A Subjective Theory of Belief in the Gambler's Fallacy*

Jawwad Noor and Fernando Payró

July 31, 2026

Abstract

Studies suggest that people often believe that random processes exhibit reversion to the mean, regardless of the size of the sample. In a canonical coin-tossing environment, a belief in mean reversion towards some bias θ^* is straightforward to define if it is assumed that the tosses are objectively i.i.d. with bias θ^*. In this paper we define and represent a belief in mean reversion without assuming the existence of an objective random process. This decouples the notion of a belief in mean reversion from the notion of misspecified beliefs that is predominantly used in the literature.

JEL Classification: D01, D9.

Keywords: Belief Biases, Mean Reversion, Gambler's Fallacy.

***Noor is at the Department of Economics, Boston University, 270 Bay State Road, Boston MA 02215. Payró is at the Department of Economics and Economic History, Universitat Autònoma de Barcelona, Barcelona School of Economics and Center for the Study of Organizations and Decisions in Economics, Campus de la UAB, Edifici B, Bellaterra, 08193.** **Payró gratefully acknowledges financial support from the Ministerio de Economía y Competitividad and Feder (PID2023-147183NB-I00), the Generalitat de Catalunya (2021 SGR 00194) and the Spanish Agencia Estatal de Investigación (AEI), and the Severo Ochoa Programme for Centres of Excellence in R\&D (Barcelona School of Economics CEX2024-001476-S), funded by MCIN/AEI/10.13039/501100011033.**

1 Introduction

Studies show that people tend to believe that streaks are likely to break in sequences generated by an i.i.d. process. This so-called *belief in the Gambler's Fallacy* is supported by two kinds of evidence:

(i) Settings with an objective random process: Experiments inform subjects that sequences of heads and tails are generated by i.i.d. tosses of a coin with bias (say) $\theta^* = \frac{1}{2}$, and yet subjects believe, for instance, that the probability of tails is higher following a streak of heads: for instance, $P(HHHT) > P(HHHH)$. Observing gamblers playing roulette in a casino setting, Croson and Sundali (2005) document that over 65% of subjects bet against the color that appeared 5 times in a row. Studying betting on lottery numbers, Terrell (1994) demonstrates that people are significantly less likely to bet on a lottery number if it was recently a winning number (see Suetens et al. (2016) for a more recent study).

(ii) Settings without an objective random process: Using data from a field experiment, Chen, Moskowitz and Shue (2016) document that loan officers are more likely to reject (resp. accept) a loan application after they have accepted (resp. rejected) the previous application, and estimate that in some treatments up to 9% of decisions are erroneous due to such a sequencing effect. They find a similar bias in decisions by US judges in refugee asylum cases. Jin and Peng (2022) show that various puzzles in finance (such as the disposition effect, where traders hold on to recently losing stock and sell recently winning stock) can be explained by a disbelief in streaks of the type exhibited in the Gambler's Fallacy.

Motivated by evidence that mostly comes from settings of type (i), the literature in economics and psychology (Tversky and Kahneman (1974), Rabin (2002), Noor and Payró (2026)) predominantly takes the view that subjects' beliefs about randomness are misspecified, that is, subjects' beliefs about i.i.d. processes are simply incorrect. However, one may wish to avoid philosophical questions regarding whether it is justified, or even necessary for the practice of economics, that the analyst gets to ordain whether an agent's beliefs are wrong or irrational.[1] Besides, most economic settings are closer to type (ii), where there is no known objective i.i.d. process relative to which beliefs can be judged as correct or incorrect. This leads us to the question of how to define the Gambler's Fallacy without any reference to an objective random process. The realization of this possibility is our objective here. By showing what it means for a belief to display mean reversion without referring to an objective benchmark, we separate the empirical content of the Gambler's Fallacy from the interpretation that agents misunderstand randomness. This allows us to extend the study of such beliefs to a broader class of environments where there is no known process that makes such expectations objectively wrong.

2 Preliminaries

2.1 Primitives

Consider a canonical coin-tossing environment: the possible realizations of a coin toss in any period i are $\Omega_i = \Omega = \{0, 1\}$. We interpret 1 as heads H, and 0 as tails T, and sometimes use the notation H, T instead of $1, 0$. The space of all realizations of any $1 \leq n \leq \infty$ tosses is $\Omega^n = \prod_{i=1}^{n} \Omega_i$. Adopt the convention that $\Omega^0 = \emptyset$, and use $x = (x_1, x_2, ...) \in \Omega^\infty$ to denote an infinite sequence and $x^n = (x_1, ..., x_n) \in \Omega^n$ to denote a finite sequence of length n. Our primitive is the agent's belief P modeled as a probability measure over the usual infinite horizon sample space $(\Omega^\infty, \Sigma^\infty)$.[2] Although P is subjective, we take it as observable. The interpretation is that it has been elicited from choices over a sufficiently rich collection of contingent claims.

The *sample mean* at any point n of a sequence $x = (x_1, x_2, ...)$ is the mean

[1]While misspecification is often viewed as arising due to faulty intuition or limited cognitive abilities, Noor and Yang (2026) show that apparent misspecification is also rationalized by any mistrust the agent may harbor toward the information they are served about the data generating process.

