The Impact of Observation and Action Errors on Informational Cascades

Vijay G Subramanian

ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
UNIVERSITY OF MICHIGAN

Joint work with Tho Le & Randall Berry, Northwestern University

Supported by NSF via grant IIS-1219071

CSP Seminar
November 6, 2014

Anecdote[1]

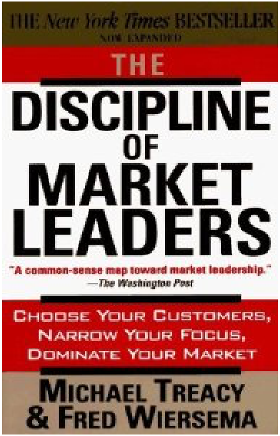

- In 1995 M. Treacy & F. Wiersema published book

[1] Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, Bikhchandani, Hirshleifer & Welch, Journal of Economic Perspectives, 1998

Anecdote[1]

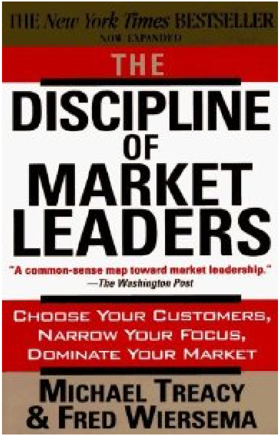

- In 1995 M. Treacy & F. Wiersema published book
- Despite average reviews
  - 15 weeks on *NYTimes* bestseller list
  - *Bloomberg Businessweek* bestseller list
  - $\sim 250K$ copies sold by 2012

[1] Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, Bikhchandani, Hirshleifer & Welch, Journal of Economic Perspectives, 1998

Anecdote[1]

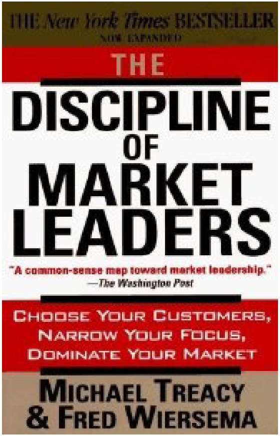

- In 1995 M. Treacy & F. Wiersema published book
- Despite average reviews
  - 15 weeks on *NYTimes* bestseller list
  - *Bloomberg Businessweek* bestseller list
  - $\sim 250K$ copies sold by 2012
- W. Stern of *Bloomberg Businessweek* in Aug'95:
  Authors bought $\sim 10K$ initial copies to make *NYTimes* list
  Increased speaking contracts & fees!

[1] Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, Bikhchandani, Hirshleifer & Welch, Journal of Economic Perspectives, 1998

Anecdote[1]

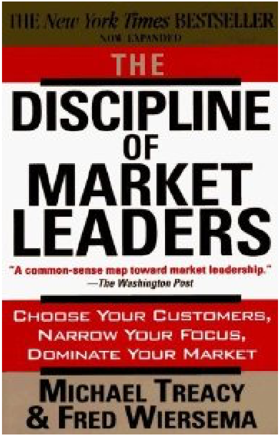

- In 1995 M. Treacy & F. Wiersema published book
- Despite average reviews
  - 15 weeks on *NYTimes* bestseller list
  - *Bloomberg Businessweek* bestseller list
  - $\sim 250K$ copies sold by 2012
- W. Stern of *Bloomberg Businessweek* in Aug'95:
  Authors bought $\sim 10K$ initial copies to make *NYTimes* list
  Increased speaking contracts & fees!
- *NYTimes* changed best-seller list policies in response

[1] Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, Bikhchandani, Hirshleifer & Welch, Journal of Economic Perspectives, 1998

Anecdote[1]

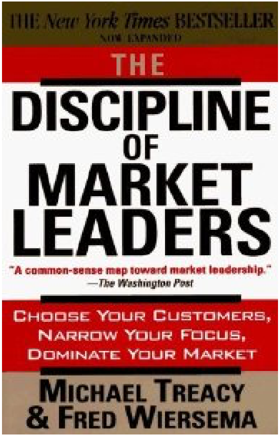

- In 1995 M. Treacy & F. Wiersema published book
- Despite average reviews
  - 15 weeks on *NYTimes* bestseller list
  - *Bloomberg Businessweek* bestseller list
  - $\sim 250K$ copies sold by 2012
- W. Stern of *Bloomberg Businessweek* in Aug'95:
  Authors bought $\sim 10K$ initial copies to make *NYTimes* list
  Increased speaking contracts & fees!
- *NYTimes* changed best-seller list policies in response

Audience greatly influenced by NYTimes' ratings of book

[1] Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, Bikhchandani, Hirshleifer & Welch, Journal of Economic Perspectives, 1998

Motivation

E-commerce, online reviews, collaborative filtering

Motivation

E-commerce, online reviews, collaborative filtering

- E-commerce sites make it easy to find out the actions/opinions of others.
- Future customers can use this information to make their decisions/purchases

Motivation

E-commerce, online reviews, collaborative filtering

- E-commerce sites make it easy to find out the actions/opinions of others.
- Future customers can use this information to make their decisions/purchases

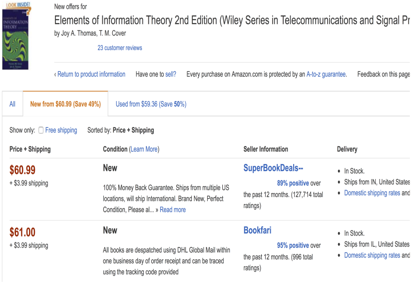

Motivation

E-commerce, online reviews, collaborative filtering

- E-commerce sites make it easy to find out the actions/opinions of others.
- Future customers can use this information to make their decisions/purchases

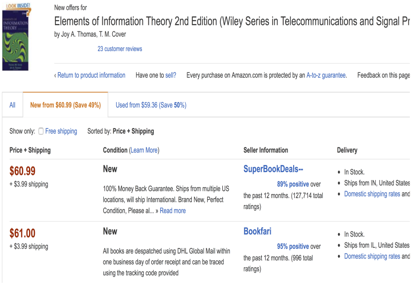

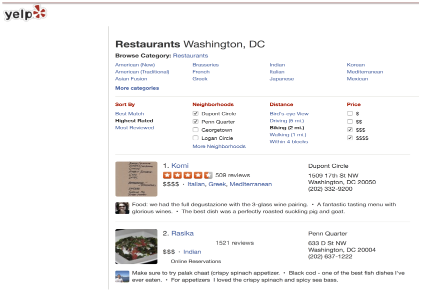

Design Questions

- What is the best information to display?

