Some Issues in Individual and Social Belief Change

Matthew Rabin

BREAD, "@ Stanford"

January 29, 2021 AD
Day 9 AT (AI)

- Thank you for inviting me.
  - Weather here in Stanford: Low of $9^{\odot}$ (Fahrenheit), wind 20 mph.
  - In foreign measurements: Butt-ass cold (Celsius),and like 700 kph.
- Plan: Hang out with a bunch of cool development economists in CA
- Reality:Stuck "with" only a few cool development economists in MA

↬

Some Issues

Today: "Overviewish" talk shamelessly (mostly) my papers.

- Limited rationality in belief updating
    - Trying to tie them together, express some themes:
    - All these papers build on Kahneman, Tversky et. al

Beliefs on a fixed question

- Title: Belief *change*, not "learning"

Learning has some connotations that don't capture what happens:

1. That beliefs change (often don't when should)
2. That change directionally (often bounce around excessively)
3. That in rich environments, reach certainty(often stop short)
4. If change directionally and towards certainty ...
    1. it is towards the truth(often, not so much)

Going to think about these issues in individual updating ... "non-volitional" data.And in social-inference context.

Some Issues

Individual Updating:

- Biased Beliefs about Sampling:
  - ⌈Inference by LSNers (2002), LSN & HH (2010), with Dimitri Vayanos
  - | NBLLN (2016), with Dan Benjamin and Collin Raymond
  - ⌊Beliefs re Random Samples (2017), experiment, with Dan Benjamin and Don Moore
- Misguided updating:
  - ⌈A Model of Confirmatory Bias (1999) (with Joel Schrag)
  - ⌊Base-Rate Neglect (2019), with DB and Aaron Bodoh-Creed
- Identifying non-Bayesian updating:
  - ⌈Belief Movement, Uncertainty Reduction, & Updating (forthcoming) (w/ Ned Augenblick)
  - ⌊Restrictions Asset-Price Movement (2018), Augenblick and Eben Lazarus (no Rabin)
- See Dan Benjamin chapter on judgmental bias (2019)

Some Issues

Social Context:

- "Cursed" failure to extract information from others
  - ⌈Cursed Equilibrium (2005) with Erik Eyster
  - | Neglect of Info in Asset Prices (2019) with Eyster and Dimitri Vayanos
  - ⌊Disagreement Doesn't Induce Trade (in progress) with Eyster and Tristan G-B
- But neglecting redundancy when do extract information
  - ⌈Naive Herding in Rich-Info Settings (2010) with Erik Eyster
  - | Extensive Imitation Irrational and Harmful (2014) with Eyster
  - | Experiment on Social Mislearning w/Eyster and Georg Weizsacker
  - ⌊Social Mislearning and Unlearning (2016) with Tristan Gagnon-Bartsch
- Correct learning is hard.
  - ⌈Taste Projection in Social Learning by Tristan Gagnon-Bartsch
  - | Selective Attention and Learning (2014) by Josh Schwartzstein
  - ⌊Channeled Attention and Stable Errors (2018) with TG-B and JS

Some Issues

Some general themes from the psychology of belief formation

- Psychology of error is often not about complexity
  - People neglect some simple insights (logic of markets, herds)
  - People have bad intuitions that lead them astray
  - *Optimal* attention would solve all at low cost
  - Often implausible that real costs of attention to well-understood costs of inattention.
  - E.g., stock market. seems (obvious) to me that that irrational attention is much bigger (trillions $ worth) problem than rational inattention.
  - Gambler's fallacy requires more attention than the right answer.

↬

Some Issues

- Those entertaining non-optimality ...
- Not just overinterpretation, over-inference, and overconfidence
  - Also under-interpretation, underinference, and underconfidence
  - **And**: False consciousness of attributing errors to "overconfidence":
- Overconfidence often a restatement of fact that you made an error.
  - If losing money, whatever it is you're doing, you're overconfident it is the right strategy.
  - Trade in market
- When people use factor A more than B, we shouldn't assume overweighting A.
  - It could be underweighting B. Consumer reports vs. friend?
  - And that distinction matters.
  - In updating: Overweighting most recent info, or downweighting earlier?
  - Makes a massive difference.