[2]For any $A^n \subset \Omega^n$, define an n-cylinder by the event $A^n\Omega^\infty := \{(x^n z) \in \Omega^\infty : x^n \in A^n \text{ and } z \in \Omega^\infty\}$ in the infinite horizon space Ω^∞. Then $\Sigma^\infty = \sigma(\cup_{n=1}^{\infty} \Sigma^n)$ is the σ-algebra generated by all the n-cylinders, $n = 1, 2, ...$

number of heads:

$$\overline{x}^n := \frac{1}{n}\sum_{i\leq n} x_i.$$

We maintain throughout that P is non-trivial in that it expresses some mean reversion: there exist histories x^n and y^m such that

$$P(1 \mid x^n) > P(0 \mid x^n) \quad \text{and} \quad P(0 \mid y^m) > P(1 \mid y^m).$$

2.2 Considerations

We wish to define what it means for the agent to believe in the Gambler's Fallacy. The term is used to describe beliefs following a streak, for instance, the agent expects $HHHT$ to be more likely than $HHHH$. In this example, four heads is viewed as a streak that is "too long" and not likely to persist. Accordingly, Noor and Payró (2026) define

Definition 1. *P exhibits the Gambler's Fallacy if for each outcome $\omega \in \Omega$ and history x^n there exists M such that for all $m \geq M$,*

$$P(1-\omega|x^n\omega^m) > P(\omega|x^n\omega^m).$$

The definition states the existence of a minimal length M for which a streak ω^m is perceived to be "too long". The length M can depend on x^n and ω in an arbitrary way. Although the definition is stated in terms of observable conditional beliefs, its existential quantifier makes it difficult to verify directly. Further, it is silent on what gives rise to the perception of a streak being too long. One would like a condition that identifies what drives this perception, and one would like for it not to involve an existential quantifier. Our hypothesis is the following:

While the Gambler's Fallacy suggests that tails is more likely after HHH, our intuition is that tails may also be considered more likely after $HHHHTH$, that is, $P(HHHHTHT) > P(HHHHTHH)$. In this example there is no streak of heads preceding the last outcome. Instead, there is a *concentration* *of heads* in the history up to that point, which may lead one to believe that tails is more likely. This suggests that a driver of the Gambler's Fallacy is not the presence of a streak per se, but rather the perception that "too many heads have occurred".

This hypothesis is reminiscent of the "Law of Small Numbers" which posits that the heads and tails in any sequence are expected to match large sample proportions (Tversky and Kahneman (1971)). Thus, if i.i.d. tosses of a coin are thought to have bias θ^*, then it is expected that any finite sequence is most likely to have a proportion closest to θ^* heads, even if the sequence is short. An implication is that the agent possesses a belief that random processes are "self-correcting" (Tversky and Kahneman (1974)), that is, if the proportion of heads (say) exceeds θ^* then tails is considered more likely on the next toss in order to correct for the imbalance. Noor and Payró (2026) model the latter formally, and refer to it as a *belief in mean reversion*. However, while this literature (Noor and Payró (2026) and the experiments supporting the Law of Small Numbers) presumes that the bias θ^* is objectively given, in this paper we seek to work in a purely subjective environment. Thus, we need a way to subjectively define what it means for the agent to perceive that "too many heads have occurred".

The nontriviality of the problem is best illustrated by considering some seemingly plausible candidates. A natural first thought would be to posit that if the agent thinks head is more likely after x^{n-1} and tails is more likely after y^{m-1}, then there must be a lower concentration of heads in x^{n-1} compared to y^{m-1}. Concentration of heads is naturally defined as the sample mean. Then, the postulate can be stated as follows: for any distinct x^{n-1} and y^{m-1},[3]

$$P(1|x^{n-1}) \geq P(0|x^{n-1}) \text{ and } P(1|y^{m-1}) \leq P(0|y^{m-1})$$

$$\implies \overline{x^{n-1}} \leq \overline{y^{m-1}},$$

with strict inequality in the conclusion if there is at least one strict inequality

[3] If we allowed the condition to hold when x^{n-1} and y^{m-1} are not distinct then the condition would imply $P(1|x^{n-1}) = P(0|x^{n-1})$ for all x^{n-1}.

in the hypothesis. This, however, has non-intuitive implications. To illustrate, take $x^{n-1} = HHT$ and $y^{m-1} = HHHHTT$. Both histories have the same sample mean, namely 2/3, and so the axiom rules out the possibility of either $P(1|x^{n-1}) > P(0|x^{n-1})$ or $P(1|y^{m-1}) < P(0|y^{m-1})$. However, adding one more toss affects the sample mean differently in the two cases. After HHT, the two possible next-period sample means are 1/2 and 3/4. After $HHHHTT$, they are 4/7 and 5/7. Hence, for any subjective benchmark $\theta^* \in (5/8, 9/14)$, beliefs that mean revert toward θ^* would assign strictly higher probability to heads after HHT, and strictly higher probability to tails after $HHHHTT$. This illustrates that the sample mean of a history alone does not capture the sense in which a sequence contains "too many heads".