Design Questions

- What is the best information to display?
- How should one optimally use this information?

Design Questions

- What is the best information to display?
- How should one optimally use this information?
- Can pathological phenomena emerge?

Design Questions

- What is the best information to display?
- How should one optimally use this information?
- Can pathological phenomena emerge?
- What if information is *noisy*?

Design Questions

- What is the best information to display?
- How should one optimally use this information?
- Can pathological phenomena emerge?
- What if information is *noisy*?

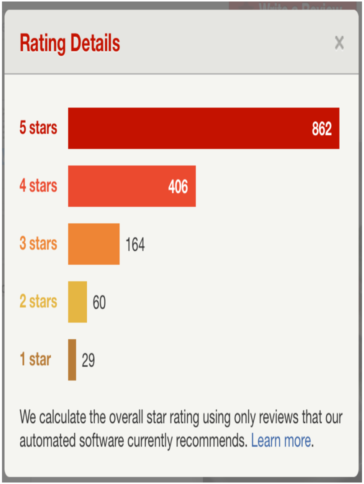

Bayesian Observational Learning

- Model this as a problem of *social learning* or *Bayesian observational learning*

Bayesian Observational Learning

- Model this as a problem of *social learning* or *Bayesian observational learning*
- Studied in economics literature as a *dynamic game with incomplete information*
  - Bikhchandani, Hirshleifer and Welch 1992 [BHW], Banerjee 1992, Smith and Sorensen 2000, Acemoglu *et al.* 2011

Bayesian Observational Learning

- Model this as a problem of *social learning* or *Bayesian observational learning*
- Studied in economics literature as a *dynamic game with incomplete information*
  - Bikhchandani, Hirshleifer and Welch 1992 [BHW], Banerjee 1992, Smith and Sorensen 2000, Acemoglu *et al.* 2011
- Connected to *sequential detection/hypothesis testing*
  - Cover 1969, HellmanCover 1970

BHW model

- An item is available in a market at cost $1/2$

BHW model

- An item is available in a market at cost $1/2$
  - Item's value (V) equally likely Good (1) or Bad (0)

BHW model

- An item is available in a market at cost $1/2$
  - Item's value (V) equally likely Good (1) or Bad (0)
- Agents sequentially decide to Buy or Not Buy the item
  - $A_i = Y$ or $A_i = N$

BHW model

- An item is available in a market at cost $1/2$
  - Item's value (V) equally likely Good (1) or Bad (0)
- Agents sequentially decide to Buy or Not Buy the item
  - $A_i = Y$ or $A_i = N$
- These decisions are recorded via a database

BHW model

- An item is available in a market at cost $1/2$
    - Item's value (V) equally likely Good (1) or Bad (0)
- Agents sequentially decide to Buy or Not Buy the item
    - $A_i = Y$ or $A_i = N$
- These decisions are recorded via a database
- Agent i's payoff, π_i:

Action A_i
- N: payoff $\pi_i = 0$
- Y:
    - payoff $\pi_i = -\frac{1}{2}$ if $V = 0$
    - payoff $\pi_i = +\frac{1}{2}$ if $V = 1$

Information Structure

- Agent i ($i = 1, 2, \ldots$) receives i.i.d. *private signal*, S_i

Information Structure

- Agent i ($i = 1, 2, \ldots$) receives i.i.d. *private signal*, S_i
- Obtained from V via a BSC$(1-p)$

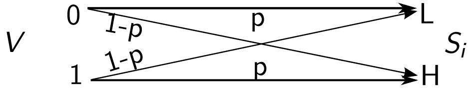

Information Structure

- Agent i ($i = 1, 2, \ldots$) receives i.i.d. *private signal*, S_i
- Obtained from V via a BSC$(1 - p)$

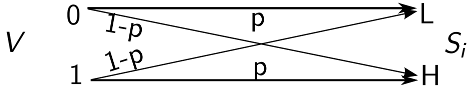

  - Assume $0.5 < p < 1$: Private signal is *informative*, but *non-revealing*

Information Structure

- Agent i ($i = 1, 2, \ldots$) receives i.i.d. *private signal*, S_i
- Obtained from V via a BSC($1 - p$)

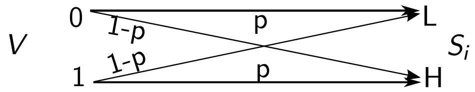

  - Assume $0.5 < p < 1$: Private signal is *informative*, but *non-revealing*

- Agent $i >= 2$ observes actions $A_1, \ldots, A_{i-1}$ in addition to S_i

  Database provides this information

Information Structure

- Agent i ($i = 1, 2, ...$) receives i.i.d. *private signal*, S_i
- Obtained from V via a BSC$(1 - p)$

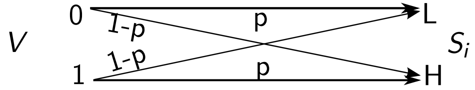

  - Assume $0.5 < p < 1$: Private signal is *informative*, but *non-revealing*

- Agent $i >= 2$ observes actions $A_1, ..., A_{i-1}$ in addition to S_i

  Database provides this information

- Denote the information set as $I_i = \{S_i, A_1, ..., A_{i-1}\}$

Information Structure

- Agent i ($i = 1, 2, ...$) receives i.i.d. *private signal*, S_i
- Obtained from V via a BSC$(1 - p)$

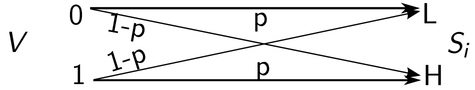

  - Assume $0.5 < p < 1$: Private signal is *informative*, but *non-revealing*

- Agent $i >= 2$ observes actions $A_1, ..., A_{i-1}$ in addition to S_i
  Database provides this information
- Denote the information set as $I_i = \{S_i, A_1, ..., A_{i-1}\}$
- Distribution of value and signals are common knowledge.

