Behavioral Economics

Eyster and Rabin (2010), "Naive Herding in Rich-Information Settings"

Sergey Alexeev

May 8, 2017

University of Technology Sydney

1. There is a decision to be made – for example, whether to adopt a new technology, wear a new style of clothing, eat in a new restaurant, or support a particular political position;
2. People make the decision sequentially, and each person can observe the choices made by those who acted earlier;
3. Each person has some private information that helps guide their decision;
4. A person can't directly observe the private information that other people know, but he or she can make inferences about this private information from what they do.

Anderson and Holt (1996)

I: Classroom Games: Information Cascades

An urn that could be majority red or majority blue equally likely:

(a) The 1st student tells what what she sees. The guess conveys perfect information;

(b) If the 2nd student sees the same color – guess the same. If different – brakes the tie. The guess again conveys perfect information;

(c) If the 3rd Student sees opposite guess, then he guess what he sees. If the same then he saw three independent draws and ignores his information;

(d) If the 4th student (and onward) sees three identical guesses in a row and knows that the first 2 were true, while the 3rd is not informative, then she ignores her private information;

- Notes on: full rationality, potential non-optimality, fragility.

- Chances maximizing rule: $\Pr[\mathrm{mb}|\mathrm{s\&h}] > \frac{1}{2}$
- Priors: $\Pr[\mathrm{mb}] = \Pr[\mathrm{mr}] = \frac{1}{2}$
- Conditionals: $\Pr[\mathrm{b}|\mathrm{mb}] = \Pr[\mathrm{r}|\mathrm{mr}] = \frac{2}{3}$
- The 1st student:

$$\begin{aligned} \Pr[\mathrm{mb}|\mathrm{b}] &= \frac{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b}|\mathrm{mb}]}{\Pr[\mathrm{b}]} \\ &= \frac{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b}|\mathrm{mb}]}{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b}|\mathrm{mb}]+\Pr[\mathrm{mr}]\cdot\Pr[\mathrm{b}|\mathrm{mr}]} = \frac{2}{3} \end{aligned}$$

- The 2nd student: ...
- The 3rd student:

$$\begin{aligned} \Pr[\mathrm{mb}|\mathrm{b},\mathrm{b},\mathrm{r}] &= \frac{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b,b,r}|\mathrm{mb}]}{\Pr[\mathrm{b,b,r}]} \\ &= \frac{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b,b,r}|\mathrm{mb}]}{\Pr[\mathrm{mb}]\cdot\Pr[\mathrm{b,b,r}|\mathrm{mb}]+\Pr[\mathrm{mr}]\cdot\Pr[\mathrm{b,b,r}|\mathrm{mr}]} = \frac{2}{3} \end{aligned}$$

- Learning does not have to be Bayesian in the first place, but if it is:

$$B(h^*|e_1, e_2, ...) = \frac{\pi(e_1, e_2, ...|h^*)\pi(h^*)}{\sum_{h \in H} \pi(e_1, e_2, ...|h)\pi(h)}$$

  it could be biased in many ways, e.g.:
- Law of Small Numbers;
- Non-belief in LLN;
- Base-Rate Neglect.

Eyster and Rabin (2010)

I: Naive Inference

- Best response trailing naive inference (BRTNI);
- Builds off of a weaker form of concept of “cursed equilibrium”;
- Inferentially naive players infer “too much”:
    - Inferential naivety push toward overweighting prior signals, its essential property, which drives the central results, is that herders end up placing far too much weight on early relative to late signals.
- The relative weight placed on different predecessors’ signals *vs.* relative weight each person places on her own versus others’ signals.

Eyster and Rabin (2010)

II: Rational and Naive Learning in a Rich Setting

- Rational herders either converge to only weak public beliefs or only very infrequently herd on the wrong action;
- Consider:
  - $\omega \in \{0, 1\}$;
  - $\Pr[\omega = 1] = \pi$;
  - I_t information (private and public) of t;
  - $Q_t \equiv E[\omega | I_t] = \Pr[\omega = 1 | I_t]$;
- Implies upper bound of the posterior belief:

$$\Pr[Q_t \geq q | \omega = 0] \leq \frac{\pi}{1-\pi} \frac{1-q}{q} \tag{P1}$$

- The maximum probability that t can hold information causing him to believe that $\omega = 1$ with at least probability q, when, in fact, $\omega = 0$;
- The bound holds in any binary-state Bayesian learning environment.