Cognitive Biases

Stylized versions of individual (and social) biases:
Updating: p_t priors, σ_t new information, and p_{t+1} posteriors.

- Bayesian: $\ln(\frac{p_{t+1}}{1-p_{t+1}}) = \ln(\frac{p_t}{1-p_t}) + \ln(\frac{\sigma_t}{1-\sigma_t})$
- Non-Bayesian: $\ln(\frac{p_{t+1}}{1-p_{t+1}}) = \alpha \ln(\frac{p_t}{1-p_t}) + \beta \ln(\frac{\sigma_t}{1-\sigma_t})$
- $\alpha < 1$ is base-rate neglect (Saki)
- $\beta > 1$ when σ_t is small sample is LSN (Freddy)
- $\beta < 1$ when σ_t is large sample is NBLLN (Barney)
- And confirmatory bias (Joely) is ... $\alpha > 1$ (?)
  - when σ_t is ambiguous?
  - See AR (forthcoming) for justification, and more natural variant.

Now: Two features of belief change:
$\looparrowright$

Cognitive Biases

- For binary events, look at changes in beliefs over time.

Define (expected) belief **movement**:

- $m_t \equiv E\left\{\sum(p_{t+1} - p_t)^2\right\}$

Define (expected) **uncertainty reduction**:

- $r_t \equiv E\left\{\sum[p_t(1 - p_t) - p_{t+1}(1 - p_{t+1})]\right\}$.

Let's talk about these for a while.

- (Conspicuously missing for a while ... whether beliefs are correct.)

Cognitive Biases

What only recently I've seen the relevance of these per se:

- Scattered in papers on quasi-Bayesian models ...
- Trying to characterize the effects of biases
  - Main object has been: long-run overconfidence or underconfidence
  - "Socratic" or "Bernoullian" failure.
- m and r are objects of direct interest ...
  - Are people's beliefs due to bias too volatile? too docile?
  - Do they figure things out too slow? "too fast?" (and hence wrongly)

Claim: We can associate common errors with patterns of how movement and uncertainty resolution compare to Bayesian:

$\looparrowright$

Movement and Uncertainty Reduction

Movement vs. Uncertainty Resolution (Claims not quite true)

-
  - See AR (forthcoming) for simple variants that *are* true:
  - Comparing these errors to Bayesian (i.e, "Tommy Bayes"):

	movement	resolution
Tommy (Bayes)	$m^T = m^T$	$r^T = r^T$
Saki (BRN)	$m^S > m^T$	$r^S < r^T$
Freddy (LSN)	$m^F > m^T$	$r^F > r^T$
Barney (NBLLN)	$m^B < m^T$	$r^B < r^T$
Joely (Confirm, herd)	$m^J < m^T$	$r^J > r^T$

Note: most of research in this area is like this chart ...

- Compare error to Bayesian when we know right answer.
- **Aside**: classical examples of BRN (inferring too much false positive on medical test) misleading ...
  - People too slow to update on average.
- AR (forthcoming) different take.
  - If **don't know** m^T, r^T, are there markers of rationality?

Movement and Uncertainty Reduction

Harken back to 2012 ... how innocent we were ... what passed for bad back then? Remember when Rick Perry was stupid?When Newt Gingrich was an extremist blowhard? When Rick Santorum had ugly social views? When Michelle Bachmann was a bit of a nutter?And when (before loyalty demanded self-harm and other-harm) Herman Cain was alive?

- Each had their day in the betting-market sun.

