THE REVIEW OF **ECONOMIC STUDIES**

Review of Economic Studies (2010) **77**, 730–778
0034-6527/10/00411011$02.00
doi: 10.1111/j.1467-937X.2009.00582.x

The Gambler's and Hot-Hand Fallacies: Theory and Applications

MATTHEW RABIN
University of California, Berkeley

and

DIMITRI VAYANOS
London School of Economics, CEPR and NBER

First version received September 2007; *final version accepted May* 2009 *(Eds.)*

We develop a model of the gambler's fallacy—the mistaken belief that random sequences should exhibit systematic reversals. We show that an individual who holds this belief and observes a sequence of signals can exaggerate the magnitude of changes in an underlying state but underestimate their duration. When the state is constant, and so signals are i.i.d., the individual can predict that long streaks of similar signals will continue—a hot-hand fallacy. When signals are serially correlated, the individual typically under-reacts to short streaks, over-reacts to longer ones, and under-reacts to very long ones. Our model has implications for a number of puzzles in finance, e.g. the active-fund and fund-flow puzzles, and the presence of momentum and reversal in asset returns.

1. INTRODUCTION

Many people fall under the spell of the "gambler's fallacy", expecting outcomes in random sequences to exhibit systematic reversals. When observing flips of a fair coin, for example, people believe that a streak of heads makes it more likely that the next flip will be a tail. The gambler's fallacy is commonly interpreted as deriving from a fallacious belief in the "law of small numbers" or "local representativeness": people believe that a small sample should resemble closely the underlying population, and hence believe that heads and tails should balance even in small samples. On the other hand, people also sometimes predict that random sequences will exhibit excessive persistence rather than reversals. While basketball fans believe that players have "hot hands", being more likely than average to make the next shot when currently on a hot streak, several studies have shown that no perceptible streaks justify such beliefs.[1]

At first blush, the hot-hand fallacy appears to directly contradict the gambler's fallacy because it involves belief in excessive persistence rather than reversals. Several researchers

1. The representativeness bias is perhaps the most commonly explored bias in judgement research. Section 2 reviews evidence on the gambler's fallacy, and a more extensive review can be found in Rabin (2002). For evidence on the hot-hand fallacy, see, for example, Gilovich, Vallone and Tversky (1985) and Tversky and Gilovich (1989*a*,*b*). See also Camerer (1989) who shows that betting markets for basketball games exhibit a small hot-hand bias.

have, however, suggested that the two fallacies might be related, with the hot-hand fallacy arising as a consequence of the gambler's fallacy.[2] Suppose that an investor prone to the gambler's fallacy observes the performance of a mutual fund, which can depend on the manager's ability and luck. Convinced that luck should reverse, the investor underestimates the likelihood that a manager of average ability will exhibit a streak of above- or below-average performances. Following good or bad streaks, therefore, the investor over-infers that the current manager is above or below average, and so in turn predicts continuation of unusual performances.

This paper develops a model to examine the link between the gambler's fallacy and the hot-hand fallacy, as well as the broader implications of the fallacies for people's predictions and actions in economic and financial settings. In our model, an individual observes a sequence of signals that depend on an unobservable underlying state. We show that because of the gambler's fallacy, the individual is prone to exaggerate the magnitude of changes in the state, but underestimate their duration. We characterize the individual's predictions following streaks of similar signals, and determine when a hot-hand fallacy can arise. Our model has implications for a number of puzzles in finance, e.g. the active-fund and fund-flow puzzles, and the presence of momentum and reversal in asset returns.

While providing extensive motivation and elaboration in Section 2, we now present the model itself in full. An agent observes a sequence of signals whose probability distribution depends on an underlying state. The signal s_t in period $t = 1, 2, \ldots$ is

$$s_t = \theta_t + \epsilon_t, \tag{1}$$

where θ_t is the state and ϵ_t an i.i.d. normal shock with mean zero and variance $\sigma_\epsilon^2 > 0$. The state evolves according to the auto-regressive process

$$\theta_t = \rho\theta_{t-1} + (1-\rho)(\mu + \eta_t), \tag{2}$$

where $\rho \in [0, 1]$ is the persistence parameter, μ the long-run mean, and η_t an i.i.d. normal shock with mean zero, variance σ_η^2, and independent of ϵ_t. As an example that we shall return to often, consider a mutual fund run by a team of managers. We interpret the signal as the fund's return, the state as the managers' average ability, and the shock ϵ_t as the managers' luck. Assuming that the ability of any given manager is constant over time, we interpret $1-\rho$ as the rate of managerial turnover, and σ_η^2 as the dispersion in ability across managers.[3]

We model the gambler's fallacy as the mistaken belief that the sequence $\{\epsilon_t\}_{t\geq 1}$ is not i.i.d., but rather exhibits reversals: according to the agent,

$$\epsilon_t = \omega_t - \alpha_\rho \sum_{k=0}^{\infty} \delta_\rho^k \epsilon_{t-1-k}, \tag{3}$$

where the sequence $\{\omega_t\}_{t\geq 1}$ is i.i.d. normal with mean zero and variance σ_ω^2, and $(\alpha_\rho, \delta_\rho)$ are parameters in $[0, 1)$ that can depend on ρ.[4] Consistent with the gambler's fallacy, the agent

2. See, for example, Camerer (1989) and Rabin (2002). The causal link between the gambler's fallacy and the hot-hand fallacy is a common intuition in psychology. Some suggestive evidence comes from an experiment by Edwards (1961), in which subjects observe a very long binary series and are given no information about the generating process. Subjects seem, by the evolution of their predictions over time, to come to believe in a hot hand. Since the actual generating process is independent and identically distributed (i.i.d.), this is suggestive that a source of the hot hand is the perception of too many long streaks.

3. Alternatively, we can assume that a fraction $1-\rho$ of existing managers get a new ability "draw" in any given period. Ability could be time-varying if, for example, managers' expertise is best suited to specific market conditions. We use the managerial-turnover interpretation because it is easier for exposition.

4. We set $\epsilon_t = 0$ for $t \leq 0$, so that all terms in the infinite sum are well defined.

believes that high realizations of $\epsilon_{t'}$ in period $t' < t$ make a low realization more likely in period t. The parameter α_ρ measures the strength of the effect, while δ_ρ measures the relative influence of realizations in the recent and more distant past. We mainly focus on the case where $(\alpha_\rho, \delta_\rho)$ depend linearly on ρ, i.e. $(\alpha_\rho, \delta_\rho) = (\alpha\rho, \delta\rho)$ for $\alpha, \delta \in [0, 1)$. Section 2 motivates linearity based on the assumption that people expect reversals only when all outcomes in a random sequence are drawn from the same distribution, e.g. the performances of a fund manager whose ability is constant over time. This assumption rules out belief in mean reversion across distributions; e.g. a good performance by one manager does not make another manager due for a bad performance next period. If managers turn over frequently (ρ is small) therefore, the gambler's fallacy has a small effect, consistent with the linear specification. Section 2 discusses alternative specifications for $(\alpha_\rho, \delta_\rho)$ and the link with the evidence. Appendix A shows that the agent's error patterns in predicting the signals are very similar across specifications.

Section 3 examines how the agent uses the sequence of past signals to make inferences about the underlying parameters and to predict future signals. We assume that the agent infers as a fully rational Bayesian and fully understands the structure of his environment, except for a mistaken and dogmatic belief that $\alpha > 0$. From observing the signals, the agent learns about the underlying state θ_t, and possibly about the parameters of his model $(\sigma_\eta^2, \rho, \sigma_\omega^2, \mu)$ if these are unknown.[5]

In the benchmark case where the agent is certain about all model parameters, his inference can be treated using standard tools of recursive (Kalman) filtering, where the gambler's fallacy essentially expands the state vector to include not only the state θ_t but also a statistic of past luck realizations. If instead the agent is uncertain about parameters, recursive filtering can be used to evaluate the likelihood of signals conditional on parameters. An appropriate version of the law of large numbers (LLN) implies that after observing many signals, the agent converges with probability 1 to parameter values that maximize a limit likelihood. While the likelihood for $\alpha = 0$ is maximized for limit posteriors corresponding to the true parameter values, the agent's abiding belief that $\alpha > 0$ leads him generally to false limit posteriors. Identifying when and how these limit beliefs are wrong is the crux of our analysis.[6]

Section 4 considers the case where signals are i.i.d. because $\sigma_\eta^2 = 0$.[7] If the agent is initially uncertain about parameters and does not rule out any possible value, then he converges to the belief that $\rho = 0$. Under this belief, he predicts the signals correctly as i.i.d., despite the gambler's fallacy. The intuition is that he views each signal as drawn from a new distribution;

5. When learning about model parameters, the agent is limited to models satisfying equations (1) and (2). Since the true model belongs to that set, considering models outside the set does not change the limit outcome of rational learning. An agent prone to the gambler's fallacy, however, might be able to predict the signals more accurately using an incorrect model, as the two forms of error might offset each other. Characterizing the incorrect model that the agent converges to is central to our analysis, but we restrict such a model to satisfy equations (1) and (2). A model not satisfying equations (1) and (2) can help the agent make better predictions when signals are serially correlated. See footnote 24.

6. Our analysis has similarities to model mis-specification in econometrics. Consider, for example, the classic omitted-variables problem, where the true model is $Y = \alpha + \beta_1 X_1 + \beta_2 X_2 + \epsilon$ but an econometrician estimates $Y = \alpha + \beta_1 X_1 + \epsilon$. The omission of the variable X_2 can be interpreted as a dogmatic belief that $\beta_2 = 0$. When this belief is incorrect because $\beta_2 \neq 0$, the econometrician's estimate of β_1 is biased and inconsistent (e.g. Greene, 2008, Ch. 7). Mis-specification in econometric models typically arises because the true model is more complicated than the econometrician's model. In our setting the agent's model is the more complicated because it assumes negative correlation where there is none.

7. When $\sigma_\eta^2 = 0$, the shocks η_t are equal to zero, and therefore equation (2) implies that in steady state θ_t is constant.

e.g. new managers run the fund in each period. Therefore, his belief that any given manager's performance exhibits mean reversion has no effect.