Proof of P1

- e.g. maximum probability that herders can be 99% confidence
- In richer environment confident-yet-mistaken herd is even more limited;
- Consider:
  - $A = \left\{0, {}^{1}/_{n}, {}^{2}/_{n}, ..., {}^{(n-1)}/_{n}, 1\right\}$, $n+1$ set of actions
  - Assume $g_t(a;\omega) = -(a_t - \omega)^2$, with $\arg\max a_t = E[\omega|I_t]$
  - S is a denumerable set of signals;
  - $r \equiv \inf_{s \in S}\{\Pr[\omega = 1|s]\}$;
  - All predecessors' are observed and actions converge.

$$\Pr[\lim_{x\to\infty} a_t = 1|\omega = 0] \leq \frac{r}{1-r}\frac{1}{2n-1} \qquad \text{(C2)}$$

- Set $n = 1$ and $r = 0.05$ and apply (P1) and (C2);
- Finer action spaces reduce mistaken herding not by improving players' inference, but by increasing the strength of pubic beliefs necessary for a herd.

Proof of C2

Consider continuous-signal and continuous-action model:

- $\omega \in \{0, 1\}$, ex ante equally likely ;
- t in a countable infinite sequence receives $s_t \in [0, 1]$ which are *i.i.d* conditional on the state;
- Signal have continuous densities f_0 and f_1;
- Before taking action in $[0, 1]$, t observes signal and *all* actions of previous players;
- For simplicity: for each $s \in [0, 1]$, $f_0(s) = f_1(1 - s)$ and $L(s) \equiv f_1(s)/f_0(s)$ with image $\mathcal{R}_+$ and $L'(s) > 0$;
- Simplifications allow to normalize $s = \Pr[\omega = 1|s]$;
- Let $a_t(a_t, ..., a_{t-1}; s_t)$ be an action of t. Rich action space ensures that each player's action fully reveals her beliefs;
- Let $E[\omega|I_t] = \Pr[\omega = 1|I_t]$ a probabilistic belief of t with I_t that $\omega = 1$;
- Assume $g_t(a; \omega) = -(a_t - \omega)^2$, with $\arg\max a_t = E[\omega|I_t]$;
- t takes $a_t = 0$ if $E[\omega|I_t] = 0$ and $a_t = 1$ if $E[\omega|I_t] = 1$;

cf. simple model

The analyzes of a rational player:

- P1 chooses $\ln(a_1/(1-a_1)) = \ln(s_1/(1-s_1))$
- P2 combines P1's action with his private information:

$$\ln\left(\frac{a_2}{1-a_2}\right) = \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) = \ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right)$$

- Interpretation: since agents share a common prior, P2 can adopt P1's posterior as his own prior before incorporating his private signal;
- That's why P3 does not benefit from observing P1's if P2's is seen;
- In general: $\ln(a_t/(1-a_t)) = \sum_{\tau \le t} \ln(s_\tau/(1-s_\tau))$
- Behaviorally: $\ln(a_t/(1-a_t)) = \ln(a_{t-1}/(1-a_{t-1})) + \ln(s_t/(1-s_t))$
- A note on t's unbounded likelihood ratio and continuum of actions.

- BRTNI neglect their predecessors' inferences;
- P1 is not effected (no inference involved);
- P2 correctly infers her signal from her action (typical BNE):

$$\begin{aligned}\ln\left(\frac{a_2}{1-a_2}\right) &= \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) \\ &= \ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right)\end{aligned}$$

- P3 neglects how P2 incorporates P1 signal into his action:

$$\begin{aligned}\ln\left(\frac{a_3}{1-a_3}\right) &= \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{a_2}{1-a_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right) \\ &= \ln\left(\frac{s_1}{1-s_1}\right) + \left(\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right)\right) + \ln\left(\frac{s_3}{1-s_3}\right) \\ &= 2\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right)\end{aligned}$$