We now show that the sample mean after a history *and* outcome is also insufficient. Hypothesize that for any distinct x^{n-1} and y^{m-1},[4]

$$P(1|x^{n-1}) \geq P(0|x^{n-1}) \text{ and } P(1|y^{m-1}) \leq P(0|y^{m-1})$$

$$\implies \overline{x^{n-1}1} \leq \overline{y^{m-1}0},$$

with strict inequality in the conclusion if there is at least one strict inequality in the hypothesis. This too has non-intuitive implications. For instance, assume beliefs exhibit mean reversion with respect to $\theta^* = \frac{1}{2}$. Then it is easy to see that if $x^{n-1} = TTH$ then $P(1|x^{n-1}) > P(0|x^{n-1})$, and $y^{m-1} = HHT$ then $P(1|y^{m-1}) < P(0|y^{m-1})$. However, the axiom requires $\frac{1}{2} = \overline{x^{n-1}1} < \overline{y^{m-1}0} = \frac{1}{2}$, a contradiction. [5]

[4]If we allowed the condition to hold even when x^{n-1} and y^{m-1} are not distinct then the condition would imply $P(1|x^{n-1}) = P(0|x^{n-1})$ for all x^{n-1}.

[5]Weakening the strict inequality part of the axiom does not fix this issue. Take $\theta^* = \frac{2}{3}$, $x^{n-1} = HHT$ and $y^{m-1} = HHHHHHTT$. Then $P(1|x^{n-1}) > P(0|x^{n-1})$ and $P(1|y^{m-1}) < P(0|y^{m-1})$ but $\overline{x^{n-1}1} = \frac{3}{4} > \frac{2}{3} = \overline{y^{m-1}0}$.

3 Self-Correcting Bias

3.1 Axioms

The previous discussion demonstrates that neither the sample mean of a history nor the sample mean after the next toss captures the sense in which a sequence contains "too many heads". To identify the relevant statistic, fix a history x^{n-1}. If the agent believes in mean reversion toward some (unobserved) parameter θ, then heads is *corrective* after x^{n-1} whenever the sample mean after $x^{n-1}1$ is closer to θ than the sample mean after $x^{n-1}0$. Similarly, tails is corrective whenever the sample mean after $x^{n-1}0$ is closer to θ than the sample mean after $x^{n-1}1$. We observe that, given x^{n-1}, the value of θ for which the two continuations are equally close is

$$\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2}.$$

We call this value the *cutoff* associated with history x^{n-1}. If the agent's subjective parameter lies above this cutoff, then heads is corrective; if it lies below, then tails is corrective.

Now consider two histories x^{n-1} and y^{m-1}, which may have different lengths. Because of this, comparing their sample means directly does not provide a natural way to rank which history contains "more heads" relative to the agent's subjective benchmark. The cutoffs do provide such a ranking. For instance, if

$$P(1 \mid x^{n-1}) > P(0 \mid x^{n-1}) \quad \text{and} \quad P(1 \mid y^{m-1}) < P(0 \mid y^{m-1}),$$

then heads is corrective after x^{n-1}, whereas tails is corrective after y^{m-1}. This means that x^{n-1} contains too few heads relative to the subjective benchmark, while y^{m-1} contains too many, as the subjective parameter θ^* must lie between the two cutoffs:

$$\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} < \theta^* < \frac{\overline{y^{m-1}0} + \overline{y^{m-1}1}}{2}.$$

This observation leads to our key axiom.

Axiom 1. *(Subjective Mean Reversion) For any* x^{n-1} *and* y^{m-1}*,*

$$P(1|x^{n-1}) > P(0|x^{n-1}) \textit{ and } P(1|y^{m-1}) < P(0|y^{m-1})$$

$$\implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} < \frac{\overline{y^{m-1}0} + \overline{y^{m-1}1}}{2}.$$

A natural strengthening is:

Axiom 2. *(Strong Subjective Mean Reversion) For any* x^{n-1} *and* y^{m-1}*,*

$$P(1|x^{n-1}) \geq P(0|x^{n-1}) \textit{ and } P(1|y^{m-1}) \leq P(0|y^{m-1})$$

$$\implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}}1}{2} \leq \frac{\overline{y^{m-1}0} + \overline{y^{m-1}}1}{2}.$$

Moreover, strict inequality in the conclusion holds if there is at least one strict inequality in the hypothesis.

Define the set of average sample means that correspond to a strict belief in tails, and the corresponding set for a strict belief in heads, as follows:

$$K = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}}1}{2} : \text{for some } x^{n-1} \text{ s.t. } P(1|x^{n-1}) < P(0|x^{n-1})\}$$

$$L = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}}1}{2} : \text{for some } x^{n-1} \text{ s.t. } P(1|x^{n-1}) > P(0|x^{n-1})\}.$$

The following is a mild continuity condition, requiring that the limit of a sequence corresponding to a strict belief in tails (resp. heads) should involve a weak belief in tails (resp. heads).

Axiom 3. *(Mean Reversion Continuity) If a sequence in* K *(resp.* L*) converges to* $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}}1}{2}$ *for some* z^{k-1} *, then* $P(1|z^{k-1}) \leq$ *(resp.* $\geq$*)*$P(0|z^{k-1})$*.*

3.2 Results

We show that our subjective notion of the Gambler's Fallacy is connected with the notion of mean reversion studied in Noor and Payró (2026). For any parameter θ^*, we denote the distance between the sample mean and θ^* by

$$d_{\theta^*}(x^n) = |\overline{x}^n - \theta^*|.$$

Adopt the convention that $d_{\theta^*}(x^{i-1}) = \overline{x}^{i-1} = 0$ when $i = 1$. For any $\theta^* \in (0,1)$, say that an outcome x_n is *mean reverting at* x^{n-1} if, at history x^{n-1}, outcome x_n induces a smaller distance to θ^* than the other outcome $1 - x_n$,

$$d_{\theta^*}(x^{n-1}x_n) \leq d_{\theta^*}(x^{n-1}(1 - x_n)).$$

A very basic notion of belief in mean reversion would require simply that, at every history, the agent assigns a higher conditional probability to the mean reverting outcome:

Axiom 4. *(Mean Reversion wrt θ^*) For any $n \geq 1$, sequence $x^{n-1} \in \Omega^{n-1}$ and outcomes $x_n, y_n \in \Omega$,*

$$d_{\theta^*}(x^{n-1}x_n) \leq d_{\theta^*}(x^{n-1}y_n) \implies P(x_n|x^{n-1}) \geq P(y_n|x^{n-1}).$$

A natural strengthening is the "if and only if" counterpart:

Axiom 5. *(Strong Mean Reversion wrt θ^*) For any $n \geq 1$, sequence $x^{n-1} \in \Omega^{n-1}$ and outcomes $x_n, y_n \in \Omega$,*

$$d_{\theta^*}(x^{n-1}x_n) \leq d_{\theta^*}(x^{n-1}y_n) \iff P(x_n|x^{n-1}) \geq P(y_n|x^{n-1}).$$

These conditions are reminiscent of a "belief that a random process is self-correcting" (Tversky and Kahneman (1974)). We model the notion of "self-correction" as a belief that is driven by distance $d_{\theta^*}(x^n)$.[6] Note well that at history x^{n-1}, the belief that outcome x_n will obtain on the next toss does *not* depend on the mean at the history, $\overline{x}^{n-1}$, but rather the mean achieved at the

[6]The use of absolute differences implies that deviations from θ^* from above or from below are treated symmetrically. This is natural for a coin-tossing setup but may not be in others. For instance, symmetry may be violated if there is some utility attached to each outcome - noise in a forest may be associated with danger and silence with safety, or positive vs negative news about a favored political party may not be treated symmetrically.

next toss, $\overline{x^{n-1}x_n}$. Thus, this notion of self-correction has a *forward-looking flavor*.

Say that P has *full support* if $P(A^n) > 0$ for every nonempty $A^n \subset \Omega^n$. Our main result is:

Theorem 1. *A full support belief P satisfies Subjective Mean Reversion and Mean Reversion Continuity if and only if there exists some $\theta^* \in (0,1)$ with respect to which beliefs exhibit Mean Reversion.*

When Subjective Mean Reversion and Mean Reversion Continuity hold, then the agent behaves as if she believes the coin to possess some bias θ^* towards which the sample mean corrects at each toss n. This bias is *stationary* in the sense that the agent expects mean reversion towards it regardless of n.

"Self-correction" is defined by a belief that is driven by distance $d_{\theta^*}(x^n)$.[7] Note well that at history x^{n-1}, the belief that outcome x_n will obtain on the next toss does *not* depend on the mean at the history, $\overline{x}^{n-1}$, but rather on the mean achieved after the possible outcomes $\overline{x^{n-1}1}$ and $\overline{x^{n-1}0}$. Thus, self-correction has a forward-looking flavor.

The subjective bias θ^* is not unique in general. The next result shows that strengthening Subjective Mean Reversion yields uniqueness of θ^*. It also demonstrates how θ^* can be computed, using either of the sets, K or L.

Theorem 2. *A full support belief P satisfies Strong Subjective Mean Reversion and Mean Reversion Continuity if and only if there exists some $\theta^* \in (0,1)$ with respect to which beliefs exhibit Strong Mean Reversion. Moreover, θ^* is unique and satisfies $\theta^* = \inf K = \sup L$.*

As a corollary, we obtain a simple representation result and establish that its strong version implies the Gambler's Fallacy:

Proposition 1. *A full support belief P satisfies Subjective Mean Reversion (resp. Strong Subjective Mean Reversion) and Mean Reversion Continuity if*

[7] Embodied in the use of absolute differences is that deviations from θ^* from above or from below are treated symmetrically. This is natural for a coin-tossing setup, though may not be in others.

and only if there exists $\theta^* \in (0,1)$ *such that for all* n *and* $x^n \in \Omega^n$*,*

$$P(x^n) = \prod_{i=1}^{n} (\theta_{x^{i-1}})^{x_i}(1-\theta_{x^{i-1}})^{1-x_i},$$

where $(\theta_{x^{i-1}})^{x_i}(1-\theta_{x^{i-1}})^{1-x_i} = g^i(d_{\theta^*}(x^{i-1}x_i), x^{i-1})$ *for some function* $g^i(\cdot, x^{i-1}) : [0,1] \to (0,1]$ *that is continuous and weakly (resp. strictly) decreasing for each* i *and* x^{i-1}*. Moreover, in the strong case,* P *exhibits the Gambler's Fallacy.*

The (history x^{i-1}-dependent) bias towards heads is $\theta_{x^{i-1}} = g^i(d_{\theta^*}(x^{i-1}1), x^{i-1})$ and towards tails is $1-\theta_{x^{i-1}} = g^i(d_{\theta^*}(x^{i-1}0), x^{i-1})$. As g^i is weakly decreasing, the model requires that the bias is weakly self-correcting: if $d_{\theta^*}(x^{i-1}1) \leq d_{\theta^*}(x^{i-1}0)$ then $\theta_{x^{i-1}} \geq 1-\theta_{x^{i-1}}$. While $g^n(\cdot, x^{n-1})$ is a function on $[0,1]$, it is uniquely pinned down only for the values of d_{θ^*} that are achievable in toss n. Specifically, on such values of d_{θ^*},

$$g^n(d_{\theta^*}(x^n), x^{n-1}) = \frac{g^n(d_{\theta^*}(x^n), x^{n-1}) \times \prod_{i=1}^{n-1} g^i(d_{\theta^*}(x^i), x^{i-1})}{\prod_{i=1}^{n-1} g^i(d_{\theta^*}(x^i), x^{i-1})}$$

$$= \frac{P(x^n)}{P(x^{n-1})} = P(x_n|x^{n-1}),$$

that is, the function $g^n(\cdot, x^{n-1})$ reflects conditional beliefs.