We next assume that the agent knows on prior grounds that $\rho > 0$; e.g. is aware that managers stay in the fund for more than one period. Ironically, the agent's correct belief $\rho > 0$ can lead him astray. This is because he cannot converge to the belief $\rho = 0$, which, while incorrect, enables him to predict the signals correctly. Instead, he converges to the smallest value of ρ to which he gives positive probability. He also converges to a positive value of σ_η^2, believing falsely that managers differ in ability, so that (given turnover) there is variation over time in average ability. This belief helps him to explain the incidence of streaks despite the gambler's fallacy: a streak of high returns, for example, can be readily explained through the belief that good managers might have joined the fund recently. Of course, the agent thinks that the streak might also have been due to luck, and expects a reversal. We show that the expectation of a reversal dominates for short streaks, but because reversals that do not happen make the agent more confident the managers have changed, he expects long streaks to continue. Thus, predictions following long streaks exhibit a hot-hand fallacy.[8]

Section 5 relaxes the assumption that $\sigma_\eta^2 = 0$, to consider the case where signals are serially correlated. As in the i.i.d. case, the agent underestimates ρ and overestimates the variance $(1 - \rho)^2\sigma_\eta^2$ of the shocks to the state. He does not converge, however, all the way to $\rho = 0$ because he must account for the signals' serial correlation. Because he views shocks to the state as overly large in magnitude, he treats signals as very informative, and tends to overreact to streaks. For very long streaks, however, there is under-reaction because the agent's underestimation of ρ means that he views the information learned from the signals as overly short-lived. Under-reaction also tends to occur following short streaks because of the basic gambler's fallacy intuition.

In summary, Sections 4 and 5 confirm the oft-conjectured link from the gambler's to the hot-hand fallacy, and generate novel predictions which can be tested in experimental or field settings. We summarize the predictions at the end of Section 5.

We conclude this paper in Section 6 by exploring finance applications. One application concerns the belief in financial expertise. Suppose that returns on traded assets are i.i.d., and while the agent does not rule out that expected returns are constant, he is confident that any shocks to expected returns should last for more than one period ($\rho > 0$). He then ends up believing that returns are predictable on the basis of past history, and exaggerates the value of financial experts if he believes that the experts' advantage derives from observing market information. This could help explain the active-fund puzzle, namely why people invest in actively managed funds in spite of the evidence that these funds underperform their passively managed counterparts. In a related application we show that the agent not only exaggerates the value added by active managers but also believes that this value varies over time. This could help explain the fund-flow puzzle, namely that flows into mutual funds are positively correlated with the funds' lagged returns, and yet lagged returns do not appear to predict future returns. Our model could also speak to other finance puzzles, such as the presence of momentum and reversals in asset returns and the value premium.

Our work is related to Rabin's (2002) model of the law of small numbers. In Rabin, an agent draws balls from an urn with replacement but believes that replacement occurs only every odd period. Thus, the agent overestimates the probability that the ball drawn in an even period

8. We are implicitly defining the hot-hand fallacy as a belief in the continuation of streaks. An alternative and closely related definition involves the agent's assessed auto-correlation of the signals. See footnote 22 and the preceding text.

is of a different colour than the one drawn in the previous period. Because of replacement, the composition of the urn remains constant over time. Thus, the underlying state is constant, which corresponds to $\rho = 1$ in our model. We instead allow ρ to take any value in $[0, 1)$ and show that the agent's inferences about ρ significantly affect his predictions of the signals. Additionally, because $\rho < 1$, we can study inference in a stochastic steady state and conduct a true dynamic analysis of the hot-hand fallacy (showing how the effects of good and bad streaks alternate over time). Because of the steady state and the normal-linear structure, our model is more tractable than Rabin's. In particular, we characterize fully predictions after streaks of signals, while Rabin can do so only with numerical examples and for short streaks. The finance applications in Section 6 illustrate further our model's tractability and applicability, while also leading to novel insights such as the link between the gambler's fallacy and belief in outperformance of active funds.

Our work is also related to the theory of momentum and reversals of Barberis, Shleifer and Vishny (1998). In BSV, investors do not realize that innovations to a company's earnings are i.i.d. Rather, they believe them to be drawn either from a regime with excess reversals or from one with excess streaks. If the reversal regime is the more common, the stock price under-reacts to short streaks because investors expect a reversal. The price over-reacts, however, to longer streaks because investors interpret them as sign of a switch to the streak regime. This can generate short-run momentum and long-run reversals in stock returns, consistent with the empirical evidence, surveyed in BSV. It can also generate a value premium because stocks with long positive (negative) streaks become overvalued (undervalued) relative to earnings and yield low (high) expected returns. Our model has similar implications because in i.i.d. settings the agent can expect short streaks to reverse and long streaks to continue. But while the implications are similar, our approach is different. BSV provide a psychological foundation for their assumptions by appealing to a combination of biases: the conservatism bias for the reversal regime and the representativeness bias for the streak regime. Our model, by contrast, not only derives such biases from the single underlying bias of the gambler's fallacy, but in doing so provides predictions as to which biases are likely to appear in different informational settings.[9]

2. MOTIVATION FOR THE MODEL

Our model is fully described by equations (1) to (3) presented in the Introduction. In this section, we motivate the model by drawing the connection with the experimental evidence on the gambler's fallacy. A difficulty with using this evidence is that most experiments concern sequences that are binary and i.i.d., such as coin flips. Our goal, by contrast, is to explore the implications of the gambler's fallacy in richer settings. In particular, we need to consider non-i.i.d. settings since the hot-hand fallacy involves a belief that the world is non-i.i.d. The experimental evidence gives little direct guidance on how the gambler's fallacy would manifest itself in non-binary, non-i.i.d. settings. In this section, however, we argue that our model represents a natural extrapolation of the gambler's fallacy "logic" to such settings. Of course, any such extrapolation has an element of speculativeness. But, if nothing else, our specification of the gambler's fallacy in the new settings can be viewed as a working hypothesis about the

9. Even in settings where error patterns in our model resemble those in BSV, there are important differences. For example, the agent's expectation of a reversal can increase with streak length for short streaks, while in BSV it unambiguously decreases (as it does in Rabin (2002) because the expectation of a reversal lasts for only one period). This increasing pattern is a key finding of an experimental study by Asparouhova, Hertzel and Lemmon (2008). See also Bloomfield and Hales (2002).

broader empirical nature of the phenomenon that both highlights features of the phenomenon that seem to matter and generates testable predictions for experimental and field research.

Experiments documenting the gambler's fallacy are mainly of three types: *production* tasks, where subjects are asked to produce sequences that look to them like random sequences of coin flips; *recognition* tasks, where subjects are asked to identify which sequences look like coin flips; and *prediction* tasks, where subjects are asked to predict the next outcome in coin-flip sequences. In all types of experiments, the typical subject identifies a switching (i.e. reversal) rate greater than 50% to be indicative of random coin flips.[10] The most carefully reported data for our purposes comes from the production-task study of Rapoport and Budescu (1997). Using their Table 7, we estimate in Table 1 the subjects' assessed probability that the next flip of a coin will be heads given the last three flips.[11]

According to Table 1, the average effect of changing the most recent flip from heads (H) to tails (T) is to raise the probability that the next flip will be H from 40.1 ($= \frac{30\%+38\%+41.2\%+51.3\%}{4}$) to 59.9%, i.e. an increase of 19.8%. This corresponds well to the general stylized fact in the literature that subjects tend to view randomness in coin-flip sequences as corresponding to a switching rate of 60% rather than 50%. Table 1 also shows that the effect of the gambler's fallacy is not limited to the most recent flip. For example, the average effect of changing the second most recent flip from H to T is to raise the probability of H from 43.9 to 56.1%, i.e. an increase of 12.2%. The average effect of changing the third most recent flip from H to T is to raise the probability of H from 45.5 to 54.5%, i.e. an increase of 9%.

10. See Bar-Hillel and Wagenaar (1991) for a review of the literature, and Rapoport and Budescu (1992, 1997) and Budescu and Rapoport (1994) for more recent studies. The experimental evidence has some shortcomings. For example, most prediction-task studies report the fraction of subjects predicting a switch but not the subjects' assessed probability of a switch. Thus, it could be that the vast majority of subjects predict a switch, and yet their assessed probability is only marginally larger than 50%. Even worse, the probability could be exactly 50%, since under that probability subjects are indifferent as to their prediction.

Some prediction-task studies attempt to measure assessed probabilities more accurately. For example, Gold and Hester (2008) find evidence in support of the gambler's fallacy in settings where subjects are given a choice between a sure payoff and a random payoff contingent on a specific coin outcome. Supporting evidence also comes from settings outside the laboratory. For example, Clotfelter and Cook (1993) and Terrell (1994) study pari-mutuel lotteries, where the winnings from a number are shared among all people betting on that number. They find that people avoid systematically to bet on numbers that won recently. This is a strict mistake because the numbers with the fewest bets are those with the largest expected winnings. See also Metzger (1984), Terrell and Farmer (1996), and Terrell (1998) for evidence from horse and dog races, and Croson and Sundali (2005) for evidence from casino betting.

11. Rapoport and Budescu report relative frequencies of short sequences of heads (H) and tails (T) within the larger sequences (of 150 elements) produced by the subjects. We consider frequencies of four-element sequences, and average the two "observed" columns. The first four lines of Table 1 are derived as follows:

$$\text{Line 1} = \frac{f(\text{HHHH})}{f(\text{HHHH}) + f(\text{HHHT})},$$
$$\text{Line 2} = \frac{f(\text{THHH})}{f(\text{THHH}) + f(\text{HTTH})},$$
$$\text{Line 3} = \frac{f(\text{HTHH})}{f(\text{HTHH}) + f(\text{HTHT})},$$
$$\text{Line 4} = \frac{f(\text{HHTH})}{f(\text{HHTH}) + f(\text{HHTT})}.$$

(The denominator in Line 2 involves HTTH rather than the equivalent sequence THHT, derived by reversing H and T, because Rapoport and Budescu group equivalent sequences together.) The last four lines of Table 1 are simply transformations of the first four lines, derived by reversing H and T. While our estimates are derived from relative frequencies, we believe that they are good measures of subjects' assessed probabilities. For example, a subject believing that HHH should be followed by H with 30% probability could be choosing H after HHH 30% of the time when constructing a random sequence.