- Generally:

$$\ln\left(\frac{a_t}{1-a_t}\right) = \left[\sum\nolimits_{\tau<t} 2^{t-1-\tau}\ln\left(\frac{s_\tau}{1-s_\tau}\right)\right] + \ln\left(\frac{s_t}{1-s_t}\right)$$

- BRTNI play allows a failure of learning of true state even with unbounded signal strength and arbitrary large number of Ps;

$$\textit{In BRTI play, for each } r < 1, \textit{there exist } \delta > 0, \textit{such that}: \\ \Pr[a_t > r, \forall t | \omega = 0] > \delta \tag{P3}$$

- Even when $\omega = 0$ it is possible that BRTNI in an infinite sequence chooses and action that exceeds any given threshold;
- If the first couple of agents receive signals high enough to take actions above r, then with positive probability no agent ever takes an action below r;

Proof of P3

- Unlike rational beliefs, BRTNI beliefs do not form a martingale:
  - When public belief $P_t > 1/2$, then $E[P_{t+1}|P_t] > P_t$
  - When public belief $P_t < 1/2$, then $E[P_{t+1}|P_t] < P_t$
- Beliefs drift in this predictable way because BRTNI players in future periods reweigh information already contained in current beliefs and become fully confident in the wrong state.

BRTI actions and beliefs converge almost surely to 0 *or* 1 (P4)

Proof of P4

Simulations when $\omega = 1$ with $f_o(s) = 2(1-s)$ and $f_1(s) = 2s$

	BNE			BRTNI		
Player	$a \leq 0.05$	$0.05 < a \leq 0.95$	$a > 0.95$	$a \leq 0.05$	$0.05 < a \leq 0.95$	$a > 0.95$
1	0.0026	0.8998	0.0976	0.0025	0.8998	0.0977
2	0.0060	0.6905	0.3035	0.0058	0.6912	0.3030
3	0.0070	0.5059	0.4871	0.0216	0.3819	0.5965
4	0.0069	0.3684	0.6247	0.0483	0.1877	0.7640
5	0.0060	0.2708	0.7232	0.0739	0.0929	0.8332
6	0.0051	0.1995	0.7954	0.0914	0.0463	0.8623
7	0.0041	0.1482	0.8477	0.1016	0.023	0.8754
8	0.0033	0.111	0.8857	0.1068	0.0117	0.8815
9	0.0026	0.0826	0.9148	0.1098	0.0057	0.8845
10	0.0020	0.0624	0.9356	0.1115	0.0029	0.8856

- In rare cases beliefs converge slowly;

For each interval $[c, d] \subset (1/2, 1)$*, there exist* $t \in \mathcal{N}$ *such that if for each* $t \in \{1, ..., T\}, a_t \in [c, d]$ *under BRITNI play, then:* (P5)
$\Pr[\omega = 0|(a_1, ..., a_T)] > \Pr[\omega = 1|(a_1, ..., a_t)]$

- It happens if a player sees a long run of high action (and infer high signals), but receives a low signal herself;
- So she do not converge to 99$ confidence in $\omega = 1$;

Hence, the only way that a large number of players can take actions above 90% percent without any single one of them reaching 99% is that if after a few pieces of evidence supporting $\omega = 1$, all subsequent signals point towards $\omega = 0$, overall indicating $\omega = 0$ more likely.

Proof of P5

- (P3) and (P4) show that BRTNI play converge on the wrong limiting action with positive probability;
- On contrary rational players almost surely converge on the right action;
- Hence, naive Ps obtain strictly lower long-run average payoffs;
- Also if expected cost of overconfidence exceeds the added information in others' actions, BRTNI players can be harmed;
- $g_k(a_k, \omega) = -(a - \omega)^{2n}$, higher n more costly it is to chose action distant from the true state;
- In simulation if $n = 1$ about 11% of time BRTNI converges to wrong liming beliefs and actions with a loss of $-(1)^{2n} = -1$;
- A lower bound on average payoff is $-(1/2)^{2n}$, so if $n = 1$ learning is good, while for $n \geq 2$ no so much, since $-(1/2)^{2n} \geq -1/9$.

Eyster and Rabin (2010)

III: Harmful Learning with Long-Run Agents

- People may choose actions repeatedly
- Consider:
    - Player $\{A, B, C\}$ move in sequence $A, B, C, A, B, C, A \ldots$;
    - Each player's growing collection of private signals almost surely reveals the state;
    - Rational and naive would choose the right action if acted solely, yet...