Intrade beliefs that Romney would win the 2012 Republican nomination:

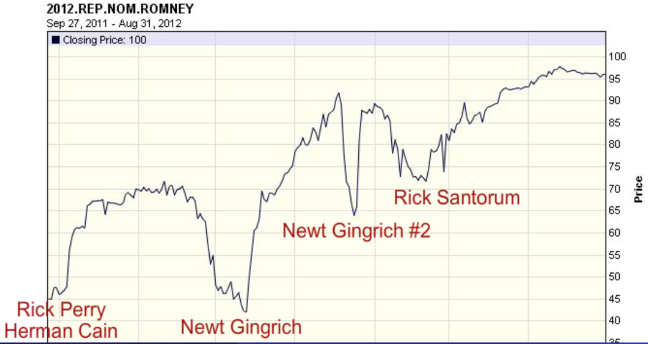

Movement and Uncertainty Reduction

- People overreacting to information?
- Perhaps not.
  - Perhaps what happened was truly surprising?
  - We'll have nothing to say distinguishing unlikely-to-be-rational beliefs vs. truly unlikely events.
  - But we can identify unlikely belief movement, whatever its source.
- Several interrelated angles on project
  - Assess rationality of updating without knowing right answer?
- How do we assess markers of irrationality without comparing to right answer?We can compare
  - not to known probabilities
  - not to market fundamentals
  - not to Nate Silver

Movement and Uncertainty Reduction

- If too much volatility happens over and over,
  - might start to question "Bayesianess."
  - [Bayesian = rational expectations given information]
- Fundamental principle of rational belief updating,
  - Independent of strong views of normative beliefs ...
  - If you are changing your mind a lot, you had better (on average) be resolving uncertainty.
  - A surprising and simple equality captures relationship.

Movement and Uncertainty Reduction

Reminder: For binary events, look at changes in beliefs over time.

(expected) belief **movement**:

- $m_t \equiv E\left\{\sum(p_{t+1} - p_t)^2\right\}$

(expected) **uncertainty reduction**:

- $r_t \equiv E\left\{\sum[p_t(1-p_t) - p_{t+1}(1-p_{t+1})]\right\}$.

Lemma:**All Bayesian processes** beginning from correct priors:

- $m_t = r_t$

In fact, **all** martingales have this property.

- And can look for this to see whether anything is a martingale.
- Except how do you know if seen too much movement?
  - **Since uncertainty bounded, expected movement bounded.**
  - So we can look at the world and see if m is too big.
  - And note: This has nothing to do with duration of belief stream.
  - (AR: Boundedness as rendered by finite variance (not finite range)

Movement and Uncertainty Reduction

Key intuition on rationality of updating:

- If beliefs changing, better be learning something:
- If a lot of movement, better be resolution of uncertainty.
  - Agnostic whether updating should be fast or slow

When "resolving streams"

- $E[\sum_{t=1}^{\infty}(\pi_t - \pi_{t-1})^2] = \pi_0(1 - \pi_0) = r_0$.
  - So even single stream can be too long.
  - In principle, watching just one stream could reject greater than .05 likelihood of Bayesian having $m > 5$.
  - (In fact, $m > 2.16$ too ...)
- But in practice, very unlikely to observe such wild movement.
  - Need big data set.
  - And definitely need big data set (and more assumptions) to determine $m < r$.

There are other such equivalencies ... with different definitions of movement and uncertainty resolution.

- Corresponding to different measures in information theory.
- This is probably the simplest and most familiar
- But many metrics of movement cannot be characterized as excessive.
  - Recall, e.g., that $E[\sum_{t=1}^{\infty} |\pi_t - \pi_{t-1}|]$ can be infinite!

But more than the equivalence, noting that $\pi_0(1 - \pi_0)$ is (very) finite:

- $E[\sum_{t=1}^{\infty} (\pi_t - \pi_{t-1})^2]$ must be (very) finite.

Movement and Uncertainty Reduction

- Ultimately tests we propose are variants of martingale tests ...
  - calibration, autocorrelation tests, etc. others
- **But**
  - But comparing movement vs. u-resolution has nice properties.
  - AR compare in paper to other tests, show advantages (far fewer false rejections, yet high-powered)
  - beliefs are **bounded martingales ...** testing for excess movement
  - both in test design and effectiveness, test for excess movement
  - No analog with unbounded martingales.