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set
- Without any information:

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set
- Without any information:
  - Expected payoff $E[\pi_i] = 0$ since $\mathbb{P}[V = 1] = \mathbb{P}[V = 0] = \frac{1}{2}$

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
    - Given her infromation set
- Without any information:
    - Expected payoff $E[\pi_i] = 0$ since $\mathbb{P}[V = 1] = \mathbb{P}[V = 0] = \frac{1}{2}$
- With *only* private signal:

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set
- Without any information:
  - Expected payoff $E[\pi_i] = 0$ since $\mathbb{P}[V = 1] = \mathbb{P}[V = 0] = \frac{1}{2}$
- With *only* private signal:
  - Update posterior probability:
    $Pr(V = G|S_i = H) = Pr(V = B|S_i = L) = p > 0.5$

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set
- Without any information:
  - Expected payoff $E[\pi_i] = 0$ since $\mathbb{P}[V = 1] = \mathbb{P}[V = 0] = \frac{1}{2}$
- With *only* private signal:
  - Update posterior probability:
    $Pr(V = G|S_i = H) = Pr(V = B|S_i = L) = p > 0.5$
  - Optimal Action: Buy if and only if $S_i = H$.

Bayesian Rational Agents

- Suppose each agent seeks to maximize her expected pay-off.
  - Given her infromation set
- Without any information:
  - Expected payoff $E[\pi_i] = 0$ since $\mathbb{P}[V = 1] = \mathbb{P}[V = 0] = \frac{1}{2}$
- With *only* private signal:
  - Update posterior probability:
    $Pr(V = G|S_i = H) = Pr(V = B|S_i = L) = p > 0.5$
  - Optimal Action: Buy if and only if $S_i = H$.
  - Pay-off: $E[\pi_i] = \frac{1}{2}\left(\frac{2p-1}{2}\right) + \frac{1}{2}(0) = \frac{2p-1}{4} > 0$

Bayesian Rational Agents cont'd.

- With private signal S_i and actions $A_1, \ldots, A_{i-1}$:

Bayesian Rational Agents cont'd.

- With private signal S_i and actions $A_1, ..., A_{i-1}$:
  - Update posterior probability $\mathbb{P}[V = 1 | I_i] = \frac{\mathbb{P}[I_i | V=1]}{\mathbb{P}[I_i | V=1] + \mathbb{P}[I_i | V=0]}$

Bayesian Rational Agents cont'd.

- With private signal S_i <u>and</u> actions $A_1, ..., A_{i-1}$:
  - Update posterior probability $\mathbb{P}[V = 1 | I_i] = \frac{\mathbb{P}[I_i | V=1]}{\mathbb{P}[I_i | V=1] + \mathbb{P}[I_i | V=0]}$
  - Decision:

$$\text{Action } A_i \begin{cases} Y \text{ if } \mathbb{P}[V = 1 | I_i] > \frac{1}{2} \\ N \text{ if } \mathbb{P}[V = 1 | I_i] < \frac{1}{2} \\ \text{follow own signal if } \mathbb{P}[V = 1 | I_i] = \frac{1}{2} \end{cases}$$

Bayesian Rational Agents cont'd.

- With private signal S_i and actions $A_1, \ldots, A_{i-1}$:
  - Update posterior probability $\mathbb{P}[V = 1 | I_i] = \frac{\mathbb{P}[I_i | V=1]}{\mathbb{P}[I_i | V=1] + \mathbb{P}[I_i | V=0]}$
  - Decision:

Action A_i →
- Y if $\mathbb{P}[V = 1 | I_i] > \frac{1}{2}$
- N if $\mathbb{P}[V = 1 | I_i] < \frac{1}{2}$
- follow own signal if $\mathbb{P}[V = 1 | I_i] = \frac{1}{2}$

- Can now iteratively calculate the actions of each agent for a given realization of V and $\{S_i\}$.

BHW'92 Analysis

- First agent always follows their own signal.

---

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
  - Two possibilities:
    - Observation and signal match.
    - Observation and signal differ.

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
  - Two possibilities:
    - Observation and signal match.
    - Observation and signal differ.
  - In second case, agent is indifferent between following signal or not.[2]

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
    - Two possibilities:
        - Observation and signal match.
        - Observation and signal differ.
    - In second case, agent is indifferent between following signal or not.[2]
- Third agent?

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
  - Two possibilities:
    - Observation and signal match.
    - Observation and signal differ.
  - In second case, agent is indifferent between following signal or not.[2]
- Third agent?
  - Interesting cases: $I_3 = \{H, N, N\}$ or $\{L, Y, Y\}$.

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
  - Two possibilities:
    - Observation and signal match.
    - Observation and signal differ.
  - In second case, agent is indifferent between following signal or not.[2]
- Third agent?
  - Interesting cases: $I_3 = \{H, N, N\}$ or $\{L, Y, Y\}$.
  - In these cases, optimal action is to "follow the crowd"

[2] Here assume they always follow signal in this case.

BHW'92 Analysis

- First agent always follows their own signal.
- Consider second agent.
  - Two possibilities:
    - Observation and signal match.
    - Observation and signal differ.
  - In second case, agent is indifferent between following signal or not.[2]
- Third agent?
  - Interesting cases: $I_3 = \{H, N, N\}$ or $\{L, Y, Y\}$.
  - In these cases, optimal action is to "follow the crowd"
- Subsequent agents?

[2] Here assume they always follow signal in this case.

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal
    - Consequences:

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal
  - Consequences:
    - Non-zero probability of herding in wrong decision

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal
  - Consequences:
    - Non-zero probability of herding in wrong decision
    - Private signal information lost

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal
  - Consequences:
    - Non-zero probability of herding in wrong decision
    - Private signal information lost
  - BHW'92, Banerjee'92, Welch'92: Agents eventually exhibit herding

Herding

- Definition: *Herding* (informational cascade) occurs when it is *optimal* for an agent to *take a fixed action* based on other agents' actions, *regardless* of her own signal
  - Consequences:
    - Non-zero probability of herding in wrong decision
    - Private signal information lost

  - BHW'92, Banerjee'92, Welch'92: Agents eventually exhibit herding
  - BHW'92: herding as soon as $|\#Y's - \#N's| = 2$ in the history.

    Once herding starts, all agents follow suit.

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
    - Two urns with mix of red and blue balls.

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
  - Two urns with mix of red and blue balls.
  - One has a majority of blue/one majority red.
  - One Urn selected and identity kept secret.

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
  - Two urns with mix of red and blue balls.
  - One has a majority of blue/one majority red.
  - One Urn selected and identity kept secret.
  - Students take turns drawing one ball from the selected urn, then guessing which urn it is.