TABLE 1

Assessed probability that the next flip of a coin will be heads (H) given the last three flips being heads or tails (T). Based on Rapoport and Budescu (1997, Table 7, p. 613)

3rd-to-last	2nd-to-last	Very last	Prob. next will be H (%)
H	H	H	30.0
T	H	H	38.0
H	T	H	41.2
H	H	T	48.7
H	T	T	62.0
T	H	T	58.8
T	T	H	51.3
T	T	T	70.0

How would a believer in the gambler's fallacy, exhibiting behaviour such as in Table 1, form predictions in non-binary, non-i.i.d. settings? Our extrapolation approach consists in viewing the richer settings as combinations of coins. We first consider settings that are non-binary but i.i.d. Suppose that in each period a large number of coins are flipped simultaneously and the agent observes the sum of the flips, where we set H $= 1$ and T $= -1$. For example, with 100 coins, the agent observes a signal between 100 and -100, and a signal of 10 means that 55 coins came H and 45 came T. Suppose that the agent applies his gambler's fallacy reasoning to each individual coin (i.e. to 100 separate coin-flip sequences), and his beliefs are as in Table 1. Then, after a signal of 10, he assumes that the 55 H coins have probability 40.1% to come H again, while the 45 T coins have probability 59.9% to switch to H. Thus, he expects on average $40.1\% \times 55 + 59.9\% \times 45 = 49.01$ coins to come H, and this translates to an expectation of $49.01 - (100 - 49.01) = -1.98$ for the next signal.

The "multiple-coin" story shares many of our model's key features. To explain why, we specialize the model to i.i.d. signals, taking the state θ_t to be known and constant over time. We generate a constant state by setting $\rho = 1$ in equation (2). For simplicity, we normalize the constant value of the state to zero. For $\rho = 1$, equation (3) becomes

$$\epsilon_t = \omega_t - \alpha \sum_{k=0}^{\infty} \delta^k \epsilon_{t-1-k}, \tag{4}$$

where $(\alpha, \delta) \equiv (\alpha_1, \delta_1)$. When the state is equal to zero, equation (1) becomes $s_t = \epsilon_t$. Substituting into equation (4) and taking expectations conditional on Period $t - 1$, we find

$$E_{t-1}(s_t) = -\alpha \sum_{k=0}^{\infty} \delta^k s_{t-1-k}. \tag{5}$$

Comparing equation (5) with the predictions of the multiple-coin story, we can calibrate the parameters (α, δ) and test our specification of the gambler's fallacy. Suppose that $s_{t-1} = 10$ and that signals in prior periods are equal to their average of zero. Equation (5) then implies that $E_{t-1}(s_t) = -10\alpha$. According to the multiple-coin story, $E_{t-1}(s_t)$ should be -1.98, which yields $\alpha = 0.198$. To calibrate δ, we repeat the exercise for s_{t-2} (setting $s_{t-2} = 10$ and $s_{t-i} = 0$ for $i = 1$ and $i > 2$) and then again for s_{t-3}. Using the data in Table 1, we find $\alpha\delta = 0.122$ and $\alpha\delta^2 = 0.09$. Thus, the decay in influence between the most and second most recent signal is $0.122/0.198 = 0.62$, while the decay between the second and third most recent signal is $0.09/0.122 = 0.74$. The two rates are not identical as our model assumes, but are quite close. Thus, our geometric-decay specification seems reasonable, and we can take $\alpha = 0.2$ and $\delta = 0.7$ as a plausible calibration. Motivated by the evidence, we impose from now on the restriction $\delta > \alpha$, which simplifies our analysis.

Several other features of our specification deserve comment. One is normality: since ω_t is normal, equation (4) implies that the distribution of $s_t = \epsilon_t$ conditional on period $t-1$ is normal. The multiple-coin story also generates approximate normality if we take the number of coins to be large. A second feature is linearity: if we double s_{t-1} in equation (5), holding other signals to zero, then $E_{t-1}(s_t)$ doubles. The multiple-coin story shares this feature: a signal of 20 means that 60 coins came H and 40 came T, and this doubles the expectation of the next signal. A third feature is additivity: according to equation (5), the effect of each signal on $E_{t-1}(s_t)$ is independent of the other signals. Table 1 generates some support for additivity. For example, changing the most recent flip from H to T increases the probability of H by 20.8% when the second and third most recent flips are identical (HH or TT) and by 18.7% when they differ (HT or TH). Thus, in this experiment the effect of the most recent flip depends only weakly on prior flips.

We next extend our approach to non-i.i.d. settings. Suppose that the signal the agent is observing in each period is the sum of a large number of independent coin flips, but where coins differ in their probabilities of H and T, and are replaced over time randomly by new coins. Signals are thus serially correlated: they tend to be high at times when the replacement process brings many new coins biased towards H, and vice versa. If the agent applies his gambler's fallacy reasoning to each individual coin, then this will generate a gambler's fallacy for the signals. The strength of the latter fallacy will depend on the extent to which the agent believes in mean reversion across coins: if a coin comes H, does this make its replacement coin more likely to come T? For example, in the extreme case where the agent does not believe in mean reversion across coins and each coin is replaced after one period, the agent will not exhibit the gambler's fallacy for the signals.

There seems to be relatively little evidence on the extent to which people believe in mean reversion across random devices (e.g. coins). Because the evidence suggests that belief in mean reversion is moderated but not eliminated when moving across devices, we consider the two extreme cases both where it is eliminated and where it is not moderated.[12] In the next two paragraphs we show that in the former case the gambler's fallacy for the signals takes the form of equation (3), with $(\alpha_\rho, \delta_\rho)$ linear functions of ρ. We use this linear specification in Sections 3–6. In Appendix A we study the latter case, where $(\alpha_\rho, \delta_\rho)$ are independent of ρ. The agent's error patterns in predicting the signals are very similar across the two cases.[13]

To derive the linear specification, we shift from coins to normal distributions. Consider a mutual fund that consists of a continuum with mass one of managers, and suppose that a random fraction $1-\rho$ of managers are replaced by new ones in each period. Suppose that the fund's

12. Evidence that people believe in mean reversion across random devices comes from horse and dog races. Metzger (1994) shows that people bet on the favourite horse significantly less when the favourites have won the previous two races (even though the horses themselves are different animals). Terrell and Farmer (1996) and Terrell (1998) show that people are less likely to bet on repeat winners by post position: if, e.g. the dog in post-position 3 won a race, the (different) dog in post-position 3 in the next race is significantly underbet. Gold and Hester (2008) find that belief in mean reversion is moderated when moving across random devices. They conduct experiments where subjects are told the recent flips of a coin, and are given a choice with payoffs contingent on the next flip of the same or of a new coin. Subjects' choices reveal a strong prediction of reversal for the old coin, but a much weaker prediction for the new coin.

13. One could envision alternative specifications for $(\alpha_\rho, \delta_\rho)$. For example, α_ρ could be assumed decreasing in ρ, i.e. if the agent believes that the state is less persistent, he expects more reversals conditional on the state. Appendix A extends some of our results to a general class of specifications.

return s_t is an average of returns attributable to each manager, and a manager's return is the sum of ability and luck, both normally distributed. Ability is constant over time for a given manager, while luck is i.i.d. Thus, a manager's returns are i.i.d. conditional on ability, and the manager can be viewed as a "coin" with the probability of H and T corresponding to ability. To ensure that aggregate variables are stochastic despite the continuum assumption, we assume that ability and luck are identical within the cohort of managers who enter the fund in a given period.[14]

We next show that if the agent applies his gambler's fallacy reasoning to each manager, per our specification in equation (4) for $\rho = 1$, and rules out mean reversion across managers, then this generates a gambler's fallacy for fund returns, per our specification in equation (3) with $(\alpha_\rho, \delta_\rho) = (\alpha\rho, \delta\rho)$. Denoting by $\epsilon_{t,t'}$ the luck in period t of the cohort entering in period $t' \leq t$, we can write equation (4) for a manager in that cohort as

$$\epsilon_{t,t'} = \omega_{t,t'} - \alpha \sum_{k=0}^{\infty} \delta^k \epsilon_{t-1-k,t'}, \tag{6}$$

where $\{\omega_{t,t'}\}_{t \geq t' \geq 0}$ is an i.i.d. sequence and $\epsilon_{t'',t'} \equiv 0$ for $t'' < t'$. To aggregate equation (6) for the fund, we note that in period t the average luck ϵ_t of all managers is

$$\epsilon_t = (1 - \rho) \sum_{t' \leq t} \rho^{t-t'} \epsilon_{t,t'}, \tag{7}$$

since $(1 - \rho)\rho^{t-t'}$ managers from the cohort entering in period t' are still in the fund. Combining equations (6) and (7) and setting $\omega_t \equiv (1 - \rho) \sum_{t' \leq t} \rho^{t-t'} \omega_{t,t'}$, we find equation (3) with $(\alpha_\rho, \delta_\rho) = (\alpha\rho, \delta\rho)$. Since $(\alpha_\rho, \delta_\rho)$ are linear in ρ, the gambler's fallacy is weaker the larger the managerial turnover is. Intuitively, with large turnover, the agent's belief that a given manager's performance should average out over multiple periods has little effect.

We close this section by highlighting an additional aspect of our model: the gambler's fallacy applies to the sequence $\{\epsilon_t\}_{t \geq 1}$ that generates the signals given the state, but not to the sequence $\{\eta_t\}_{t \geq 1}$ that generates the state. For example, the agent expects that a mutual fund manager who overperforms in one period is more likely to underperform in the next. He does not expect, however, that if high-ability managers join the fund in one period, low-ability managers are more likely to follow. We rule out the latter form of the gambler's fallacy mainly for simplicity. In Appendix B we show that our model and solution method generalize to the case where the agent believes that the sequence $\{\eta_t\}_{t \geq 1}$ exhibits reversals, and our main results carry through.

3. INFERENCE—GENERAL RESULTS

In this section we formulate the agent's inference problem, and establish some general results that serve as the basis for the more specific results of Sections 4–6. The inference problem consists in using the signals to learn about the underlying state θ_t and possibly about the parameters of the model. The agent's model is characterized by the variance $(1 - \rho)^2\sigma_\eta^2$ of the shocks to the state, the persistence ρ, the variance σ_ω^2 of the shocks affecting the signal noise, the long-run mean μ, and the parameters (α, δ) of the gambler's fallacy. We assume that the agent does not question his belief in the gambler's fallacy, i.e. has a dogmatic point prior on

14. The intuition behind the example would be the same, but more plausible, with a single manager in each period who is replaced by a new one with Poisson probability $1 - \rho$. We assume a continuum because this preserves normality. The assumption that all managers in a cohort have the same ability and luck can be motivated in reference to the single-manager setting.