Back to the biases: Movement vs. Uncertainty Resolution:

- (Still not quite true)

	m vs. r	move	resolve
Tommy	$m^T = r^T$	$m^T = m^T$	$r^T = r^T$
Saki	$m^S > r^S$	$m^S > m^T$	$r^S < r^T$
Freddy	$m^F > r^F$	$m^F > m^T$	$r^F > r^T$
Barney	$m^B < r^B$	$m^B < m^T$	$r^B < r^T$
Joely	$m^J < r^J$	$m^J < m^T$	$r^J > r^T$

$\looparrowright$

Movement and Uncertainty Reduction

- But when don't know for sure the right models.
- Not test of whether beliefs **correct:** test of **consistency** of beliefs
  - You have to be as surprised as you thought you were going to be.
  - But if eventually truth revealed, consistency closer to correctness.
  - Also reject when correct prior
  - The test will not reject all non-Bayesian behavior
- But switching to social learning, can also apply this.
  - It's right that nutrionists and medicine should be updating beliefs ...
    - But only so many times eggs, coconut oil, etc. can be glorified/vilified.
  - In healthy sciences, views update based on evidence.

Errors in Social Inference

Others' behavior might reveal information to us.

- How good gleaning this information?
- What systematic errors?
- Effects of these errors?

Two facts about your interactions with others:

1. Their behavior often contains information relevant to you which you do not glean in other ways.
2. Information in behavior of others will likely be correlated
   - Because they are also gleaning information from each other.

In series of papers with Erik Eyster, we propose:

$\looparrowright$

Errors in Social Inference

"Cursed Thinking" (not today):

- People under-infer information from others' behavior.
  - Winner's **curse** in common-values auctions:
  - Lemons and **financial markets**
  - Excess trade—naturally and directly

Alternative approaches?

- Overconfidence
  - People think own signals are better than they are.
  - (And maybe that others' signals worse)
- **Agreeing to disagree, non-common priors**
  - ATD directly, or to "close" overconfidence models.
  - Traders aware of disagreements; seek bets based on them!
  - Often (despite massive degrees of freedom) not plausible.

Errors in Social Inference

Eyster, Rabin, and Vayanos (2015):

- Adapt CE to asset-market model.
- Cursed traders don't think through information in market prices.
  - Don't ask: why is that person trading with me?
  - Retro! ... pre-Grossman REE revolution.
  - Old school not 100% wrong
- Some traders fully rational, some partially cursed.

↬

Errors in Social Inference

ATD driving force of small investors trading against the market?

- **Do they think they are outsmarting experts and insiders?**
  - **We say: not attending to fact they are trading against them**

ERV shows that, even if true, overconfidence/ATD unlikely to explain high volume in large, "high-information" markets.Roughly:

- As long as trader believes lots of information besides her own insights, then just like REE logic $\Longrightarrow$ little speculative trade in large markets.
- Fundamental, ubiquitous, necessary, very costly error investors make:
  - underappreciating info in market prices.

But also (per work with TG-B) can think of all the other domains:

- Are consumers thinking through that retailing and insurance companies probably know bottom line ... you must on average be losing actuarilly + transaction costs.
- What error is it to believe snake-oil salesman has cure for cancer, etc.?
    - Biological? Like thinking bark can alieve pain, or mold can save millions?(e.g., solve 500-year pandemic virtually overnight?)(Syphillis, penecillin, if spelled correctly.)
    - You should think: How would that thing work if doc didn't tell me, and so many dying of cancer?

↬

Errors in Social Inference

Now: Redundancy Neglect/Naive Inference:

- Insofar as *do* attend to information in others' behavior, insufficient attention to *redundancy* in social beliefs.