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
  - Two urns with mix of red and blue balls.
  - One has a majority of blue/one majority red.
  - One Urn selected and identity kept secret.
  - Students take turns drawing one ball from the selected urn, then guessing which urn it is.
    - Only see the color of the drawn ball
    - Students see all previous students' guesses

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
    - Two urns with mix of red and blue balls.
    - One has a majority of blue/one majority red.
    - One Urn selected and identity kept secret.
    - Students take turns drawing one ball from the selected urn, then guessing which urn it is.
        - Only see the color of the drawn ball
        - Students see all previous students' guesses
    - Experiment is repeated, each time the urn is chosen randomly.

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
  - Two urns with mix of red and blue balls.
  - One has a majority of blue/one majority red.
  - One Urn selected and identity kept secret.
  - Students take turns drawing one ball from the selected urn, then guessing which urn it is.
    - Only see the color of the drawn ball
    - Students see all previous students' guesses
  - Experiment is repeated, each time the urn is chosen randomly.
  - Students with correct guess will be rewarded after the experiment

Do real people herd?

- Ball-drawing experiment (Anderson & Holt 1997)
  - Two urns with mix of red and blue balls.
  - One has a majority of blue/one majority red.
  - One Urn selected and identity kept secret.
  - Students take turns drawing one ball from the selected urn, then guessing which urn it is.
    - Only see the color of the drawn ball
    - Students see all previous students' guesses
  - Experiment is repeated, each time the urn is chosen randomly.
  - Students with correct guess will be rewarded after the experiment
  - Result: About 80% of the cases the students copy guesses.

Discussion

Why does herding happen? OR When can learning occur?

Discussion

Why does herding happen? OR When can learning occur?

- Discrete feedback from agents is not rich enough
  - Cover1969, SmithSorensen2000: reporting posterior beliefs better
  - Cover1969, Hellman thesis: Can reduce to finite memory of display
- Likelihood ratios of private signals bounded
  - SmithSorensen2000, Sorensen2000 thesis: if unbounded, then learning occurs

Discussion

Why does herding happen? OR When can learning occur?

- Discrete feedback from agents is not rich enough
  - Cover1969, SmithSorensen2000: reporting posterior beliefs better
  - Cover1969, Hellman thesis: Can reduce to finite memory of display
- Likelihood ratios of private signals bounded
  - SmithSorensen2000, Sorensen2000 thesis: if unbounded, then learning occurs
- Bayes update plus threshold rule may not be optimal
  - Cover1969: different rule with finite memory display
  - Zhang *et al.* 2013: different sequence of thresholds gives learning

Discussion

Why does herding happen? OR When can learning occur?

- Discrete feedback from agents is not rich enough
  - Cover1969, SmithSorensen2000: reporting posterior beliefs better
  - Cover1969, Hellman thesis: Can reduce to finite memory of display
- Likelihood ratios of private signals bounded
  - SmithSorensen2000, Sorensen2000 thesis: if unbounded, then learning occurs
- Bayes update plus threshold rule may not be optimal
  - Cover1969: different rule with finite memory display
  - Zhang *et al.* 2013: different sequence of thresholds gives learning
- Information structure reinforces actions
  - Acemoglu *et al.* 2011: changing set of past agents sampled gives learning even with bounded likelihoods

Discussion

Why does herding happen? OR When can learning occur?

- Discrete feedback from agents is not rich enough
  - Cover1969, SmithSorensen2000: reporting posterior beliefs better
  - Cover1969, Hellman thesis: Can reduce to finite memory of display
- Likelihood ratios of private signals bounded
  - SmithSorensen2000, Sorensen2000 thesis: if unbounded, then learning occurs
- Bayes update plus threshold rule may not be optimal
  - Cover1969: different rule with finite memory display
  - Zhang *et al.* 2013: different sequence of thresholds gives learning
- Information structure reinforces actions
  - Acemoglu *et al.* 2011: changing set of past agents sampled gives learning even with bounded likelihoods

Why should strategic users follow any of these remedial schemes?

Noisy Observations

- Introduce i.i.d. observation errors

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

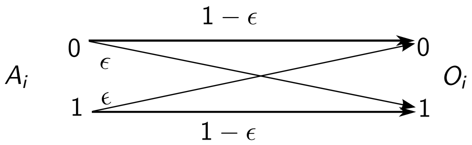

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

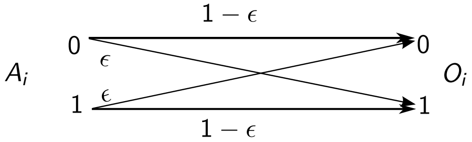

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

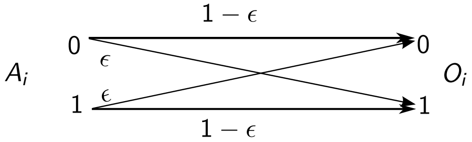

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$
- **Objective**: Study the effects of such errors on BHW model

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

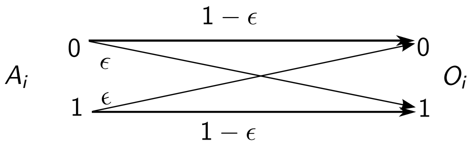

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$
- **Objective**: Study the effects of such errors on BHW model
  - Note with noisy, observations are less reliable

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

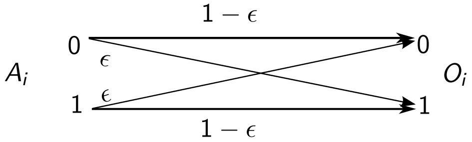

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$
- **Objective**: Study the effects of such errors on BHW model
  - Note with noisy, observations are less reliable
  - Does herding still occur?

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

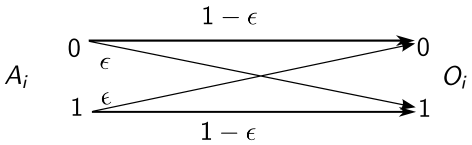

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$
- **Objective**: Study the effects of such errors on BHW model
  - Note with noisy, observations are less reliable
  - Does herding still occur?
  - How does probability of wrong herding change?