Suppose that three long-run BRTNI players $\{A, B, C\}$ *move in sequence* $A, B, C, A \ldots$ *Then for each* $r \in (0, 1)$ *there exist* $\delta > 0$ *such that* (P6)

$$\Pr\left[\left(\frac{a_t}{1-a_t}\right) > e^t\left(\frac{r}{1-r}\right), \forall t | \omega = 0\right] > \delta$$

- When $\omega = 0$, for $r > 1/2$, it happens that all long-run BRTNI players play actions above r and converge to certain beliefs that $\omega = 1$;

Proof of P6

Eyster and Rabin (2010)

IV: Discussion and Conclusion

This section should be carefully read and appreciated. So, please just read the paper

Banerjee (1992)

I: A Simple, General Cascade Model

- A simple, cascade model consists of:
  1. States of the world: $\Pr[G] = p$ and $\Pr[B] = 1 - p$;
  2. Payoffs: reject $\rightarrow 0$ or accept $\rightarrow v_g p + v_b(1-p)(= 0)$
  3. Signals: $q > \frac{1}{2}$ (e.g. more reviews for a better restaurant)

$$\Pr[H|G] = q \quad \Leftrightarrow \quad \Pr[L|G] = 1 - q$$
$$\Pr[L|B] = q \quad \Leftrightarrow \quad \Pr[H|B] = 1 - q$$

States

Signals		B	G
	L	q	$1-q$
	H	$1-q$	q

- Individual decision:
  - A high signal shifts expected payoff:
    $v_g \Pr[G] + v_b \Pr[B] = 0 \rightarrow v_g \Pr[G|H] + v_b \Pr[B|H]$

$$\begin{aligned} \Pr[G|H] &= \frac{\Pr[G]\cdot\Pr[H|G]}{\Pr[H]} \\ &= \frac{\Pr[G]\cdot\Pr[H|G]}{\Pr[G]\cdot\Pr[H|G]+\Pr[B]\cdot\Pr[H|B]} \\ &= \frac{pq}{pq+(1-p)(1-q)} > p^* \end{aligned}$$

- Multiple agents:
  - Define S as a set of signals with a high and b low signals then:

$$\begin{aligned} \Pr[G|S] &= \frac{\Pr[G]\cdot\Pr[S|G]}{\Pr[S]} \\ &= \frac{pq^a(1-q)^b}{pq^a(1-q)^b+(1-p)(1-q)^a q^b}^\dagger \end{aligned}$$

  - Implying: $$\begin{aligned} a > (<) b &\Rightarrow \Pr[G|S] > (<) \Pr[G] \\ a = b &\Rightarrow \Pr[G|S] = \Pr[G] \end{aligned}$$

---

*Note that $pq + (1-p) < pq + (1-p)q = q$

†Replace second term in denominator with $(1-p)q^a(1-q)^b$

Banerjee (1992)

II: Sequential Decision-Making and Cascades

- Recall that if P1 and P2 made opposite decisions P3 follows his signal. And future Ps know that;
- If P1 and P2 made the same decision then all do the same;
- If number of acceptance differ from number of rejections by at most one, person follows the signal;
- But once the difference is bigger, everyone follows the majority;
- The difference won't stay within $(-1, 1)$ for long:
  - Divide N into three consecutive players;
  - People in a block receive the same signal with probability: $q^3 + (1-q)^3$
  - The probability that none of these blocks consist of the same signal: $\left(1 - q^3 - (1-q)^3\right)^{N/3}$
  - And goes to 0 as $N \to \infty$

References

- Anderson, L. R. and C. A. Holt (1996). "Classroom Games: Information Cascades". In: *The Journal of Economic Perspectives* 10.4, pp. 187–193. ISSN: 08953309. URL: http://www.jstor.org/stable/2138561.
- Banerjee, A. V. (1992). "A Simple Model of Herd Behavior". In: *The Quarterly Journal of Economics* 107.3, pp. 797–817. ISSN: 00335533, 15314650. URL: http://www.jstor.org/stable/2118364.
- Eyster, E. and M. Rabin (2010). "Naive Herding in Rich-Information Settings". In: *American Economic Journal: Microeconomics* 2.4, pp. 221–243. ISSN: 19457669, 19457685. URL: http://www.jstor.org/stable/25760414.