In observational learning:

- Rationality predicts some imitation ...
  - but *precludes* extensive imitation
  - *typically, intrinsicly, virtually-assumption-freely* predicts anti-imitation too
- Redundancy neglect/naive inferences $\Longrightarrow$ ubiquitous imitation
- Big welfare cost: overconfidently wrong social beliefs

Rational and Naive Observational Learning

1. **Behavioral** (as in behavior, not mugs) implications of rational observational learning
2. Informational, societal consequences of naive herding

↪

Inferential Naivete in Observational Learning

- People may extract information (not fully cursed) from others.
- But take this information "at face value"
- Portable and pinned down universal definition
  - But now talk solely about observational learning.
- These literatures have branched (sequential learning vs. networks repeat interaction).
- I think different in lessons little to do with any difference except
  - The models that entertain basic irrationality, Degroot, seems close to naive learning.
  - But a subtle difference is huge.
  - And when you don't look at the fundamental error—redundancy neglect—you miss how network structures compare in terms of learning **bad** ideas.

Errors in Social Inference

Rational-Herding Literature:

- People infer from actions of those with similar tastes.
- Rational imitation.
- Herds may start & last on wrong choice.

Universal feature not fully appreciated #1:

- Limited inefficiency: merely non-aggregation; delusional herds can't be likely.

Universal feature not fully appreciated #2:

- Because all rationally realize others also imitating
- hence realize **inherent** redundancy in others' behavior
- Don't imitate very much.

↬

Errors in Social Inference

We claim:

- Canonical example, connotation of literature misleading.
- **Limits** to imitation far bigger punchline than the imitation itself.
- We think the non-imitation is unrealistic.
- **And we should care a lot about over-imitation:**
  - Theories that generate extensive imitation allow possibility of severe badnesses in societal beliefs ... real inefficiencies
- Extensive imitation $\implies$ not rational
- Extensive imitation $\implies$ social confirmation bias & false beliefs.

Now:

- **Behavioral implications** of rational observational learning, once move outside limited environment.
  - Also indicates likely effects redundancy neglect

Errors in Social Inference

Modification of the standard two-restaurant model of social learning.

- Two restaurants in town,
  - A and B, $p(A\ good, B\ bad) = p(B\ good, A\ bad) = .5$.
  - Two states: $\omega_A \rightarrow A$ is good, $\omega_B \rightarrow B$ is good.
  - Binary-state model universal.... and weird.

$\looparrowright$

Errors in Social Inference

- Each of ∞ diners receives private signals $\in \{\alpha, \beta, \emptyset\}$
- The signals are *i.i.d.* conditional on the state,
  - α supports ω_A,
  - β supports ω_B,
  - $\emptyset$ uninformative.
- For each Player k,
  - $\Pr[s_k = \alpha|\omega_A] = \Pr[s_k = \beta|\omega_B] = .7(1-\eta)$ and
  - $\Pr[\emptyset|\omega_A] = \Pr[\emptyset|\omega_B] = \eta$.
  - $\eta = 0$, canonical binary-signal information structure.
  - When $\eta \to 1$, information is very rare.
  - (Lots results independent of η)

$\hookrightarrow$

Errors in Social Inference

- Each Player k chooses among nine choices:
  - dine at A, dine at B, or dine at home.
  - Goes to restaurant if thinks more than 60% chance it is good;
  - stays at home if that is not true at either restaurant.
- Depending on confidence in restaurant's quality, may go alone, or take one, two, or three of her relatives.
- Superscripts for the number of people she takes:

$p(\omega_A)$	[0,10),[10,20),[20,30),[30,40)	[40,60]	(60,70],(70,80],(80,90],(90,100]
Choice	$B^{+++}, B^{++}, B^{+}, B$	H	$A, A^{+}, A^{++}, A^{+++}$

Three people choose restaurants each period,

- Signal conditionally i.i.d. given state
- Each doing so after observing her own signal,and the full actions (three locations, and party size), in order,taken in all previous periods.

Specific example clearly contrived.

- But logic very general, punchlines not based on specifics.
- Based on basic principle of statistical inference impossible to escape.
- (By contrast, canonical examples **are** misleading.)