Noisy Observations

- Introduce i.i.d. observation errors
  - Actions are recorded on *common database* via another BSC(ϵ), $0 < \epsilon < 0.5$

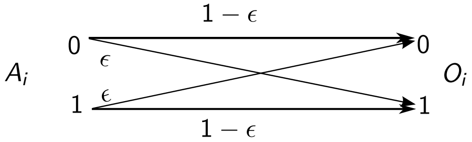

  - Information set is now $I_i = \{S_i, O_1, ..., O_{i-1}\}$
- **Objective**: Study the effects of such errors on BHW model
  - Note with noisy, observations are less reliable
  - Does herding still occur?
  - How does probability of wrong herding change?
  - Can parameters be changed to improve things?

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	herding iff $O_1 = O_2$ and $\epsilon < \epsilon^*(3, p)$

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	herding iff $O_1 = O_2$ and $\epsilon < \epsilon^*(3, p)$
Agent n		

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	herding iff $O_1 = O_2$ and $\epsilon < \epsilon^*(3, p)$
Agent n	herding iff $\|\#Y's - \#N's\| \geq 2$	

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \mid S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \mid S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	herding iff $O_1 = O_2$ and $\epsilon < \epsilon^*(3, p)$
Agent n	herding iff $\lvert \#Y's - \#N's \rvert \geq 2$	herding iff $\lvert \#Y's - \#N's \rvert \geq k$ and $\epsilon < \epsilon^*(k+1, p)$ for some integer $k \geq 2$

Herding in noiseless and noisy models

	Noiseless Model $\epsilon = 0$	**Noisy Model** $\epsilon > 0$
Available Information	$\{S_i, A_1, ..., A_{i-1}\}$	$\{S_i, O_1, ..., O_{i-1}\}$
Posterior Probability	$\mathbb{P}[V = 1 \| S_i, A_1, ..., A_{i-1}]$	$\mathbb{P}[V = 1 \| S_i, O_1, ..., O_{i-1}]$
Agent 1	Follows private signal S_1	Follows private signal S_1
Agent 2	Follows private signal S_2	Follows private signal S_2
Agent 3	herding iff $A_1 = A_2$	herding iff $O_1 = O_2$ and $\epsilon < \epsilon^*(3, p)$
Agent n	herding iff $\|\#Y's - \#N's\| \geq 2$	herding iff $\|\#Y's - \#N's\| \geq k$ and $\epsilon < \epsilon^*(k + 1, p)$ for some integer $k \geq 2$

- We can obtain closed-form expression for $\epsilon^*(k + 1, p)$ (thresholds)

Noise thresholds

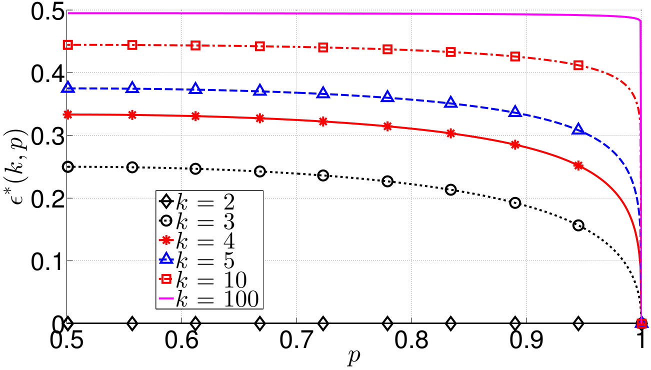

Summary of herding property

Model inherits many behaviors of noiseless model ([BHW'92], $\epsilon = 0$)

Summary of herding property

Model inherits many behaviors of noiseless model ([BHW'92], $\epsilon = 0$)

- **Property 1** Until herding occurs, each agent's Bayesian update depends only on their private signal and the difference $(\#Y's - \#N's)$ in the observation history

Summary of herding property

Model inherits many behaviors of noiseless model ([BHW'92], $\epsilon = 0$)

- **Property 1** Until herding occurs, each agent's Bayesian update depends only on their private signal and the difference ($\#Y's - \#N's$) in the observation history
- **Property 2** Once herding happens, it lasts forever

Summary of herding property

Model inherits many behaviors of noiseless model ([BHW'92], $\epsilon = 0$)

- **Property 1** Until herding occurs, each agent's Bayesian update depends only on their private signal and the difference $(\#Y's - \#N's)$ in the observation history
- **Property 2** Once herding happens, it lasts forever
- **Property 3** Given $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$, if any time in the history $|\#Y's - \#N's| \geq k$, then herding will start

Summary of herding property

Model inherits many behaviors of noiseless model ([BHW'92], $\epsilon = 0$)

- **Property 1** Until herding occurs, each agent's Bayesian update depends only on their private signal and the difference $(\#Y's - \#N's)$ in the observation history
- **Property 2** Once herding happens, it lasts forever
- **Property 3** Given $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$, if any time in the history $|\#Y's - \#N's| \geq k$, then herding will start
  - Eventually herding happens (in finite time)

Markov chain viewpoint

- Assume $V = 1$ and $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$

Markov chain viewpoint

- Assume $V = 1$ and $\epsilon^*(k, p) \le \epsilon < \epsilon^*(k+1, p)$
- State at time i is $(\#Y's - \#N's)$ seen by an agent i

Markov chain viewpoint

- Assume $V = 1$ and $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$
- State at time i is $(\#Y's - \#N's)$ seen by an agent i
- Time index = agent's index

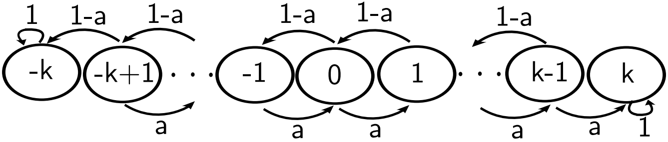

Markov chain viewpoint

- Assume $V = 1$ and $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$
- State at time i is $(\#Y's - \#N's)$ seen by an agent i
- Time index = agent's index

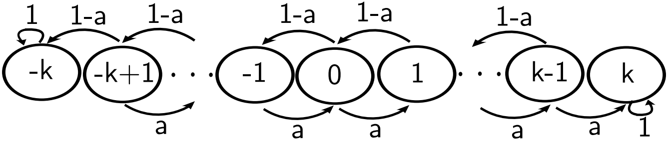

- Agent 1 starts at state 0
- $a = \mathbb{P}[One\ more\ Y\ added] = (1-\epsilon)p + \epsilon(1-p) > 0.5$, decreasing in ϵ, increasing in p

Markov chain viewpoint

- Assume $V = 1$ and $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$
- State at time i is $(\#Y's - \#N's)$ seen by an agent i
- Time index = agent's index