Technical appendix

The Proof of Proposition 1

Let $\bar{I}_t = \{I_t = (s_t; a_1, ..., a_{t-1}) : Q_t \geq q\}$. From Bayes' Rule,

$$\begin{aligned} \Pr[\omega = 1|\bar{I}_t] &= \frac{\pi}{\pi + (1-\pi)\frac{\Pr[\bar{I}_t|\omega=0]}{\Pr[\bar{I}_t|\omega=1]}} \geq q \\ \Rightarrow \frac{\Pr[\bar{I}_t|\omega=0]}{\Pr[\bar{I}_t|\omega=1]} &\leq \frac{\pi}{1-\pi}\frac{1-q}{1} \end{aligned}$$

Because $\Pr[\bar{I}_t|\omega = 1] \leq 1$, $\Pr[\bar{I}_t|\omega = 0] \leq \frac{\pi}{1-\pi}\frac{1-q}{1}$

back to P1

The Proof of Corollary 2

When public beliefs are that $\Pr[\omega = 1|(a_t, ..., a_{t-1})] = p$, player t with private belief r takes action $a_t = 1$ iff:

$$\begin{aligned} \Pr[\omega = 1|I_t] = \tfrac{pr}{pr+(1-p)(1-r)} &\geq \tfrac{2n-1}{2n} \\ p &\geq \frac{1}{1+\frac{r}{1-r}\frac{1}{2n-1}} \end{aligned}$$

Then (P1) with $q = \frac{1}{1+\frac{r}{1-r}\frac{1}{2n-1}}$ and $\pi = \frac{1}{2}$ gives (C2)

back to C2

Bayesian Updating as a Likelihood Ration (Bayes Factor)

With binary sample space the odds of E are: $O(E) = \frac{P(E)}{P(E^c)}$

- Think of a flip of a fair coin;
- $P(E) = p \Rightarrow O(E) = {}^{p}\!/_{1-p}$

Bayesian updating – in the language of odds – is prior odds updated to posterior odds:

$$\begin{aligned}
\text{Bayes factor} = O(H|D) &= \frac{P(D|H)}{P(D|H^c)} \\
&= \frac{P(D|H) \cdot P(H)}{P(D|H^c) \cdot P(H^c)} \\
&= \frac{P(D|H)}{P(D|H^c)} \cdot \frac{P(H)}{P(H^c)} \\
&= \frac{P(D|H)}{P(D|H^c)} \cdot O(H) \\
\text{posterior odds} &= \text{Bayes factor} \times \text{prior odds}
\end{aligned}$$

Log odds are more convenient in practice:

$$\begin{aligned}
O(H|D_1, D_2) &= BF_2 \cdot BF_1 \cdot O(H) \\
\ln(O(H|D_1, D_2)) &= \ln(BF_2) + \ln(BF_1) + \ln(O(H))
\end{aligned}$$

back

The Proof of Proposition 3 (beginning)

Pick $r \in (1/2, 1)$, define $R = \ln(1/(1-r)) > 0$, let P_t be a log likelihood of public belief at period t.

With BRTNI play $P_{t+1} = 2P_t + \ln(s_t/(1-s_t))$

When $\omega = 0$, with positive probability $P_2 \geq 3R$

If $\ln(s_t/(1-s_t)) > -tR \ \forall t$ then:

$$\begin{aligned} P_3 = 2P_2 + \ln(s_2/(1-s_2)) &> 2 \times 3R - 2R = 4R \text{ and} \\ P_4 = 2P_3 + \ln(s_3/(1-s_3)) &> 2 \times 4R - 3R = 5R, \text{ etc.} \end{aligned}$$

In general:

$$\begin{aligned} P_t &> (t+1)R \text{ and so} \\ \ln(a_t/(1-a_t)) = P_t + \ln(s_t/(1-s_t)) &> (t+1)R - tR = R \end{aligned}$$

Now...