To see that in the strong case P exhibits the Gambler's Fallacy, fix a history x^n and an outcome $\omega \in \{0,1\}$. For all sufficiently large m, the continuation $1-\omega$ brings the sample mean following $x^n\omega^m$ strictly closer to θ^* than does ω. Since g^{n+m+1} is strictly decreasing,

$$P(1-\omega \mid x^n\omega^m) > P(\omega \mid x^n\omega^m),$$

and hence P satisfies Definition 1.[8] See Noor and Payró (2026) for examples

[8]While the model posits that conditional beliefs depend on the distance $|\frac{\sum x^i}{i} - \theta^*|$, it subsumes the case where beliefs depend on the difference $|\sum x^i - i\theta^*|$ between actual number of heads and the expected number of heads. The reason is that g^i depends on i, and so we can think of models where $g^i(|\frac{\sum x^i}{i} - \theta^*|, x^{i-1}) = f^i(i \times |\frac{\sum x^i}{i} - \theta^*|, x^{i-1})$. Since $i \times |\frac{\sum x^i}{i} - \theta^*| = |\sum x^i - i\theta^*|$ we see that g^i now depends on the deviation in head counts.

of random processes that exhibit Mean Reversion and also for how the model can be extended to the case where the random process has more than two outcomes.

3.3 Comparison with Noor and Payró (2026)

Noor and Payró (2026) develop a theory of belief in mean reversion. In the spirit of the literature, they assume the existence of an objective random process. Formally, the sequences are assumed to be generated through i.i.d. tosses of a coin with bias $\theta^* \in (0, 1)$, leading to a probability law

$$Q_{\theta^*}(x^n) = \prod_{i=1}^{n} (\theta^*)^{x_i} (1 - \theta^*)^{1-x_i}.$$

They analyze and represent beliefs P that exhibit Mean Reversion with respect to the θ^* defined by this objective random process.

We dub this the *misspecification approach,* since it deems the agent to have correct knowledge of the bias θ^* of the objective random process, but to otherwise mistakenly believe that the coin is self-correcting. Conceptually speaking, there is an obvious distinction between an agent believing in mean reversion and an agent having an incorrect understanding of the environment. Indeed, the latter is not necessary for the former, as an agent may exhibit a belief in mean reversion due to information that the analyst does not possess. Our results decouple the notion of mean reversion from the notion of misspecification. We do so by adopting a *subjective approach* where a belief in mean reversion is an outcome of some property of the agent's beliefs which makes no reference to any objective random process. While the contribution of Noor and Payró (2026) is the formalization of psychological notions discussed in the context of the Law of Small Numbers, this paper makes the conceptual contribution of divorcing these notions from the misspecification perspective typically employed in the literature.

4 Concluding Remarks

We conclude by discussing some economic settings where agents do not typically have access to the objective random process.

Consider speculative trading. Traders may believe that, as a consequence of some complex interaction of economic variables and randomness in other traders' expectations, the sequence of positive or negative returns of an asset will correct toward some benchmark. Heterogeneity in such benchmarks, or in the intensity of the corrections, may then generate disagreement and trade even when traders observe the same history. Our subjective approach is particularly suited to such environments because it does not require the existence of a known objective process relative to which traders' beliefs can be classified as correct or incorrect – their beliefs just are what they are. We leave an analysis of the implications for contracting, trading, and prices to future research.

Next, consider a principal-agent environment in which the performance outcome of effort by the agent is generated by some process that is random conditional on the effort. The agent's effort and the conditional random process are both unobserved by the principle. Instead, the principal observes a sequence of performance outcomes and uses this history to form beliefs about future performance. If the principal believes that performance is self-correcting, a sequence of unusually good outcomes may lead her to expect weaker future performance, while a sequence of poor outcomes may lead her to expect improvement. Anticipating these beliefs, the agent may adjust her behavior across periods to accommodate or influence the principal's interpretation of her performance history. Subjective mean reversion may therefore affect the design of performance pay, retention and promotion decisions, and the extent to which compensation responds to recent outcomes.

A Proof of Proposition 1

Lemma 1. *Any full support belief P can be written as*

$$P(x^n) = \prod_{i=1}^{n} (\theta_{x^{i-1}})^{x_i} (1 - \theta_{x^{i-1}})^{1-x_i},$$

where $\theta_{x^{i-1}} = \frac{P(x^{i-1}1)}{P(x^{i-1})}$ *and* $1 - \theta_{x^{i-1}} = \frac{P(x^{i-1}0)}{P(x^{i-1})}$.

