac{\rho_\infty}{2})^M}{(1-\frac{\rho_\infty}{2})^M+(\frac{\rho_\infty}{2})^M}\right)\right)$$

- Though unravelling exacerbates this bound.

Self-Locating Information

Self-Locating Information

- A detail I have glossed over here: agents observe $\tilde{n} = \sum_{i=1}^{n-1} v_i + 1$ rather than n.
- By assumption: $\mathbb{P}(\theta = 1|\tilde{n}) = \frac{1}{2}$
- Should Bayesian Agents respond in this way?
  - This assumption reflects a *halfer* response to the *Sleeping Beauty Problem*

Sleeping Beauty Problem

- Inspired by the Absentminded Driver Paradox, the SBP begins on Sunday.

$$\mathbb{P}(H) = 0.5$$

	Heads	Tails
Monday	•	•
Tuesday		•

- For halfers $\mathbb{P}(\theta = H) = \frac{1}{2}$: what is the probability of reaching the information set in each state of the world?
- For thirders $\mathbb{P}(\theta = H) = \frac{1}{3}$: Uniform probability across awakenings.

Arguments behind $\mathbb{P}(\theta = 1|\tilde{n}) = \frac{1}{2}$

In my paper, θ is analogous to the coin toss, and n to the day of the week. Observing $\tilde{n}$ is analogous to observing that you are awake. Thus:

- If halfer: correct. (Noise agents $\Rightarrow$ Each $\tilde{n}$ exists a.s.)
- If halfer in *Duplicating Sleeping Beauty*: correct.
- More behaviourally plausible (requires less sophisticated reasoning)
- If thirder: $\mathbb{P}(\theta = 1|\tilde{n}) = \frac{\mathbb{P}(\tilde{n}|\theta=1)}{\mathbb{P}(\tilde{n}|\theta=1)+\mathbb{P}(\tilde{n}|\theta=0)} = \frac{0}{0+0}$

Could alternatively give agents an improper prior $\epsilon > 0$ on each $n \in \mathbb{N}$. What happens if we do?

Arguments behind $\mathbb{P}(\theta = 1|\tilde{n}) = \frac{1}{2}$

In my paper, θ is analogous to the coin toss, and n to the day of the week. Observing $\tilde{n}$ is analogous to observing that you are awake. Thus:

- If halfer: correct. (Noise agents $\Rightarrow$ Each $\tilde{n}$ exists a.s.)
- If halfer in *Duplicating Sleeping Beauty*: correct.
- More behaviourally plausible (requires less sophisticated reasoning)
- If thirder: $\mathbb{P}(\theta = 1|\tilde{n}) = \frac{\mathbb{P}(\tilde{n}|\theta=1)}{\mathbb{P}(\tilde{n}|\theta=1)+\mathbb{P}(\tilde{n}|\theta=0)} = \frac{0}{0+0}$

Could alternatively give agents an improper prior $\epsilon > 0$ on each $n \in \mathbb{N}$. What happens if we do?

- My benchmark theorem with no noise agents no longer holds...
- But all other results hold anyway, though we may need higher t values for unravelling.

Summary

- I introduce choice over action-visibility into sequential social learning.
- The combined effects of introducing timidity and noise depend on whether the network is sparse or dense.
    - In sparse networks they produce a form of unravelling, leading to an overrepresentation of noise and damaging learning (with enough timidity), though a little timidity can help the latter.
    - In dense networks, noise and timidity can remove cascade beliefs and facilitate learning, unravelling is also no longer a threat.

Improvement Reasoning

Decision Rules

Lemma

After any history, h_n, agent n will choose $x_n = 1$ if

$$\mathbb{P}(\theta = 1|s_n) + \mathbb{P}(\theta = 1|B(n)) > 1$$

and $x_n = 0$ if this sum is strictly less than 1. If they exactly equal 1, the agent is indifferent and assumed to randomise uniformly.

Lemma

After any history h_n, the agent chooses $v_n = 0$ if:

$$1 - (t-1)\mathbb{P}(\theta = 1|B(n))\mathbb{P}(\theta = 1|s_n) < S_n < 1 + \frac{t-1}{t}\mathbb{P}(\theta = 1|B(n))\mathbb{P}(\theta = 1|s_n)$$

for $S_n = \mathbb{P}(\theta = 1|B(n)) + \mathbb{P}(\theta = 1|s_n)$.

Back

Line Network Improvement Principle Graphs

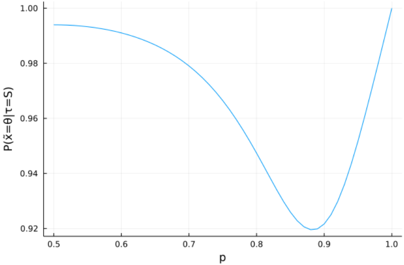

(a) Conditional Success Probability

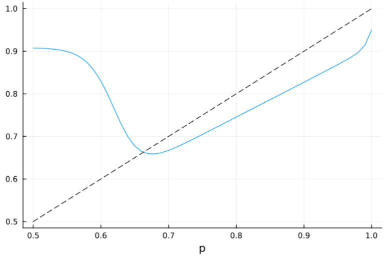

(b) Unconditional Success Probability

- The first of these shows how improvement logic still does hold conditional on the next agent being social (as the curve dominates the 45 degree line), but that providing Blackwell-better information can lead to worse performance.

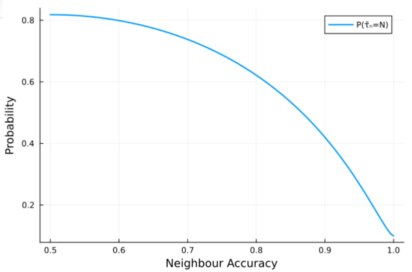

(a) Conditional Success Probability

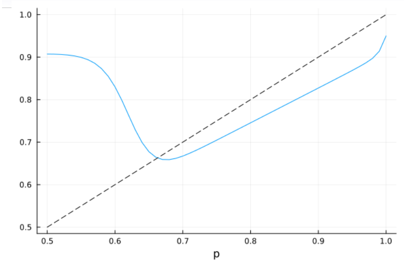

(b) Unconditional Success Probability

- Noise brings the overall graph into contact with the 45 degree line (giving the fixed point to which the game tends).
- With antisymmetric private beliefs noise is always a decreasing function like this.

Back to Sparse Back to Summary