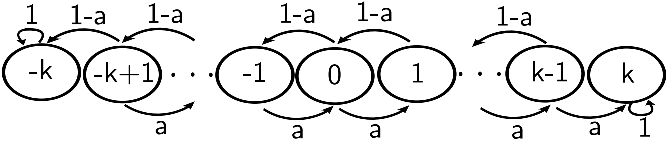

- Agent 1 starts at state 0
- $a = \mathbb{P}[\textit{One more Y added}] = (1-\epsilon)p + \epsilon(1-p) > 0.5$, decreasing in ϵ, increasing in p
- Absorbing state k: herd Y, Absorbing state $-k$: herd N

Markov Chain viewpoint (continued)

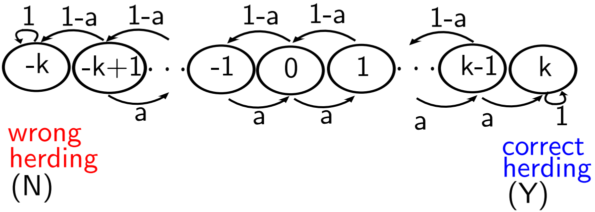

- Can exactly calculate expected payoff $E[\pi_i]$ & probability of wrong (correct) herding for any agent i

Markov Chain viewpoint (continued)

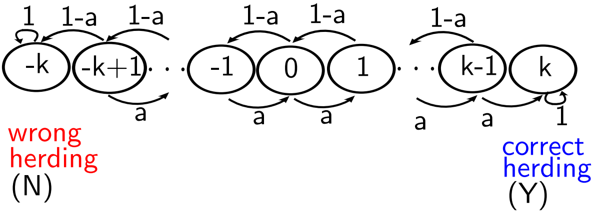

- Can exactly calculate expected payoff $E[\pi_i]$ & probability of wrong (correct) herding for any agent i
  - $E[\pi_i]$ (MC with rewards)

Markov Chain viewpoint (continued)

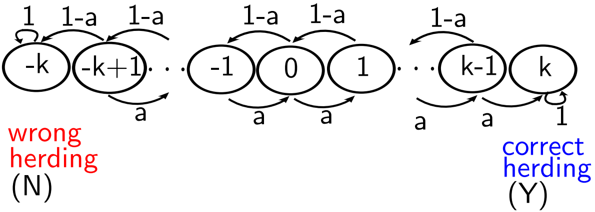

- Can exactly calculate expected payoff $E[\pi_i]$ & probability of wrong (correct) herding for any agent i
  - $E[\pi_i]$ (MC with rewards)
  - $\mathbb{P}[wrong_{i-1}] = \sum_{n=1}^{i-1} \mathbb{P}[agent\ n\ is\ the\ first\ to\ hit\ -k]$

Markov Chain viewpoint (continued)

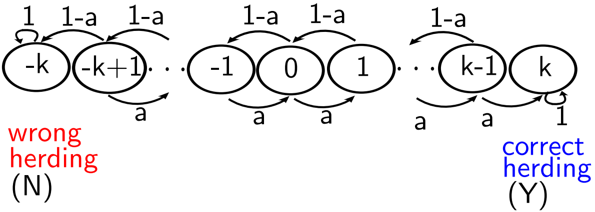

- Can exactly calculate expected payoff $E[\pi_i]$ & probability of wrong (correct) herding for any agent i
  - $E[\pi_i]$ (MC with rewards)
  - $\mathbb{P}[wrong_{i-1}] = \sum_{n=1}^{i-1} \mathbb{P}[agent\ n\ is\ the\ first\ to\ hit\ -k]$
  - $\mathbb{P}[correct_{i-1}] = \sum_{n=1}^{i-1} \mathbb{P}[agent\ n\ is\ the\ first\ to\ hit\ k]$

Markov Chain viewpoint (continued)

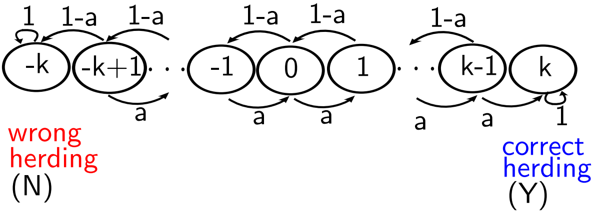

- Can exactly calculate expected payoff $E[\pi_i]$ & probability of wrong (correct) herding for any agent i
  - $E[\pi_i]$ (MC with rewards)
  - $\mathbb{P}[wrong_{i-1}] = \sum_{n=1}^{i-1} \mathbb{P}[agent\ n\ is\ the\ first\ to\ hit\ -k]$
  - $\mathbb{P}[correct_{i-1}] = \sum_{n=1}^{i-1} \mathbb{P}[agent\ n\ is\ the\ first\ to\ hit\ k]$
  - First-time hitting probabilities: Use probability generating function method [Feller'68]

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$

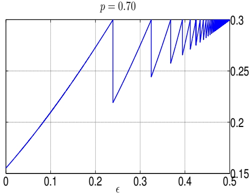

Limiting wrong herding probability

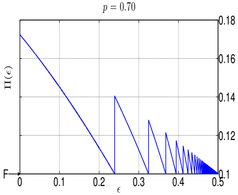

Limiting payoff $\Pi(\epsilon) = \lim_{i \to \infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed

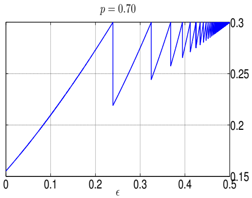

Limiting wrong herding probability

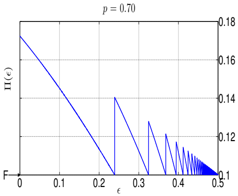

Limiting payoff $\Pi(\epsilon) = \lim_{i \to \infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \leq \epsilon < \epsilon^*(k+1,p)$

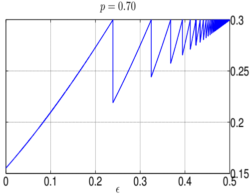

Limiting wrong herding probability

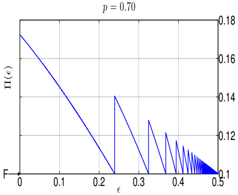

Limiting payoff $\Pi(\epsilon) = \lim_{i\to\infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \leq \epsilon < \epsilon^*(k+1,p)$
    - Probability of wrong herding increases