So instead of single file, binary action/signals, multi-file rich action/signals $\looparrowright$

What predictions does full rationality make?

- $\emptyset$ signal, observes nothing but $H \longrightarrow$ stay home.
- α or β signal, observes nothing but $H \longrightarrow$ go to restaurant.
- (alone, because beliefs exactly .7 $\longrightarrow$ alone).

Suppose in period 2 observe that exactly one person has gone to Restaurant A in period 1.

- What do as a function of your signal?
- You will realize that the three signals in period 1 were $\{\alpha, \emptyset, \emptyset\}$.
  - $\beta \longrightarrow H$.
  - $\emptyset \longrightarrow A$.
  - $\alpha \longrightarrow A^{++}$

↬

If observe:

	actions
Period 1:	$\{A, H, H\}$
Period 2:	$\{A, A, A\}$

What do (as function of signal)?

$\looparrowright$

Errors in Social Inference

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{A, A, A\}$		$\{\emptyset, \emptyset, \emptyset\}$

Period 3: $\beta \rightarrow H$, $\emptyset \rightarrow A$, $\alpha \rightarrow A^{++}$

- Key logic: guys in period 2 did *not* get any additional information.
    - (If did, would not have gone alone.)
    - Period 3: rationally realize no new information in Period-2 followers.

↬

Errors in Social Inference

If observe:

	actions
Period 1:	$\{A, H, H\}$
Period 2:	$\{A, A, A\}$
Period 3:	$\{A, A, A\}$
Period 4:	$\{A, A, A\}$
Period 5:	$\{A, A, A\}$

What do as (as function of signal)?

$\looparrowright$

Errors in Social Inference

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{A, A, A\}$		$\{\emptyset, \emptyset, \emptyset\}$
Period 3:	$\{A, A, A\}$		$\{\emptyset, \emptyset, \emptyset\}$
Period 4:	$\{A, A, A\}$		$\{\emptyset, \emptyset, \emptyset\}$
Period 5:	$\{A, A, A\}$		$\{\emptyset, \emptyset, \emptyset\}$

Period 6: $\beta \rightarrow H,\ \emptyset \rightarrow A,\ \alpha \rightarrow A^{++}$

Errors in Social Inference

If observe:

	actions
Period 1:	$\{A, H, H\}$
Period 2:	$\{A, A, H\}$

What do (as function of signal)?

$\looparrowright$

Errors in Social Inference

Herding without sufficiently increased enthusiasm is a bad sign:

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{A, A, H\}$		$\{\emptyset, \emptyset, \beta\}$
Period 3:		$\beta \to B$, $\emptyset \to H$, $\alpha \to A$	

3 A, 3 H $\to$ ω_A, ω_B equally likely!

- Do we get that?

↬

If observe:

	actions
Period 1:	$\{A, H, H\}$
Period 2:	$\{A, H, H\}$

What do (as function of signal)?

$\looparrowright$

Errors in Social Inference

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{A, H, H\}$		$\{\emptyset, \beta, \beta\}$
Period 3:		$\beta \rightarrow B^{++}, \emptyset \rightarrow B, \alpha \rightarrow H$	

You shouldn't go to A even if get α!

Errors in Social Inference

If observe:

	actions
Period 1:	$\{A, H, H\}$
Period 2:	$\{A^{++}, A, A\}$
Period 3:	$\{A^{++}, A^{++}, A\}$
Period 4:	$\{A^{++}, A^{++}, A\}$
Period 5:	$\{A^{++}, A^{++}, A^{++}\}$

What do?

$\looparrowright$

Errors in Social Inference

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{A^{++}, A, A\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 3:	$\{A^{++}, A^{++}, A\}$		$\{\emptyset, \emptyset, \beta\}$
Period 4:	$\{A^{++}, A^{++}, A\}$		$\{\alpha, \alpha, \emptyset\}$
Period 5:	$\{A^{++}, A^{++}, A^{++}\}$		$\{\beta, \beta, \beta\}$
Period 6:		$\beta \rightarrow B$, $\emptyset \rightarrow H$, $\alpha \rightarrow A$	

Will a β signal really stop the herd?