back to P3

The Proof of Proposition 3 (continuation)

$$\begin{aligned} \Pr\left[\ln\left(S_t/(1-S_t)\right) < -tR|\omega=0\right] &< \Pr\left[\left|\ln\left(S_t/(1-S_t)\right)\right| > tR\,|\omega=0\right] \\ &\overset{\ddagger}{<} 1/(tR)^2 E\left[\left|\ln\left(S_t/(1-S_t)\right)\right|^2|\omega=0\right] \end{aligned}$$

Also,

$$\begin{aligned} Q \equiv E\left[\left(\ln\left(S/(1-S)\right)\right)^2|\omega=0\right] &= \int_0^1 \left(\ln\left(s/(1-s)\right)\right)^2 f_0(s)ds \\ &\leq M\int_0^1 \left(\ln\left(s/(1-s)\right)\right)^2 ds \\ &= M\left(\pi^2/3\right) \end{aligned}$$

for $M \equiv \sup\{f_0(s) : s \in [0,1]\}$, which is finite by the continuity of f_0

‡Markov inequality: $\Pr[X \geq a] \leq E[X]/a$ if X is nonnegative r.v and $a > 0$

back to P3

The Proof of Proposition 3 (finale)

Define $\tau = \min\{t \in \mathcal{N} : Q < t^2 R^2\}$ so that for each $t \geq \tau$, $((t^2 R^2 - Q)/t^2 R^2) \in (0,1)$, and let $C(R) \equiv \prod_{t=1}^{\tau-1} (1 - F_0(-tR)) > 0$.

$$
\begin{aligned}
\Pr\left[{}^{S_t}/_{1-S_t} > e^{-tR}, \forall t \,|\omega = 0\right] \quad &> \quad C(R) \prod_{t \geq \tau} {}^{(t^2R^2-Q)}/_{(t^2R^2)} \\
&= \quad C(R) \exp\left\{\sum_{t \geq \tau} {}^{(t^2R^2-Q)}/_{(t^2R^2)}\right\} \\
&= \quad C(R) \exp\left\{\sum_{t \geq \tau} -{}^{Q}/_{z_t}\right\}
\end{aligned}
$$

for $z_t \in (t^2 R^2 - Q, t^2 R^2)$, by the Mean-Value Theorem. Then,

$$
\begin{aligned}
\Pr\left[{}^{S_t}/_{1-S_t} > e^{-tR}, \forall t \,|\omega = 0\right] \quad &> \quad C(R) \exp\left\{\sum_{t \geq \tau} -{}^{Q}/_{t^2R^2}\right\} \\
&> \quad C(R) \exp\left\{\sum_{t \geq 1} -{}^{Q}/_{t^2R^2}\right\} \\
&= \quad C(R) \exp\left\{-\left({}^{Q\pi}/_{6R^2}\right)\right\} > 0
\end{aligned}
$$

back to P3

The Proof of Proposition 4

From above, write: $2^{1-t}P_t = \sum_{\tau<t} 2^{-\tau} \ln\left({}^{s_\tau}/_{1-s_\tau}\right)$

Since the three series

$$\sum_{\tau=1}^{\infty} E\left[2^{-\tau}\ln\left({}^{s}/_{1-s}\right)|\omega=0\right] \;=^{\S}\; 2E\left[\ln\left({}^{s}/_{1-s}\right)|\omega=0\right]$$

$$\sum_{\tau=1}^{\infty} \mathrm{var}\left[2^{-\tau}\ln\left({}^{s}/_{1-s}\right)|\omega=0\right] \;=^{\P}\; {}^{1}/_{3}\mathrm{var}\left[\ln\left({}^{s}/_{1-s}\right)|\omega=0\right]$$

$$\sum_{\tau=1}^{\infty} \mathrm{var}\left[2^{-\tau}\ln\left({}^{s}/_{1-s}\right)|\geq 1\right] \;=^{\|}\; \sum_{\tau=1}^{\infty} 4^{-\tau}\mathrm{var}\left[\ln\left({}^{s}/_{1-s}\right)|\omega=0\right]$$

Kolmogorov's Three-Series Theorem implies that $2^{1-t}P_t$ converges a.s.

---

§ Follow from finiteness of the second moment (and therefore the first)

¶ See above

‖ by Chebyshev's inequality

back to P4

The Proof of Proposition 5

back to P5

The Proof of Proposition 6

back to P6