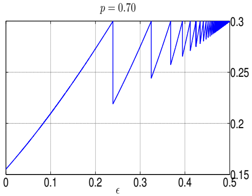

Limiting wrong herding probability

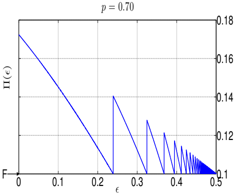

Limiting payoff $\Pi(\epsilon) = \lim_{i \to \infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$
    - Probability of wrong herding increases
    - $\Pi(\epsilon)$ decreases to F

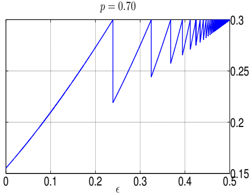

Limiting wrong herding probability

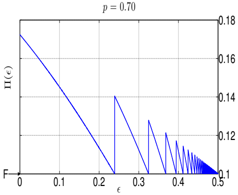

Limiting payoff $\Pi(\epsilon) = \lim_{i \to \infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \leq \epsilon < \epsilon^*(k+1,p)$
    - Probability of wrong herding increases
    - $\Pi(\epsilon)$ decreases to F
  - Probability of wrong herding jumps when k changes

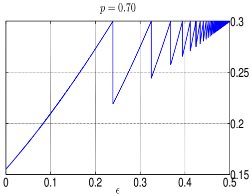

Limiting wrong herding probability

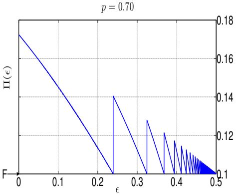

Limiting payoff $\Pi(\epsilon) = \lim_{i \to \infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \leq \epsilon < \epsilon^*(k+1,p)$
    - Probability of wrong herding increases
    - $\Pi(\epsilon)$ decreases to F
  - Probability of wrong herding jumps when k changes
  - Limiting payoff also jumps at same point $F = \Pi(\epsilon^*(k+1,p)^-) < \Pi(\epsilon^*(k+1,p)^+)$

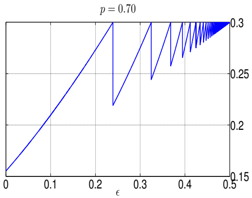

Limiting wrong herding probability

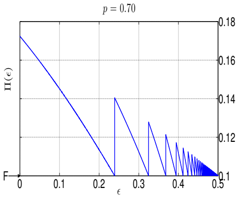

Limiting payoff $\Pi(\epsilon) = \lim_{i\to\infty} E[\pi_i]$

Results

- Payoff for agents is non-decreasing in i & at least $F = \frac{2p-1}{4} > 0$
- Limiting payoff $\Pi(\epsilon)$ & probability of wrong herding can be analyzed
  - For $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$
    - Probability of wrong herding increases
    - $\Pi(\epsilon)$ decreases to F
  - Probability of wrong herding jumps when k changes
  - Limiting payoff also jumps at same point $F = \Pi(\epsilon^*(k+1,p)^-) < \Pi(\epsilon^*(k+1,p)^+)$
- There exists a range where increasing noise improves performance!!!

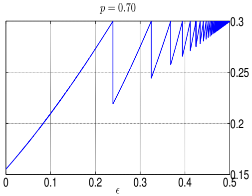

Limiting wrong herding probability

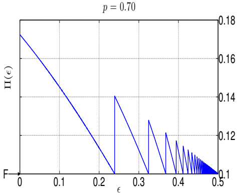

Limiting payoff $\Pi(\epsilon) = \lim_{i\to\infty} E[\pi_i]$

Results for an arbitrary agent i

Similar ordering holds for every user's payoff & probability of wrong herding

- Discontinuities and jumps at the same thresholds
- For $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$: $E[\pi_i]$ decreases in ϵ

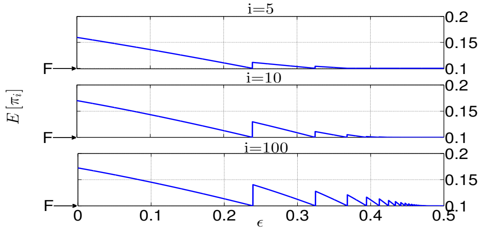

Individual payoff for signal quality p=0.70

Results for an arbitrary agent i

Similar ordering holds for every user's payoff & probability of wrong herding

- Discontinuities and jumps at the same thresholds
- For $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$: $E[\pi_i]$ decreases in ϵ
  - Proof using stochastic ordering of Markov Chains & coupling

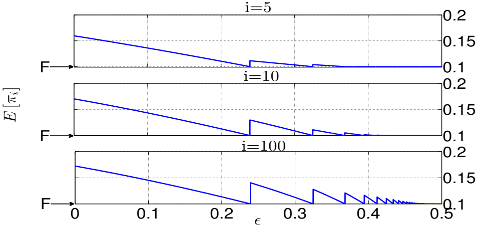

Individual payoff for signal quality p=0.70

Results for an arbitrary agent i

Similar ordering holds for every user's payoff & probability of wrong herding

- Discontinuities and jumps at the same thresholds
- For $\epsilon^*(k,p) \le \epsilon < \epsilon^*(k+1,p)$: $E[\pi_i]$ decreases in ϵ
    - Proof using stochastic ordering of Markov Chains & coupling

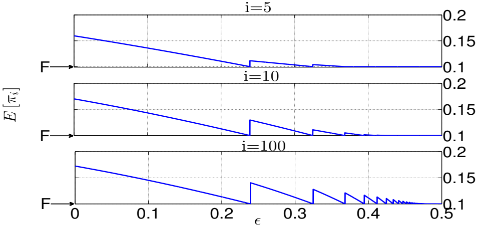

Individual payoff for signal quality p=0.70

- For given level of noise, adding more noise may not improve all agents pay-offs.

Extension: Quasi-rational agents

- Real-world agents not always rational

Extension: Quasi-rational agents

- Real-world agents not always rational
- One simple model: agents make "action errors" with some probability ϵ_1

Extension: Quasi-rational agents

- Real-world agents not always rational
- One simple model: agents make "action errors" with some probability ϵ_1
  - e.g., noisy best response, trembling hand, inconsistency in preferences

Extension: Quasi-rational agents

- Real-world agents not always rational
- One simple model: agents make "action errors" with some probability ϵ_1
  - e.g., noisy best response, trembling hand, inconsistency in preferences
- How to account for this (assuming ϵ_1 is known)?