Proof. Since beliefs have full support we can always write $P(x^n) = P(x^1) \prod_{i=2}^{n} \frac{P(x^i)}{P(x^{i-1})}$. Define $\theta_{x^{i-1}} = \frac{P(x^{i-1}1)}{P^i(x^{i-1})}$ and $\gamma_{x^{i-1}} = \frac{P(x^{i-1}0)}{P^i(x^{i-1})}$ to obtain the representation

$$P(x^n) = \prod_{i=1}^{n} (\theta_{x^{i-1}})^{x_i} (\gamma_{x^{i-1}})^{1-x_i}.$$

To conclude, note that since conditional beliefs sum to 1 we have $\frac{P(x^{i-1}1)}{P(x^{i-1})} + \frac{P(x^{i-1}0)}{P(x^{i-1})} = 1$, and equivalently $\theta_{x^{i-1}} + \gamma_{x^{i-1}} = 1$. □

Lemma 2. *For any $\theta^* \in (0, 1)$, a full support belief P satisfies Mean Reversion (resp. Strong Mean Reversion) wrt θ^* iff for each i there exists a function $g^i : [0, 1] \times \Omega^{i-1} \to (0, 1]$ that is continuous and weakly (resp. strictly) decreasing in its first argument, and moreover for all n and $x^n \in \Omega^n$,*

$$P(x^n) = \prod_{i=1}^{n} (\theta_{x^{i-1}})^{x_i} (1 - \theta_{x^{i-1}})^{1-x_i},$$

where $\theta_{x^{i-1}} := g^i(d_{\theta^*}(x^{i-1}1), x^{i-1})$ *and* $1 - \theta_{x^{i-1}} := g^i(d_{\theta^*}(x^{i-1}0), x^{i-1})$.

Proof. Mean Reversion wrt θ^* implies that for any x^{n-1}, conditional beliefs $\frac{P(x^{n-1}x_n)}{P(x^{n-1})}$ depend on x_n only through $d_{\theta^*}(x^n)$. Therefore we can write $g^n(d_{\theta^*}(x^n), x^{n-1}) = \frac{P(x^{n-1}x_n)}{P(x^{n-1})}$ and insert this in Lemma 1 to obtain a representation satisfying $d_{\theta^*}(x^{n-1}x_n) \leq d_{\theta^*}(x^{n-1}y_n) \implies g^n(d_{\theta^*}(x^{n-1}x_n), x^{n-1}) \geq g^n(d_{\theta^*}(x^{n-1}y_n), x^{n-1})$. The "iff" counterpart holds under Strong Mean Reversion. In the representation, the domain of g^n is defined only by two sequences $x^{n-1}0, x^{n-1}1$, and by the full support assumption it must be that g^n takes values in $(0,1]$. Therefore, given Mean Reversion (resp. Strong Mean Reversion), the function can be extended to $[0,1]$ continuously so that it is weakly (resp. strictly) decreasing. The necessity of Mean Reversion follows from the fact that the representation implies $g^n(d_{\theta^*}(x^{n-1}x_n), x^{n-1}) = \frac{P(x^{n-1}x_n)}{P(x^{n-1})}$. $\square$

B Proof of Theorems 1 and 2

Define the following sets:

$$K = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} : \text{for some } x^{n-1} \text{ s.t. } P(1|x^{n-1}) < P(0|x^{n-1})\}$$

$$L = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} : \text{for some } x^{n-1} \text{ s.t. } P(1|x^{n-1}) > P(0|x^{n-1})\}$$

$$M = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} : \text{for some } x^{n-1} \text{ s.t. } P(1|x^{n-1}) = P(0|x^{n-1})\}.$$

Define

$$\overline{\theta} := \inf K \text{ and } \underline{\theta} := \sup L.$$

We note that:

Lemma 3. *(a) If beliefs satisfy Subjective Mean Reversion then $\underline{\theta} \leq \overline{\theta}$. Moreover, if Mean Reversion Continuity is satisfied and if $\underline{\theta} = \overline{\theta}$, then $\theta' < \underline{\theta} = \overline{\theta} < \theta$ for any $\theta \in K$ and $\theta' \in L$.*

(b) If beliefs satisfy Strong Subjective Mean Reversion then $\underline{\theta} = \overline{\theta}$. Moreover, either $M = \emptyset$ or $M = \{\hat{\theta}\}$ where $\hat{\theta} = \underline{\theta} = \overline{\theta}$.

Proof. By Subjective Mean Reversion, every element of K is greater than every element of L. It follows that $\underline{\theta} \leq \overline{\theta}$, proving the first claim. To prove the next claim, suppose $\underline{\theta} = \overline{\theta}$. If $\underline{\theta} = \overline{\theta}$ is an irrational number then it cannot lie in either K or L since these sets consist of rational numbers. It follows that $\theta' < \underline{\theta} = \overline{\theta} < \theta$ for any $\theta \in K$ and $\theta' \in L$, as desired. Next consider the case where $\underline{\theta} = \overline{\theta}$ is a rational number. Suppose by way of contradiction that $\overline{\theta} \in K$. Then there exists z^{k-1} satisfying $P(1|z^{k-1}) < P(0|z^{k-1})$ and $\overline{\theta} = \frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2}$. Since $\underline{\theta} = \overline{\theta}$, there is a sequence θ_i in L such that $\theta_i \to \overline{\theta}$. That is there exists a sequence in $L = \{\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} : P(1|x^{n-1}) > P(0|x^{n-1})\}$ that converges to $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2}$ for z^{k-1} that satisfies $P(1|z^{k-1}) < P(0|z^{k-1})$, violating Mean Reversion Continuity. Therefore $\overline{\theta} \notin K$, that is, $\overline{\theta} < \theta$ for any $\theta \in K$. A similar argument establishes $\theta' < \underline{\theta}$ for any $\theta' \in L$.