(α, δ). He can, however, learn about the other parameters. From now on, we reserve the notation $(\sigma_\eta^2, \rho, \sigma_\omega^2, \mu)$ for the true parameter values, and denote generic values by $(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu})$. Thus, the agent can learn about the parameter vector $\tilde{p} \equiv (\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu})$.

3.1. *No parameter uncertainty*

We start with the case where the agent is certain that the parameter vector takes a specific value $\tilde{p}$. This case is relatively simple and serves as an input for the parameter-uncertainty case. The agent's inference problem can be formulated as one of recursive (Kalman) filtering. Recursive filtering is a technique for solving inference problems where (i) inference concerns a "state vector" evolving according to a stochastic process, (ii) a noisy signal of the state vector is observed in each period, and (iii) the stochastic structure is linear and normal. Because of normality, the agent's posterior distribution is fully characterized by its mean and variance, and the output of recursive filtering consists of these quantities.[15]

To formulate the recursive-filtering problem, we must define the state vector, the equation according to which the state vector evolves, and the equation linking the state vector to the signal. The state vector must include not only the state θ_t, but also some measure of the past realizations of luck since according to the agent luck reverses predictably. It turns out that all past luck realizations can be condensed into a one-dimensional statistic. This statistic can be appended to the state θ_t, and therefore, recursive filtering can be used even in the presence of the gambler's fallacy. We define the state vector as

$$x_t \equiv \left[\theta_t - \tilde{\mu}, \epsilon_t^\delta\right]',$$

where the statistic of past luck realizations is

$$\epsilon_t^\delta \equiv \sum_{k=0}^{\infty} \delta_{\tilde{\rho}}^k \epsilon_{t-k},$$

and v' denotes the transpose of the vector v. Equations (2) and (3) imply that the state vector evolves according to

$$x_t = \tilde{A} x_{t-1} + w_t, \tag{8}$$

where

$$\tilde{A} \equiv \begin{bmatrix} \tilde{\rho} & 0 \\ 0 & \delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}} \end{bmatrix}$$

and

$$w_t \equiv [(1-\tilde{\rho})\eta_t, \omega_t]'.$$

Equations (1)–(3) imply that the signal is related to the state vector through

$$s_t = \tilde{\mu} + \tilde{C} x_{t-1} + v_t, \tag{9}$$

15. For textbooks on recursive filtering see, for example, Anderson and Moore (1979) and Balakrishnan (1987). We are using the somewhat cumbersome term "state vector" because we are reserving the term "state" for θ_t, and the two concepts differ in our model.

where

$$\tilde{C} \equiv [\tilde{\rho}, -\alpha_{\tilde{\rho}}]$$

and $v_t \equiv (1-\tilde{\rho})\eta_t + \omega_t$. To start the recursion, we must specify the agent's prior beliefs for the initial state x_0. We denote the mean and variance of θ_0 by $\overline{\theta}_0$ and $\sigma^2_{\theta,0}$, respectively. Since $\epsilon_t = 0$ for $t \leq 0$, the mean and variance of ϵ^δ_0 are zero. Proposition 1 determines the agent's beliefs about the state in period t, conditional on the history of signals $\mathcal{H}_t \equiv \{s_{t'}\}_{t'=1,\ldots,t}$ up to that period.

Proposition 1. *Conditional on $\mathcal{H}_t$, x_t is normal with mean $\overline{x}_t(\tilde{p})$ given recursively by*

$$\overline{x}_t(\tilde{p}) = \tilde{A}\overline{x}_{t-1}(\tilde{p}) + \tilde{G}_t\left[s_t - \tilde{\mu} - \tilde{C}\overline{x}_{t-1}(\tilde{p})\right], \qquad \overline{x}_0(\tilde{p}) = [\overline{\theta}_0 - \tilde{\mu}, 0]', \tag{10}$$

and covariance matrix $\tilde{\Sigma}_t$ given recursively by

$$\tilde{\Sigma}_t = \tilde{A}\tilde{\Sigma}_{t-1}\tilde{A}' - \left(\tilde{C}\tilde{\Sigma}_{t-1}\tilde{C}' + \tilde{V}\right)\tilde{G}_t\tilde{G}_t' + \tilde{W}, \qquad \tilde{\Sigma}_0 = \begin{bmatrix} \sigma^2_{\theta,0} & 0 \\ 0 & 0 \end{bmatrix}, \tag{11}$$

where

$$\tilde{G}_t \equiv \frac{1}{\tilde{C}\tilde{\Sigma}_{t-1}\tilde{C}' + \tilde{V}}\left(\tilde{A}\tilde{\Sigma}_{t-1}\tilde{C}' + \tilde{U}\right), \tag{12}$$

$\tilde{V} \equiv \tilde{E}(v_t^2)$, $\tilde{W} \equiv \tilde{E}(w_t w_t')$, $\tilde{U} \equiv \tilde{E}(v_t w_t)$, *and $\tilde{E}$ is the agent's expectation operator.*

The agent's conditional mean evolves according to equation (10). This equation is derived by regressing the state vector x_t on the signal s_t, conditional on the history of signals $\mathcal{H}_{t-1}$ up to period $t-1$. The conditional mean $\overline{x}_t(\tilde{p})$ coming out of this regression is the sum of two terms. The first term, $\tilde{A}\overline{x}_{t-1}(\tilde{p})$, is the mean of x_t conditional on $\mathcal{H}_{t-1}$. The second term reflects learning in period t, and is the product of a regression coefficient $\tilde{G}_t$ times the agent's "surprise" in period t, defined as the difference between s_t and its mean conditional on $\mathcal{H}_{t-1}$. The coefficient $\tilde{G}_t$ is equal to the ratio of the covariance between s_t and x_t over the variance of s_t, where both moments are evaluated conditional on $\mathcal{H}_{t-1}$. The agent's conditional variance of the state vector evolves according to equation (11). Because of normality, this equation does not depend on the history of signals, and therefore conditional variances and covariances are deterministic. The history of signals affects only conditional means, but we do not make this dependence explicit for notational simplicity. Proposition 2 shows that when t goes to ∞, the conditional variance converges to a limit that is independent of the initial value $\tilde{\Sigma}_0$.

Proposition 2. *$Lim_{t\to\infty}\tilde{\Sigma}_t = \tilde{\Sigma}$, where $\tilde{\Sigma}$ is the unique solution in the set of positive matrices of*

$$\tilde{\Sigma} = \tilde{A}\tilde{\Sigma}\tilde{A}' - \frac{1}{\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V}}\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right)\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right)' + \tilde{W}. \tag{13}$$

Proposition 2 implies that there is convergence to a steady state where the conditional variance $\tilde{\Sigma}_t$ is equal to the constant $\tilde{\Sigma}$, the regression coefficient $\tilde{G}_t$ is equal to the constant

$$\tilde{G} \equiv \frac{1}{\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V}}\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right), \tag{14}$$

and the conditional mean of the state vector x_t evolves according to a linear equation with constant coefficients. This steady state plays an important role in our analysis: it is also the limit in the case of parameter uncertainty because the agent eventually becomes certain about the parameter values.

3.2. *Parameter uncertainty*

We next allow the agent to be uncertain about the parameters of his model. Parameter uncertainty is a natural assumption in many settings. For example, the agent might be uncertain about the extent to which fund managers differ in ability (σ_η^2) or turn over (ρ).

Because parameter uncertainty eliminates the normality that is necessary for recursive filtering, the agent's inference problem threatens to be less tractable. Recursive filtering can, however, be used as part of a two-stage procedure. In a first stage, we fix each model parameter to a given value and compute the likelihood of a history of signals conditional on these values. Because the conditional probability distribution is normal, the likelihood can be computed using the recursive-filtering formulas of Section 3.1. In a second stage, we combine the likelihood with the agent's prior beliefs, through Bayes' law, and determine the agent's posteriors on the parameters. We show, in particular, that the agent's posteriors in the limit when t goes to ∞ can be derived by maximizing a limit likelihood over all possible parameter values.

We describe the agent's prior beliefs over parameter vectors by a probability measure π_0 and denote by P the closed support of π_0.[16] As we show below, π_0 affects the agent's limit posteriors only through P. To avoid technicalities, we assume from now on that the agent rules out values of ρ in a small neighbourhood of 1. That is, there exists $\overline{\rho} \in (\rho, 1)$ such that $\tilde{\rho} \leq \overline{\rho}$ for all $(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}) \in P$.

The likelihood function $L_t(\mathcal{H}_t|\tilde{p})$ associated to a parameter vector $\tilde{p}$ and history $\mathcal{H}_t = \{s_{t'}\}_{t'=1,\ldots,t}$ is the probability density of observing the signals conditional on $\tilde{p}$. From Bayes' law, this density is

$$L_t(\mathcal{H}_t|\tilde{p}) = L_t(s_1 \cdots s_t|\tilde{p}) = \prod_{t'=1}^{t} \ell_{t'}(s_{t'}|s_1 \cdots s_{t'-1}, \tilde{p}) = \prod_{t'=1}^{t} \ell_{t'}(s_{t'}|\mathcal{H}_{t'-1}, \tilde{p}),$$

where $\ell_t(s_t|\mathcal{H}_{t-1}, \tilde{p})$ denotes the density of s_t conditional on $\tilde{p}$ and $\mathcal{H}_{t-1}$. The latter density can be computed using the recursive-filtering formulas of Section 3.1. Indeed, Proposition 1 shows that conditional on $\tilde{p}$ and $\mathcal{H}_{t-1}$, x_{t-1} is normal. Since s_t is a linear function of x_{t-1}, it is also normal with a mean and variance that we denote by $\overline{s}_t(\tilde{p})$ and $\sigma_{s,t}^2(\tilde{p})$, respectively. Thus:

$$\ell_t(s_t|\mathcal{H}_{t-1}, \tilde{p}) = \frac{1}{\sqrt{2\pi\sigma_{s,t}^2(\tilde{p})}} \exp\left(-\frac{[s_t - \overline{s}_t(\tilde{p})]^2}{2\sigma_{s,t}^2(\tilde{p})}\right),$$

and

$$L_t(\mathcal{H}_t|\tilde{p}) = \frac{1}{\sqrt{(2\pi)^t \prod_{t'=1}^{t} \sigma_{s,t'}^2(\tilde{p})}} \exp\left(-\sum_{t'=1}^{t} \frac{[s_{t'} - \overline{s}_{t'}(\tilde{p})]^2}{2\sigma_{s,t'}^2(\tilde{p})}\right). \tag{15}$$

The agent's posterior beliefs over parameter vectors can be derived from his prior beliefs and the likelihood through Bayes' law. To determine posteriors in the limit when t goes to ∞, we need to determine the asymptotic behaviour of the likelihood function $L_t(\mathcal{H}_t|\tilde{p})$. Intuitively,

16. The closed support of π_0 is the intersection of all closed sets C such that $\pi_0(C) = 1$. Any neighbourhood B of an element of the closed support satisfies $\pi_0(B) > 0$ (Billingsley, 1986, 12.9, p. 181).

this behaviour depends on how well the agent can fit the data (i.e. the history of signals) using the model corresponding to $\tilde{p}$. To evaluate the fit of a model, we consider the true model according to which the data are generated. The true model is characterized by $\alpha = 0$ and the true parameters $p \equiv (\sigma_\eta^2, \rho, \sigma_\omega^2, \mu)$. We denote by $\overline{s}_t$ and $\sigma_{s,t}^2$, respectively, the true mean and variance of s_t conditional on $\mathcal{H}_{t-1}$, and by E the true expectation operator.