$\looparrowright$

Errors in Social Inference

- Enough.
- Things far more complicated if don't observe order, don't observe all, or heterogenous preferences

But nothing makes the severe limits to imitation go away

- **Others' beliefs massively correlated**
  - $\Longrightarrow$ **musn't imitate too much.**

$\looparrowright$

E.g., same setting (same signals, players per period, etc.) but:

- Cannot observe order of play.
- **Signals rare**
- In period 3, if see
  - If see $\{H, H, H, H, H, H\}$, then believe .5
  - If see $\{A, H, H, H, H, H\}$, then believe .7
  - If see $\{A, A, H, H, H, H\}$, then believe .84
  - If see $\{A, A, A, H, H, H\}$, then believe .5
  - If see $\{A, A, A, A, H, H\}$, then believe .7
  - If see $\{A, A, A, A, A, H\}$, then believe .3

Another framing using one old and symmetric new example:
↬

	actions	response	signals
Period 1:	$\{A, H, H\}$		$\{\alpha, \emptyset, \emptyset\}$
Period 2:	$\{H, H, H\}$		$\{\beta, \beta, \beta\}$
Period 3:		$\beta \to B^{+++}, \emptyset \to B^{++}, \alpha \to B$	

	actions	response	signals
Period 1:	$\{B, H, H\}$		$\{\beta, \emptyset, \emptyset\}$
Period 2:	$\{H, H, H\}$		$\{\alpha, \alpha, \alpha\}$
Period 3:		$\beta \to A, \emptyset \to A^{++}, \alpha \to A^{+++}$	

Anti-imitation!

Errors in Social Inference

Lesson:

- Most natural learning environments...
    - Great deal redundancy.
- **Almost all non-single-file environments, rationality implies 'anti-imitation'.**
    - Intuition: if two recent guys both imitating earlier guy but not each other—imitate both,
    - **but subtract off source of correlation... earlier guy**.
- What few environments that don't demand 'anti-imitation' (e.g., single-file),imitate all you see as if seeing only one person.
- The big behavioral lesson of rational herding ought to be how **little** and how **weird** imitation it predicts, not that it predicts imitation.
    - And big welfare efficiency lesson is how hard to find rich-information environment where learning isn't close to complete(and even when can, the inefficiency is non-aggregation, not wildly bad beliefs.
- If instead people are partly naive: a chance of extreme, wrong beliefs.

Errors in Social Inference

Harder to see:

- Rational observational learning in this case:
  - Eventually will herd on $\{B^{+++}\}$ or $\{A^{+++}\}$.
  - More than 95% of time $\rightarrow$ right restaurant.
  - Intuition: any lesser certainty, contrary signal will moderate behavior.
- **When signals rare:**
  - 5% of time herd starts and stays in wrong direction.
  - 25% of time: herd in wrong direction followed by reversal...
  - somebody observing at least 50 people going to one restaurant and none to other decides stay home based on opposite signal.

Briefly ... What happens if instead people neglect redundancy?

- In all those examples, easy to see: Once people start going to A gets interpreted as lots of evidence, herd will be off to the races.

Errors in Social Inference

- ER (2010) extreme form, "Britney", behaves as if all she is observing are independent.

When signals are rare,

- 30% of time Britney herds to wrong restaurant, and they become ∼100% sure they are right.
- Eyster-Rabin (2014) ... When individual strong signals rare, tiny amounts of naive redundancy neglect can lead society wildly astray with high probability.
  - E.g., blood-letting and medicinal mercury can last for centuries as common "cures".
- TG-B and Rabin: Beliefs can often become way too extreme.
  - If restaurants can be both good or both bad, tendency to early on then people sort to quality if difference, proximity if both good.But when signals rare, behavior favoring A over B indicating a little more likely to be better correctly understood, more go to A. But if that minor heard later interpreted as needing lots of independent signals, will end up believing A is the best and B is the worst.