Turn to proving part (b). Suppose Strong Subjective Mean Reversion holds. Some immediate observations are as follows. If M is nonempty the axiom implies that it must be a singleton. In this case we write $M = \{\hat{\theta}\}$. If M is nonempty, the axiom also implies

$$\underline{\theta} \leq \hat{\theta} \leq \overline{\theta}.$$

Now suppose by way of contradiction that $\underline{\theta} < \overline{\theta}$. Since the set $\{\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} | x^{n-1} \in \Omega^{n-1} \text{ and } n > 1\}$ is dense in $(0,1)$, there exists a history z^{k-1} such that $\underline{\theta} < \frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} < \overline{\theta}$. If M is nonempty then we require in addition that $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} \neq \hat{\theta}$. It must be that $P(z^{k-1}1) \neq P(z^{k-1}0)$, since this is true anyway for all sequences when M is empty, and when M is nonempty then it follows from $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} \neq \hat{\theta}$. If $P(z^{k-1}1) < P(z^{k-1}0)$ then $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} \in K$ and we get the contradiction that $\overline{\theta} = \inf K \leq \frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} < \overline{\theta}$. Similarly, if $P(z^{k-1}1) > P(z^{k-1}0)$ then $\frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} \in L$ and we get the contradiction that $\underline{\theta} < \frac{\overline{z^{k-1}0}+\overline{z^{k-1}1}}{2} \leq \sup L = \underline{\theta}$. The desired conclusion follows. □

The next lemma establishes a useful connection between the notion of distance and the average of sample means.

Lemma 4. *For any $\theta \in (0,1)$, and x^{n-1},*

$$|\overline{x^{n-1}1} - \theta| \leq |\overline{x^{n-1}0} - \theta| \iff \theta \geq \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2}.$$

Proof. Let k denote the number of heads in x^{n-1}. Observe that

$$|\overline{x^{n-1}1} - \theta| \leq |\overline{x^{n-1}0} - \theta| \iff (\overline{x^{n-1}1} - \theta)^2 \leq (\overline{x^{n-1}0} - \theta)^2$$

$$\iff \overline{x^{n-1}1}^2 - 2\overline{x^{n-1}1}\theta \leq \overline{x^{n-1}0}^2 - 2\overline{x^{n-1}0}\theta$$
$$\iff \overline{x^{n-1}1}^2 - \overline{x^{n-1}0}^2 \leq 2\theta(\overline{x^{n-1}1} - \overline{x^{n-1}0})$$
$$\iff \frac{2k+1}{n^2} \leq \frac{2\theta}{n} \iff \overline{x^{n-1}0} + \frac{1}{2n} \leq \theta,$$

therefore establishing

$$|\overline{x^{n-1}1} - \theta| \leq |\overline{x^{n-1}0} - \theta| \iff \overline{x^{n-1}0} + \frac{1}{2n} \leq \theta.$$

Note however that $\overline{x^{n-1}1} - \overline{x^{n-1}0} = \frac{1}{n}$ and so

$$\overline{x^{n-1}0} + \frac{1}{2n} = \overline{x^{n-1}0} + \frac{\overline{x^{n-1}1} - \overline{x^{n-1}0}}{2} = \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2},$$

which completes the proof. □

Finally, we show that:

Lemma 5. *(i) Subjective Mean Reversion and Mean Reversion Continuity hold iff there exists $\theta^* \in [\underline{\theta}, \overline{\theta}]$ s.t. for all x^{n-1} and all x_n,*

$$|\overline{x^{n-1}x_n} - \theta^*| \leq |\overline{x^{n-1}(1-x_n)} - \theta^*| \implies P(x^{n-1}x_n) \geq P(x^{n-1}(1-x_n)).$$

(ii) Strong Subjective Mean Reversion and Mean Reversion Continuity holds iff for $\theta^* = \underline{\theta} = \overline{\theta}$, and for all x^{n-1} and all x_n,

$$|\overline{x^{n-1}x_n} - \theta^*| \leq |\overline{x^{n-1}(1-x_n)} - \theta^*| \iff P(x^{n-1}x_n) \geq P(x^{n-1}(1-x_n)).$$

Proof. Proof of (i): To establish sufficiency, assume that Subjective Mean Reversion holds. Note that if $P(1|x^{n-1}) > P(0|x^{n-1})$ then $\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} \in L$ and so $\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} \leq \sup L = \underline{\theta}$. Similarly if $P(1|x^{n-1}) < P(0|x^{n-1})$ then $\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} \in K$ and so $\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} \geq \inf K = \overline{\theta}$. Now consider two cases:

If $\underline{\theta} < \overline{\theta}$ (Lemma 3), then we obtain that for any $\theta^* \in (\underline{\theta}, \overline{\theta})$,

$$P(1|x^{n-1}) > P(0|x^{n-1}) \implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} \leq \underline{\theta} < \theta^* \implies |\overline{x^{n-1}1} - \theta^*| < |\overline{x^{n-1}0} - \theta^*|.$$

$$P(1|x^{n-1}) < P(0|x^{n-1}) \implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} \geq \overline{\theta} > \theta^* \implies |\overline{x^{n-1}1} - \theta^*| > |\overline{x^{n-1}0} - \theta^*|.$$

The contrapositive of these statements is the desired conclusion.