Theorem 1.

$$\lim_{t\to\infty} \frac{\log L_t(\mathcal{H}_t|\tilde{p})}{t} = -\frac{1}{2}\left(\log\left[2\pi\sigma_s^2(\tilde{p})\right] + \frac{\sigma_s^2 + e(\tilde{p})}{\sigma_s^2(\tilde{p})}\right) \equiv F(\tilde{p}) \tag{16}$$

almost surely, where

$$\sigma_s^2(\tilde{p}) \equiv \lim_{t\to\infty} \sigma_{s,t}^2(\tilde{p}),$$
$$\sigma_s^2 \equiv \lim_{t\to\infty} \sigma_{s,t}^2,$$
$$e(\tilde{p}) \equiv \lim_{t\to\infty} E\left[\overline{s}_t(\tilde{p}) - \overline{s}_t\right]^2.$$

Theorem 1 implies that the likelihood function is asymptotically equal to

$$L_t(\mathcal{H}_t|\tilde{p}) \sim \exp\left[tF(\tilde{p})\right],$$

thus growing exponentially at the rate $F(\tilde{p})$. Note that $F(\tilde{p})$ does not depend on the specific history $\mathcal{H}_t$ of signals, and is thus deterministic. That the likelihood function becomes deterministic for large t follows from the LLN, which is the main result that we need to prove the theorem. The appropriate large-numbers law in our setting is one applying to non-independent and non-identically distributed random variables. Non-independence is because the expected values $\overline{s}_t(\tilde{p})$ and $\overline{s}_t$ involve the entire history of past signals, and non-identical distributions are because the steady state is not reached within any finite time.

The growth rate $F(\tilde{p})$ can be interpreted as the fit of the model corresponding to $\tilde{p}$. Lemma 1 shows that when t goes to ∞, the agent gives weight only to values of $\tilde{p}$ that maximize $F(\tilde{p})$ over P.

Lemma 1. *The set $m(P) \equiv \text{argmax}_{\tilde{p}\in P} F(\tilde{p})$ is non-empty. When t goes to ∞, and for almost all histories, the posterior measure π_t converges weakly to a measure giving weight only to $m(P)$.*

Lemma 2 characterizes the solution to the fit-maximization problem under Assumption 1.

Assumption 1. *The set P satisfies the cone property*

$$(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}) \in P \Rightarrow (\lambda\tilde{\sigma}_\eta^2, \tilde{\rho}, \lambda\tilde{\sigma}_\omega^2, \tilde{\mu}) \in P, \quad \forall\lambda > 0.$$

Lemma 2. *Under Assumption 1, $\tilde{p} \in m(P)$ if and only if*

- $e(\tilde{p}) = \min_{\tilde{p}'\in P} e(\tilde{p}') \equiv e(P)$
- $\sigma_s^2(\tilde{p}) = \sigma_s^2 + e(\tilde{p})$.

The characterization of Lemma 2 is intuitive. The function $e(\tilde{p})$ is the expected squared difference between the conditional mean of s_t that the agent computes under $\tilde{p}$, and the true conditional mean. Thus, $e(\tilde{p})$ measures the error in the agent's predictions relative to the true model, and a model maximizing the fit must minimize this error.

A model maximizing the fit must also generate the right measure of uncertainty about the future signals. The agent's uncertainty under the model corresponding to $\tilde{p}$ is measured by $\sigma_s^2(\tilde{p})$, the conditional variance of s_t. This must equal to the true error in the agent's predictions, which is the sum of two orthogonal components: the error $e(\tilde{p})$ relative to the true model, and the error in the true model's predictions, i.e. the true conditional variance σ_s^2.

The cone property ensures that in maximizing the fit, there is no conflict between minimizing $e(\tilde{p})$ and setting $\sigma_s^2(\tilde{p}) = \sigma_s^2 + e(\tilde{p})$. Indeed, $e(\tilde{p})$ depends on $\tilde{\sigma}_\eta^2$ and $\tilde{\sigma}_\omega^2$ only through their ratio because only the ratio affects the regression coefficient $\tilde{G}$. The cone property ensures that given any feasible ratio, we can scale $\tilde{\sigma}_\eta^2$ and $\tilde{\sigma}_\omega^2$ to make $\sigma_s^2(\tilde{p})$ equal to $\sigma_s^2 + e(\tilde{p})$. The cone property is satisfied, in particular, when the set P includes all parameter values:

$$P = P_0 \equiv \left\{(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}) : \tilde{\sigma}_\eta^2 \in \mathbb{R}^+, \tilde{\rho} \in [0, \overline{\rho}], \tilde{\sigma}_\omega^2 \in \mathbb{R}^+, \tilde{\mu} \in \mathbb{R}\right\}.$$

Lemma 3 computes the error $e(\tilde{p})$. In both the lemma and subsequent analysis, we denote matrices corresponding to the true model by omitting the tilde. For example, the true-model counterpart of $\tilde{C} \equiv [\tilde{\rho}, -\alpha_{\tilde{\rho}}]$ is $C \equiv [\rho, 0]$.

Lemma 3. *The error $e(\tilde{p})$ is given by*

$$e(\tilde{p}) = \sigma_s^2 \sum_{k=1}^{\infty} (\tilde{N}_k - N_k)^2 + (N^\mu)^2 (\tilde{\mu} - \mu)^2, \tag{17}$$

where

$$\tilde{N}_k \equiv \tilde{C}(\tilde{A} - \tilde{G}\tilde{C})^{k-1}\tilde{G} + \sum_{k'=1}^{k-1} \tilde{C}(\tilde{A} - \tilde{G}\tilde{C})^{k-1-k'}\tilde{G}CA^{k'-1}G, \tag{18}$$

$$N_k \equiv CA^{k-1}G, \tag{19}$$

$$N^\mu \equiv 1 - \tilde{C}\sum_{k=1}^{\infty}(\tilde{A} - \tilde{G}\tilde{C})^{k-1}\tilde{G}. \tag{20}$$

The terms $\tilde{N}_k$ and N_k can be given an intuitive interpretation. Suppose that the steady state has been reached (i.e. a large number of periods have elapsed) and set $\zeta_t \equiv s_t - \overline{s}_t$. The shock ζ_t represents the "surprise" in period t, i.e. the difference between the signal s_t and its conditional mean $\overline{s}_t$ under the true model. The mean $\overline{s}_t$ is a linear function of the history $\{\zeta_{t'}\}_{t' \leq t-1}$ of past shocks, and N_k is the impact of ζ_{t-k}, i.e.

$$N_k = \frac{\partial \overline{s}_t}{\partial \zeta_{t-k}} = \frac{\partial E_{t-1}(s_t)}{\partial \zeta_{t-k}}. \tag{21}$$

The term $\tilde{N}_k$ is the counterpart of N_k under the agent's model, i.e.

$$\tilde{N}_k = \frac{\partial \overline{s}_t(\tilde{p})}{\partial \zeta_{t-k}} = \frac{\partial \tilde{E}_{t-1}(s_t)}{\partial \zeta_{t-k}}. \tag{22}$$

If $\tilde{N}_k \neq N_k$, then the shock ζ_{t-k} affects the agent's mean differently than the true mean. This translates into a contribution $(\tilde{N}_k - N_k)^2$ to the error $e(\tilde{p})$. Since the sequence $\{\zeta_t\}_{t\in\mathbb{Z}}$ is i.i.d., the contributions add up to the sum in equation (17).

The reason why equation (21) coincides with equation (19) is as follows. Because of linearity, the derivative in equation (21) can be computed by setting all shocks $\{\zeta_{t'}\}_{t'\leq t-1}$ to zero, except for $\zeta_{t-k} = 1$. The shock $\zeta_{t-k} = 1$ raises the mean of the state θ_{t-k} conditional on period $t-k$ by the regression coefficient G_1.[17] This effect decays over time according to the persistence parameter ρ because all subsequent shocks $\{\zeta_{t'}\}_{t'=t-k+1,\ldots,t-1}$ are zero, i.e. no surprises occur. Therefore, the mean of θ_{t-1} conditional on period $t-1$ is raised by $\rho^{k-1}G_1$, and the mean of s_t is raised by $\rho^k G_1 = CA^{k-1}G = N_k$.

The reason why equation (18) is more complicated than equation (19) is that after the shock $\zeta_{t-k} = 1$, the agent does not expect the shocks $\{\zeta_{t'}\}_{t'=t-k+1,\ldots,t-1}$ to be zero. This is both because the gambler's fallacy leads him to expect negative shocks, and because he can converge to $\tilde{G}_1 \neq G_1$, thus estimating incorrectly the increase in the state. Because, however, he observes the shocks $\{\zeta_{t'}\}_{t'=t-k+1,\ldots,t-1}$ to be zero, he treats them as surprises and updates accordingly. This generates the extra terms in equation (18). When α is small, i.e. the agent is close to rational, the updating generated by $\{\zeta_{t'}\}_{t'=t-k+1,\ldots,t-1}$ is of second order relative to that generated by ζ_{t-k}. The term $\tilde{N}_k$ then takes a form analogous to N_k:

$$\tilde{N}_k \approx \tilde{C}\tilde{A}^k\tilde{G} = \tilde{\rho}^k\tilde{G}_1 - \alpha_{\tilde{\rho}}(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^{k-1}\tilde{G}_2. \tag{23}$$

Intuitively, the shock $\zeta_{t-k} = 1$ raises the agent's mean of the state θ_{t-k} conditional on period $t-k$ by $\tilde{G}_1$, and the effect decays over time at the rate $\tilde{\rho}^k$. The agent also attributes the shock $\zeta_{t-k} = 1$ partly to luck through $\tilde{G}_2$. He then expects future signals to be lower because of the gambler's fallacy, and the effect decays at the rate $(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^k \approx \delta_{\tilde{\rho}}^k$.