Extension: Quasi-rational agents

- Real-world agents not always rational
- One simple model: agents make "action errors" with some probability ϵ_1
  - e.g., noisy best response, trembling hand, inconsistency in preferences
- How to account for this (assuming ϵ_1 is known)?
  - Nothing really new from view of other agents
  - But pay-off calculation changes

Results: Quasi-rational agents and Noise

- Consider three "errors"

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong
  - $\epsilon_3 \in (0, 0.5)$: probability social planner flips the action record

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong
  - $\epsilon_3 \in (0, 0.5)$: probability social planner flips the action record

  - Similar result as before: equivalent total noise ϵ used

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong
  - $\epsilon_3 \in (0, 0.5)$: probability social planner flips the action record

  - Similar result as before: equivalent total noise ϵ used
  - Each user's payoff is reduced by a factor $(1 - 2\epsilon_1)$

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong
  - $\epsilon_3 \in (0, 0.5)$: probability social planner flips the action record

  - Similar result as before: equivalent total noise ϵ used
  - Each user's payoff is reduced by a factor $(1 - 2\epsilon_1)$
  - There exist cases where adding more observation noise (ϵ_3) always increases limiting payoff (even if $\epsilon_2 = 0$)

Results: Quasi-rational agents and Noise

- Consider three "errors"
  - $\epsilon_1 \in (0, 0.5)$: probability agents choose sub-optimal action
  - $\epsilon_2 \in (0, 0.5)$: probability actions are recorded wrong
  - $\epsilon_3 \in (0, 0.5)$: probability social planner flips the action record

  - Similar result as before: equivalent total noise ϵ used
  - Each user's payoff is reduced by a factor $(1 - 2\epsilon_1)$
  - There exist cases where adding more observation noise (ϵ_3) always increases limiting payoff (even if $\epsilon_2 = 0$)

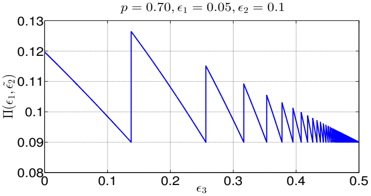

Limiting payoff, $p = 0.70$

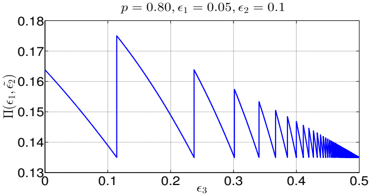

Limiting payoff, $p = 0.80$

Conclusions

- Analyzed simple Bayesian learning model with noise for herding behavior

Conclusions

- Analyzed simple Bayesian learning model with noise for herding behavior
- Noise thresholds determine the onset of herding
  - For $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$, require $|\#Y's - \#N's| \geq k$ to trigger herding.
  - Generalized BHW'92: $k = 2$ for noiseless model

Conclusions

- Analyzed simple Bayesian learning model with noise for herding behavior
- Noise thresholds determine the onset of herding
  - For $\epsilon^*(k, p) \leq \epsilon < \epsilon^*(k+1, p)$, require $|\#Y's - \#N's| \geq k$ to trigger herding.
  - Generalized BHW'92: $k = 2$ for noiseless model
- With noisy observations, sometimes it is better to increase the noise

Conclusions

- Analyzed simple Bayesian learning model with noise for herding behavior
- Noise thresholds determine the onset of herding
  - For $\epsilon^*(k,p) \leq \epsilon < \epsilon^*(k+1,p)$, require $|\#Y's - \#N's| \geq k$ to trigger herding.
  - Generalized BHW'92: $k = 2$ for noiseless model
- With noisy observations, sometimes it is better to increase the noise
  - Probability of wrong herding decreases
  - Asymptotic individual expected welfare increases
  - Average social welfare increases

Future directions

- Heterogeneous private signal qualities and noises

Future directions

- Heterogeneous private signal qualities and noises
- Possibility of more actions, richer responses
  - Combination with Sgroi'02 (guinea pigs)
    Force M initial agents to use private signals
  - Investment in private signal when facing high wrong herding probability

Future directions

- Heterogeneous private signal qualities and noises
- Possibility of more actions, richer responses
  - Combination with Sgroi'02 (guinea pigs)
    Force M initial agents to use private signals
  - Investment in private signal when facing high wrong herding probability
- Different network structures

Future directions

- Heterogeneous private signal qualities and noises
- Possibility of more actions, richer responses
  - Combination with Sgroi'02 (guinea pigs)
    Force M initial agents to use private signals
  - Investment in private signal when facing high wrong herding probability
- Different network structures
- Strategic agents in endogenous time

Future directions

- Heterogeneous private signal qualities and noises
- Possibility of more actions, richer responses
  - Combination with Sgroi'02 (guinea pigs)
    Force M initial agents to use private signals
  - Investment in private signal when facing high wrong herding probability
- Different network structures
- Strategic agents in endogenous time
- Achieve learning with agents incentivized to participate

References

- W. Feller, *An introduction to Probability Theory and Its Applications*, vol. I, 3rd ed., New York Wiley, 1968.
- T. M. Cover, *Hypothesis Testing with Finite Statistics*, Ann. Math. Stat., 40(3):828–835, June 1969.
- M. E. Hellman and T. M. Cover, *Learning with Finite Memory*, Ann. Math. Stat., 41(3):765–782, June 1970.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. *A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades*, J. Polit. Econ., vol. 100, No. 5, pp. 992-1026, 1992.
- D. Sgroi, *Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare*, Games and Economic Behaviour, vol. 39, pp. 137-166, 2002.
- L. Smith, P. Sorensen, *Pathological Outcomes of Observational Learning*, Econometrica, vol. 68, pp. 371-398, 2000.
- D. Easley, J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*, Cambridge University Press, 2010.
- D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Review of Economic Studies, vol. 78, pp. 1201-1236, 2011.
- Z. Zhang, E. K. P. Chong, A. Pezeshki, and W. Moran, *Hypothesis testing in feedforward networks with broadcast failures*, IEEE Journal of Selected Topics in Signal Processing, special issue on Learning-Based Decision Making in Dynamic Systems under Uncertainty, vol. 7, no. 5, pp. 797–810, October 2013.
- T. Le, V. Subramanian, R. Berry, *The Value of Noises for Informational Cascades, ISIT 2014*.
- T. Le, V. Subramanian, R. Berry, *The Impact of Observation and Action Errors on Informational Cascades*, to appear *CDC 2014*.

Thank you!