If on the other hand $\underline{\theta} = \overline{\theta}$, then by the implication of Mean Reversion Continuity in Lemma 3(a), we obtain the same conclusion by setting $\theta^* = \underline{\theta} = \overline{\theta}$. This proves sufficiency.

To prove necessity, it is immediate (given Lemma 4) that

$$P(1|x^{n-1}) > P(0|x^{n-1}) \text{ and } P(1|y^{m-1}) < P(0|y^{m-1})$$

$$\implies |\overline{x^{n-1}1} - \theta^*| < |\overline{x^{n-1}0} - \theta^*| \text{ and } |\overline{y^{m-1}0} - \theta^*| < |\overline{y^{m-1}1} - \theta^*|$$

$$\implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} < \theta^* < \frac{\overline{y^{m-1}0} + \overline{y^{m-1}1}}{2}$$

$$\implies \frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} < \frac{\overline{y^{m-1}0} + \overline{y^{m-1}1}}{2},$$

which establishes Subjective Mean Reversion. To establish Mean Reversion Continuity, consider a sequence in $L = \{\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2}$: for some x^{n-1} s.t. $P(1|x^{n-1}) > P(0|x^{n-1})\}$ that converges to a value $\frac{\overline{z^{k-1}0} + \overline{z^{k-1}1}}{2}$ generated by some z^{k-1} (the argument for a sequence in K is analogous). By Mean Reversion and Lemma 4, $\frac{\overline{x^{n-1}0} + \overline{x^{n-1}1}}{2} < \theta^*$ hold for all values in L and therefore it must be that $\frac{\overline{z^{k-1}0} + \overline{z^{k-1}1}}{2} \leq \theta^*$, and in turn, $P(1|z^{k-1}) \geq P(0|z^{k-1})$, as desired.

Proof of (ii): Consider sufficiency. Strong Subjective Mean Reversion and Mean Reversion Continuity hold iff $\underline{\theta} = \overline{\theta}$, set θ^* equal to this number. Since the claim in part (i) is implied, we need only show that

$$|\overline{x^{n-1}x_n} - \theta^*| < |\overline{x^{n-1}(1 - x_n)} - \theta^*| \implies P(x^{n-1}x_n) > P(x^{n-1}(1 - x_n)).$$

To prove the contrapositive, suppose $P(x^{n-1}x_n) \leq P(x^{n-1}(1 - x_n))$. We need to show that $|\overline{x^{n-1}x_n} - \theta^*| \geq |\overline{x^{n-1}(1-x_n)} - \theta^*|$. If this is strict then the proof in part (i) yields $|\overline{x^{n-1}x_n} - \theta^*| > |\overline{x^{n-1}(1-x_n)} - \theta^*|$, as desired. If $P(x^{n-1}x_n) = P(x^{n-1}(1-x_n))$ then M is not empty and by Lemma 3 we have $\frac{\overline{x^{n-1}0}+\overline{x^{n-1}1}}{2} = \hat{\theta} = \theta^*$ and so by Lemma 4 we obtain $|\overline{x^{n-1}x_n} - \theta^*| = |\overline{x^{n-1}(1-x_n)} - \theta^*|$, as desired.

The proof of necessity is analogous to the necessity argument in part (i) and therefore omitted. □

References

[1] Benjamin, D. (2019): "Errors in Probabilistic Reasoning and Judgment Biases." Chapter for the *Handbook of Behavioral Economics* (edited by Doug Bernheim, Stefano DellaVigna, and David Laibson). Elsevier Press.

[2] Chen, D., T. Moskowitz and K. Shue (2016): "Decision Making under the Gambler's Fallacy: Evidence from Asylum Judges, Loan Officers and Baseball Umpires", *Quarterly Journal of Economics* 131, pp. 1181-1242.

[3] Croson, Rachel and James Sundali (2005): "The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos", *Journal of Risk and Uncertainty* 30 (3), pp. 195-209.

[4] Jin, L. and C. Peng (2022): "Trading under the Law of Small Numbers", mimeo.

[5] Noor, J. and F. Payró (2026): "Belief in Mean Reversion and the Law of Small Numbers", *Theoretical Economics (forthcoming)*.

[6] Noor, J. and Y. Yang (2026): "Belief Updating without Trust", mimeo.

[7] Rabin, M. (2002), "Inference by Believers in the Law of Small Numbers", *Quarterly Journal of Economics* 117, pp. 775–816.

[8] Suetens, Sigrid, Claus Galbo-Jorgensen and Jean-Robert Tyran (2016): "Predicting Lotto Numbers: a Natural Experiment on the Gambler's Fal-

lacy and the Hot-Hand Fallacy", *Journal of the European Economic Association* 14, pp. 584-607.

[9] Terrell, Dek (1994): "A Test of the Gambler's Fallacy: Evidence from Pari-Mutuel Games", *Journal of Risk and Uncertainty* 8, pp. 309-317.

[10] Tversky, Amos and Daniel Kahneman (1971), "Belief in the Law of Small Numbers", *Psychological Bulletin* 76(2), pp. 105–110.

[11] Tversky, Amos and Daniel Kahneman (1974), "Judgement under Uncertainty: Heuristics and biases", *Science* 185, pp. 1124–1131.