We close this section by determining limit posteriors under rational updating ($\alpha = 0$). We examine, in particular, whether limit posteriors coincide with the true parameter values when prior beliefs give positive probability to all values, i.e. $P = P_0$.

Proposition 3. *Suppose that* $\alpha = 0$.

- *If* $\sigma_\eta^2 > 0$ *and* $\rho > 0$, *then* $m(P_0) = \{(\sigma_\eta^2, \rho, \sigma_\omega^2, \mu)\}$.
- *If* $\sigma_\eta^2 = 0$ *or* $\rho = 0$, *then*

$$m(P_0) = \left\{\begin{array}{r}(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}) : [\tilde{\sigma}_\eta^2 = 0, \tilde{\rho} \in [0, \overline{\rho}], \tilde{\sigma}_\omega^2 = \sigma_\eta^2 + \sigma_\omega^2, \tilde{\mu} = \mu] \\ \text{or } [\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \sigma_\eta^2 + \sigma_\omega^2, \tilde{\rho} = 0, \tilde{\mu} = \mu]\end{array}\right\}.$$

Proposition 3 shows that limit posteriors under rational updating coincide with the true parameter values if $\sigma_\eta^2 > 0$ and $\rho > 0$. If $\sigma_\eta^2 = 0$ or $\rho = 0$, however, then limit posteriors include both the true model and a set of other models. The intuition is that in both cases signals are i.i.d. in steady state: either because the state is constant ($\sigma_\eta^2 = 0$) or because it is not persistent ($\rho = 0$). Therefore, it is not possible to identify which of σ_η^2 or ρ is zero. Of course, the failure to converge to the true model is inconsequential because all models in the limit set predict correctly that signals are i.i.d.

17. In a steady-state context, the coefficients (G_1, G_2) and $(\tilde{G}_1, \tilde{G}_2)$ denote the elements of the vectors G and $\tilde{G}$, respectively.

4. INDEPENDENT SIGNALS

In this section we consider the agent's inference problem when signals are i.i.d. because $\sigma_\eta^2 = 0$.[18] Proposition 4 characterizes the beliefs that the agent converges to when he initially entertains all parameter values ($P = P_0$).

Proposition 4. *Suppose that* $\alpha > 0$ *and* $\sigma_\eta^2 = 0$. *Then* $e(P_0) = 0$ *and*

$$m(P_0) = \{(\tilde{\sigma}_\eta^2, 0, \tilde{\sigma}_\omega^2, \mu) : \tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \sigma_\omega^2\}.$$

Since $e(P_0) = 0$, the agent ends up predicting the signals correctly as i.i.d., despite the gambler's fallacy. The intuition derives from our assumption that the agent expects reversals only when all outcomes in a random sequence are drawn from the same distribution. Since the agent converges to the belief that $\tilde{\rho} = 0$, i.e. the state in one period has no relation to the state in the next, he assumes that a different distribution generates the signal in each period. In the mutual fund context, the agent assumes that managers stay in the fund for only one period. Therefore, his fallacious belief that a given manager's performance should average out over multiple periods has no effect. Formally, when $\tilde{\rho} = 0$, the strength $\alpha_{\tilde{\rho}}$ of the gambler's fallacy in equation (3) is $\alpha_{\tilde{\rho}} = \alpha\tilde{\rho} = 0$.[19]

We next allow the agent to rule out some parameter values based on prior knowledge. Ironically, prior knowledge can hurt the agent. Indeed, suppose that he knows with confidence that ρ is bounded away from zero. Then, he cannot converge to the belief $\tilde{\rho} = 0$, and consequently cannot predict the signals correctly. Thus, prior knowledge can be harmful because it reduces the agent's flexibility to come up with the incorrect model that eliminates the gambler's fallacy.

A straightforward example of prior knowledge is when the agent knows with confidence that the state is constant over time: this translates to the dogmatic belief that $\rho = 1$. A prototypical occurrence of such a belief is when people observe the flips of a coin they know is fair. The state can be defined as the probability distribution of heads and tails, and is known and constant.

If in our model the agent has a dogmatic belief that $\rho = 1$, then he predicts reversals according to the gambler's fallacy. This is consistent with the experimental evidence presented in Section 2. Of course, our model matches the evidence by construction, but we believe that this is a strength of our approach (in taking the gambler's fallacy as a primitive bias and examining whether the hot-hand fallacy can follow as an implication). Indeed, one could argue that the hot-hand fallacy is a primitive bias, either unconnected to the gambler's fallacy or perhaps even generating it. But then one would have to explain why such a primitive bias does not arise in experiments involving fair coins.

The hot-hand fallacy tends to arise in settings where people are uncertain about the mechanism generating the data, and where a belief that an underlying state varies over time is plausible *a priori*. Such settings are common when human skill is involved. For example, it is plausible—and often true—that the performance of a basketball player can fluctuate systematically over time because of mood, well-being, etc. Consistent with the evidence, we

18. As pointed out in the previous section, signals can be i.i.d. because $\sigma_\eta^2 = 0$ or $\rho = 0$. The source of i.i.d.-ness matters when the agent has prior knowledge on the true values of the parameters. We focus on the case $\sigma_\eta^2 = 0$ to study the effects of prior knowledge that ρ is bounded away from zero.

19. The result that the agent predicts the signals correctly as i.i.d. extends beyond the linear specification, but with a different intuition. See Appendix A.

show below that our approach can generate a hot-hand fallacy in such settings, provided that people are also confident that the state exhibits some persistence.[20]

More specifically, we assume that the agent allows for the possibility that the state varies over time, but is confident that ρ is bounded away from zero. For example, he can be uncertain as to whether fund managers differ in ability ($\sigma_\eta^2 > 0$), but know with certainty that they stay in a fund for more than one period. We take the closed support of the agent's priors to be

$$P = P_{\underline{\rho}} \equiv \left\{(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}) : \tilde{\sigma}_\eta^2 \in \mathbb{R}^+, \tilde{\rho} \in [\underline{\rho}, \overline{\rho}], \tilde{\sigma}_\omega^2 \in \mathbb{R}^+, \tilde{\mu} \in \mathbb{R}\right\},$$

where $\underline{\rho}$ is a lower bound strictly larger than zero and smaller than the true value ρ.

To determine the agent's convergent beliefs, we must minimize the error $e(\tilde{p})$ over the set $P_{\underline{\rho}}$. The problem is more complicated than in Propositions 3 and 4: it cannot be solved by finding parameter vectors $\tilde{p}$ such that $e(\tilde{p}) = 0$ because no such vectors exist in $P_{\underline{\rho}}$. Instead, we need to evaluate $e(\tilde{p})$ for all $\tilde{p}$ and minimize over $P_{\underline{\rho}}$. Equation (17) shows that $e(\tilde{p})$ depends on the regression coefficient $\tilde{G}$, which in turn depends on $\tilde{p}$ in a complicated fashion through the recursive-filtering formulas of Section 3.1. This makes it difficult to solve the problem in closed form. But a closed-form solution can be derived for small α, i.e. the agent close to rational. We next present this solution because it provides useful intuition and has similar properties to the numerical solution for general α.

Proposition 5. *Suppose that $\sigma_\eta^2 = 0$ and $\rho \geq \underline{\rho} > 0$. When α converges to zero, the set*

$$\left\{\left(\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) : \left(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) \in m(P_{\underline{\rho}})\right\}$$

converges (in the set topology) to the point $(z, \underline{\rho}, \sigma_\omega^2, \mu)$, where

$$z \equiv \frac{(1+\underline{\rho})^2}{1-\underline{\rho}^2\delta}. \tag{24}$$

Proposition 5 implies that the agent's convergent beliefs for small α are $\tilde{p} \approx (\alpha z\sigma_\omega^2, \underline{\rho}, \sigma_\omega^2, \mu)$. Convergence to $\tilde{\rho} = \underline{\rho}$ is intuitive. Indeed, Proposition 4 shows that the agent attempts to explain the absence of systematic reversals by underestimating the state's persistence $\tilde{\rho}$. The smallest value of $\tilde{\rho}$ consistent with the prior knowledge that $\rho \in [\underline{\rho}, \overline{\rho}]$ is $\underline{\rho}$.

The agent's belief that $\tilde{\rho} = \underline{\rho}$ leaves him unable to explain fully the absence of reversals. To generate a fuller explanation, he develops the additional fallacious belief that $\tilde{\sigma}_\eta^2 \approx \alpha z\sigma_\omega^2 > 0$, i.e. the state varies over time. Thus, in a mutual fund context, he overestimates both the extent of managerial turnover and the differences in ability. Overestimating turnover helps him to explain the absence of reversals in fund returns because he believes that reversals concern only the performance of individual managers. Overestimating differences in ability helps him further because he can attribute streaks of high or low fund returns to individual managers being above or below average. We show below that this belief in the changing state can generate a hot-hand fallacy.

20. Evidence linking the hot-hand fallacy to a belief in time-varying human skill comes from the casino-betting study of Croson and Sundali (2005). They show that consistent with the gambler's fallacy, individuals avoid betting on a colour with many recent occurrences. Consistent with the hot-hand fallacy, however, individuals raise their bets after successful prior bets.

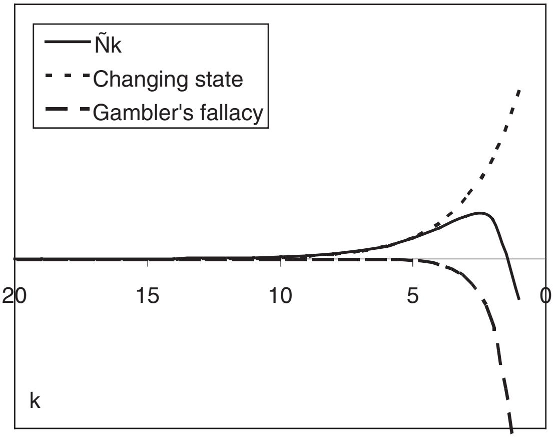

FIGURE 1
Effect of a signal in period $t-k$ on the agent's expectation $\tilde{E}_{t-1}(s_t)$, as a function of k. The dotted line represents the belief that the state has changed, the dashed line represents the effect of the gambler's fallacy, and the solid line is $\tilde{N}_k$, the sum of the two effects. The figure is drawn for $(\sigma_\eta^2/\sigma_\omega^2, \underline{\rho}, \alpha, \delta) = (0, 0.6, 0.2, 0.7)$

The error-minimization problem has a useful graphical representation. Consider the agent's expectation of s_t conditional on period $t-1$ as a function of the past signals. Equation (22) shows that the effect of the signal in period $t-k$, holding other signals to their mean, is $\tilde{N}_k$. Equation (23) expresses $\tilde{N}_k$ as the sum of two terms. Subsequent to a high s_{t-k}, the agent believes that the state has increased, which raises his expectation of s_t (term $\tilde{\rho}^k \tilde{G}_1$). But he also believes that luck should reverse, which lowers his expectation (term $-\alpha_{\tilde{\rho}}(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^{k-1}\tilde{G}_2$). Figure 1 plots these terms (dotted and dashed line, respectively) and their sum $\tilde{N}_k$ (solid line), as a function of the lag k.[21] Since signals are i.i.d. under the true model, $N_k = 0$. Therefore, minimizing the infinite sum in equation (17) amounts to finding $\left(\frac{\tilde{\sigma}_\eta^2}{\tilde{\sigma}_\omega^2}, \tilde{\rho}\right)$ that minimize the average squared distance between the solid line and the x-axis.

Figure 1 shows that $\tilde{N}_k$ is not always of the same sign. Thus, a high past signal does not lead the agent to predict always a high or always a low signal. Suppose instead that he always predicts a low signal because the gambler's fallacy dominates the belief that the state has increased (i.e. the dotted line is uniformly closer to the x-axis than the dashed line). This means that he converges to a small value of $\tilde{\sigma}_\eta^2$, believing that the state's variation is small, and treating signals as not very informative. But then, a larger value of $\tilde{\sigma}_\eta^2$ would shift the dotted line up, reducing the average distance between the solid line and the x-axis.

The change in $\tilde{N}_k$'s sign is from negative to positive. Thus, a high signal in the recent past leads the agent to predict a low signal, while a high signal in the more distant past leads him to predict a high signal. This is because the belief that the state has increased decays at the rate $\tilde{\rho}^k$, while the effect of the gambler's fallacy decays at the faster rate $(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^k = \tilde{\rho}^k(\delta - \alpha)^k$. In other words, after a high signal the agent expects luck to reverse quickly but views the increase in the state as more long-lived. The reason why he expects luck to reverse quickly

21. Since equation (23) holds approximately, up to second-order terms in α, the solid line is only an approximate sum of the dotted and dashed lines. The approximation error is, however, negligible for the parameter values for which Figures 1 and 2 are drawn.

relative to the state is that he views luck as specific to a given state (e.g. a given fund manager).

We next draw the implications of our results for the hot-hand fallacy. To define the hot-hand fallacy in our model, we consider a streak of identical signals between periods $t-k$ and $t-1$, and evaluate its impact on the agent's expectation of s_t. Viewing the expectation $\tilde{E}_{t-1}(s_t)$ as a function of the history of past signals, the impact is

$$\tilde{\Delta}_k \equiv \sum_{k'=1}^{k} \frac{\partial \tilde{E}_{t-1}(s_t)}{\partial s_{t-k'}}.$$

If $\tilde{\Delta}_k > 0$, then the agent expects a streak of k high signals to be followed by a high signal, and vice versa, for a streak of k low signals. This means that the agent expects streaks to continue and conforms to the hot-hand fallacy.[22]

Proposition 6. *Suppose that α is small, $\sigma_\eta^2 = 0$, $\rho \geq \underline{\rho} > 0$, and the agent considers parameter values in the set $P_{\underline{\rho}}$. Then, in steady state, $\tilde{\Delta}_k$ is negative for $k = 1$ and becomes positive as k increases.*

Proposition 6 shows that the hot-hand fallacy arises after long streaks while the gambler's fallacy arises after short streaks. This is consistent with Figure 1 because the effect of a streak is the sum of the effects $\tilde{N}_k$ of each signal in the streak. Since $\tilde{N}_k$ is negative for small k, the agent predicts a low signal following a short streak. But as streak length increases, the positive values of $\tilde{N}_k$ overtake the negative values, generating a positive cumulative effect.

Propositions 5 and 6 make use of the closed-form solutions derived for small α. For general α, the fit-maximization problem can be solved through a simple numerical algorithm and the results confirm the closed-form solutions: the agent converges to $\tilde{\sigma}_\eta^2 > 0$, $\tilde{\rho} = \underline{\rho}$, and $\tilde{\mu} = \mu$, and his predictions after streaks are as in Proposition 6.[23]

5. SERIALLY CORRELATED SIGNALS

In this section we relax the assumption that $\sigma_\eta^2 = 0$, to consider the case where signals are serially correlated. Serial correlation requires that the state varies over time ($\sigma_\eta^2 > 0$) and is persistent ($\rho > 0$). To highlight the new effects relative to the i.i.d. case, we assume that the agent has no prior knowledge on parameter values.

Recall that with i.i.d. signals and no prior knowledge, the agent predicts correctly because he converges to the belief that $\tilde{\rho} = 0$, i.e. the state in one period has no relation to the state in the next. When signals are serially correlated, the belief $\tilde{\rho} = 0$ obviously generates incorrect predictions. But predictions are also incorrect under a belief $\tilde{\rho} > 0$ because the gambler's

22. Our definition of the hot-hand fallacy is specific to streak length, i.e. the agent might conform to the fallacy for streaks of length k but not $k' \neq k$. An alternative and closely related definition can be given in terms of the agent's assessed auto-correlation of the signals. Denote by $\tilde{\Gamma}_k$ the correlation that the agent assesses between signals k periods apart. This correlation is closely related to the effect of the signal s_{t-k} on the agent's expectation of s_t, with the two being identical up to second-order terms when α is small. (The proof is available upon request.) Therefore, under the conditions of Proposition 6, the cumulative auto-correlation $\sum_{k'=1}^{k} \tilde{\Gamma}_k$ is negative for $k = 1$ and becomes positive as k increases. The hot-hand fallacy for lag k can be defined as $\sum_{k'=1}^{k} \tilde{\Gamma}_k > 0$.

23. The result that $\tilde{\sigma}_\eta^2 > 0$ can be shown analytically. The proof is available upon request.

fallacy then takes effect. Therefore, there is no parameter vector $\tilde{p} \in P_0$ achieving zero error $e(\tilde{p})$.[24]

We solve the error-minimization problem in closed form for small α and compare with the numerical solution for general α. In addition to α, we take σ_η^2 to be small, meaning that signals are close to i.i.d. We set $\nu \equiv \sigma_\eta^2/(\alpha\sigma_\omega^2)$ and assume that α and σ_η^2 converge to zero holding ν constant. The case where σ_η^2 remains constant while α converges to zero can be derived as a limit for $\nu = \infty$.

Proposition 7. *Suppose that $\rho > 0$. When α and σ_η^2 converge to zero, holding ν constant, the set*

$$\left\{\left(\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) : \left(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) \in m(P_0)\right\}$$

converges (in the set topology) to the point $(z, r, \sigma_\omega^2, \mu)$, where

$$z \equiv \frac{\nu\rho(1-\rho)(1+r)^2}{r(1+\rho)(1-\rho r)} + \frac{(1+r)^2}{1-r^2\delta} \tag{25}$$

and r solves

$$\frac{\nu\rho(1-\rho)(\rho-r)}{(1+\rho)(1-\rho r)^2} H_1(r) = \frac{r^2(1-\delta)}{(1-r^2\delta)^2} H_2(r), \tag{26}$$

for

$$H_1(r) \equiv \frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho r)} + \frac{r(1-\delta)\left[2-\rho r(1+\delta) - r^2\delta + \rho^2 r^4\delta^2\right]}{(1-r^2\delta)^2(1-\rho r\delta)^2},$$

$$H_2(r) \equiv \frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho r)} + \frac{r(1-\delta)\left(2-r^2\delta^2 - r^4\delta^3\right)}{(1-r^2\delta)(1-r^2\delta^2)^2}.$$

Because $H_1(r)$ and $H_2(r)$ are positive, equation (26) implies that $r \in (0, \rho)$. Thus, the agent converges to a persistence parameter $\tilde{\rho} = r$ that is between zero and the true value ρ. As in Section 4, the agent underestimates $\tilde{\rho}$ in his attempt to explain the absence of systematic reversals. But he does not converge all the way to $\tilde{\rho} = 0$ because he must explain the signals' serial correlation. Consistent with intuition, $\tilde{\rho}$ is close to zero when the gambler's fallacy is strong relative to the serial correlation (ν small), and is close to ρ in the opposite case.

Consider next the agent's estimate $(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2$ of the variance of the shocks to the state. Section 4 shows that when $\sigma_\eta^2 = 0$, the agent can develop the fallacious belief that $\tilde{\sigma}_\eta^2 > 0$

24. The proof of this result is available upon request. While models satisfying equations (1) and (2) cannot predict the signals correctly, models outside that set could. For example, when $(\alpha_\rho, \delta_\rho)$ are independent of ρ, correct predictions are possible under a model where the state is the sum of two auto-regressive processes: one with persistence $\tilde{\rho}_1 = \rho$, matching the true persistence, and one with persistence $\tilde{\rho}_2 = \delta - \alpha \equiv \delta_1 - \alpha_1$, offsetting the gambler's fallacy effect. Such a model would not generate correct predictions when $(\alpha_\rho, \delta_\rho)$ are linear in ρ because gambler's fallacy effects would decay at the two distinct rates $\delta_{\tilde{\rho}_i} - \alpha_{\tilde{\rho}_i} = \tilde{\rho}_i(\delta - \alpha)$ for $i = 1, 2$. Predictions might be correct, however, under more complicated models. Our focus is not as much to derive these models, but to characterize the agent's error patterns when inference is limited to a simple class of models that includes the true model.

as a way to counteract the effect of the gambler's fallacy. When σ_η^2 is positive, we find the analogous result that the agent overestimates $(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2$. Indeed, he converges to

$$(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2 \approx (1-r)^2\alpha z\sigma_\omega^2 = \frac{(1-r)^2 z}{(1-\rho)^2 \nu}(1-\rho)^2\sigma_\eta^2,$$

which is larger than $(1-\rho)^2\sigma_\eta^2$ because of equation (25) and $r < \rho$. Note that $(1-r)^2 z$ is decreasing in r. Thus, the agent overestimates the variance of the shocks to the state partly as a way to compensate for underestimating the state's persistence $\tilde{\rho}$.

The error-minimization problem can be represented graphically. Consider the agent's expectation of s_t conditional on period $t-1$ as a function of the past signals. The effect of the signal in period $t-k$, holding other signals to their mean, is $\tilde{N}_k$. Figure 2 plots $\tilde{N}_k$ (solid line) as a function of k. It also decomposes $\tilde{N}_k$ to the belief that the state has increased (dotted line) and the effect of the gambler's fallacy (dashed line). The new element relative to Figure 1 is that an increase in s_{t-k} also affects the expectation $E_{t-1}(s_t)$ under rational updating ($\alpha = 0$). This effect, N_k, is represented by the solid line with diamonds. Minimizing the infinite sum in equation (17) amounts to finding $(\frac{\tilde{\sigma}_\eta^2}{\tilde{\sigma}_\omega^2}, \tilde{\rho})$ that minimize the average squared distance between the solid line and the solid line with diamonds.

For large k, $\tilde{N}_k$ is below N_k, meaning that the agent under-reacts to signals in the distant past. This is because he underestimates the state's persistence parameter $\tilde{\rho}$, thus believing that the information learned from signals about the state becomes obsolete overly fast. Note that under-reaction to distant signals is a novel feature of the serial-correlation case. Indeed, with i.i.d. signals, the agent's underestimation of $\tilde{\rho}$ does not lead to under-reaction because there is no reaction under rational updating.

The agent's reaction to signals in the more recent past is in line with the i.i.d. case. Since $\tilde{N}_k$ cannot be below N_k uniformly (otherwise $e(\tilde{p})$ could be made smaller for a larger value of $\tilde{\sigma}_\eta^2$), it has to exceed N_k for smaller values of k. Thus, the agent over-reacts to signals in the more recent past. The intuition is as in Section 4: in overestimating $(1-\rho)^2\sigma_\eta^2$, the

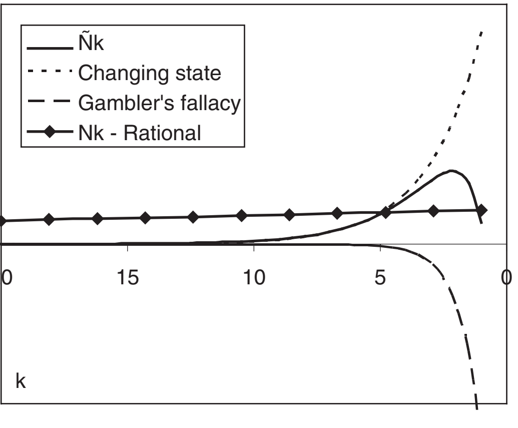

FIGURE 2

Effect of a signal in period $t-k$ on the agent's expectation $\tilde{E}_{t-1}(s_t)$, as a function of k. The dotted line represents the belief that the state has changed, the dashed line represents the effect of the gambler's fallacy, and the solid line is $\tilde{N}_k$, the sum of the two effects. The solid line with diamonds is N_k, the effect on the expectation $E_{t-1}(s_t)$ under rational updating ($\alpha = 0$). The figure is drawn for $(\sigma_\eta^2/\sigma_\omega^2, \rho, \alpha, \delta) = (0.001, 0.98, 0.2, 0.7)$

agent exaggerates the signals' informativeness about the state. Finally, the agent under-reacts to signals in the very recent past because of the gambler's fallacy.

We next draw the implications of our results for predictions after streaks. We consider a streak of identical signals between periods $t-k$ and $t-1$, and evaluate its impact on the agent's expectation of s_t, and on the expectation under rational updating. The impact for the agent is $\tilde{\Delta}_k$, and that under rational updating is

$$\Delta_k \equiv \sum_{k'=1}^{k} \frac{\partial E_{t-1}(s_t)}{\partial s_{t-k'}}.$$

If $\tilde{\Delta}_k > \Delta_k$, then the agent expects a streak of k high signals to be followed by a higher signal than under rational updating, and vice versa for a streak of k low signals. This means that the agent over-reacts to streaks.

Proposition 8. *Suppose that α and σ_η^2 are small, $\rho > 0$, and the agent has no prior knowledge ($P = P_0$). Then, in steady state $\tilde{\Delta}_k - \Delta_k$ is negative for $k = 1$, becomes positive as k increases, and then becomes negative again.*

Proposition 8 shows that the agent under-reacts to short streaks, over-reacts to longer streaks, and under-reacts to very long streaks. The under-reaction to short streaks is because of the gambler's fallacy. Longer streaks generate over-reaction because the agent overestimates the signals' informativeness about the state. But he also underestimates the state's persistence, thus under-reacting to very long streaks.

The numerical results for general α confirm most of the closed-form results. The only exception is that $\tilde{N}_k - N_k$ can change sign only once, from positive to negative. Under-reaction then occurs only to very long streaks. This tends to happen when the agent underestimates the state's persistence significantly (because α is large relative to $\tilde{\sigma}_\eta^2$). As a way to compensate for his error, he overestimates $(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2$ significantly, viewing signals as very informative about the state. Even very short streaks can then lead him to believe that the change in the state is large and dominates the effect of the gambler's fallacy.

We conclude this section by summarizing the main predictions of the model. These predictions could be tested in controlled experimental settings or in field settings. Prediction 1 follows from our specification of the gambler's fallacy.

Prediction 1 *When individuals observe i.i.d. signals and are told this information, they expect reversals after streaks of any length. The effect is stronger for longer streaks.*

Predictions 2 and 3 follow from the results of Sections 4 and 5. Both predictions require individuals to observe long sequences of signals so that they can learn sufficiently about the signal-generating mechanism.

Prediction 2 *Suppose that individuals observe i.i.d. signals, but are not told this information, and do not exclude on prior grounds the possibility that the underlying distribution might be changing over time. Then, a belief in continuation of streaks can arise. Such a belief should be observed following long streaks, while belief in reversals should be observed following short streaks. Both beliefs should be weaker if individuals believe on prior grounds that the underlying distribution might be changing frequently.*

Prediction 3 *Suppose that individuals observe serially correlated signals. Then, relative to the rational benchmark, they over-react to long streaks, but under-react to very long streaks and possibly to short ones as well.*

6. FINANCE APPLICATIONS

In this section we explore the implications of our model for financial decisions. Our goal is to show that the gambler's fallacy can have a wide range of implications, and that our normal-linear model is a useful tool for pursuing them.

6.1. *Active investing*

A prominent puzzle in finance is why people invest in actively managed funds in spite of the evidence that these funds underperform their passively managed counterparts. This puzzle has been documented by academics and practitioners, and has been the subject of two presidential addresses to the American Finance Association (Gruber, 1996; French, 2008). Gruber (1996) finds that the average active fund underperforms its passive counterpart by 35–164 basis points (bps, hundredths of a percent) per year. French (2008) estimates that the average investor would save 67 bps per year by switching to a passive fund. Yet, despite this evidence, passive funds represent only a small minority of mutual fund assets.[25]

Our model can help explain the active-fund puzzle. We interpret the signal as the return on a traded asset (e.g. a stock), and the state as the expected return. Suppose that expected returns are constant because $\sigma_\eta^2 = 0$, and so returns are i.i.d. Suppose also that an investor prone to the gambler's fallacy is uncertain about whether expected returns are constant, but is confident that if expected returns do vary, they are serially correlated ($\rho > 0$). Section 4 then implies that the investor ends up believing in return predictability. The investor would therefore be willing to pay for information on past returns, while such information has no value under rational updating.

Turning to the active-fund puzzle, suppose that the investor is unwilling to continuously monitor asset returns, but believes that market experts observe this information. Then, he would be willing to pay for experts' opinions or for mutual funds operated by the experts. The investor would thus be paying for active funds in a world where returns are i.i.d. and active funds have no advantage over passive funds. In summary, the gambler's fallacy can help explain the active-fund puzzle because it can generate an incorrect and confident belief that active managers add value.[26,27]

We illustrate our explanation with a simple asset-market model, to which we return in Section 6.2. Suppose that there are N stocks, and the return on stock $n = 1, \ldots, N$ in period t is s_{nt}. The return on each stock is generated by equations (1) and (2). The parameter σ_η^2 is equal to zero for all stocks, meaning that returns are i.i.d. over time. All stocks have the

25. Gruber (1996) finds that the average active fund investing in stocks underperforms the market by 65–194 bps per year, where the variation in estimates is due to different methods of risk adjustment. He compares these estimates to an expense ratio of 30 bps for passive funds. French (2008) aggregates the costs of active strategies over all stock market participants (mutual funds, hedge funds, institutional asset management, etc.) and finds that they are 79 bps, while the costs of passive strategies are 12 bps. Bogle (2005) reports that passive funds represent one-seventh of mutual fund equity assets as of 2004.

26. Our explanation of the active-fund puzzle relies not on the gambler's fallacy *per se*, but on the more general notion that people recognize deterministic patterns in random data. We are not aware of models of false pattern recognition that pursue the themes in this section.

27. Identifying the value added by active managers with their ability to observe past returns can be criticized on the grounds that past returns can be observed at low cost. Costs can, however, be large when there is a large number of assets, as in the model that we consider next. Stepping slightly outside of the model, suppose that the investor is confident that random data exhibit deterministic patterns, but does not know what the exact patterns are. (In our model this would mean that the investor is confident that signals are predictable, but does not know the model parameters.) Then, the value added by active managers would derive not only from their ability to observe past returns, but also from their knowledge of the patterns.

same expected return μ and variance σ_ϵ^2, and are independent of each other. An active-fund manager chooses an all-stock portfolio in each period by maximizing expected return subject to a tracking-error constraint. This constraint requires that the standard deviation of the manager's return relative to the equally weighted portfolio of all stocks does not exceed a bound TE. Constraints of this form are common in asset management (e.g. Roll, 1992). We denote by $s_t \equiv \sum_{n=1}^N s_{nt}/N$ the return on the equally weighted portfolio of all stocks.

The investor starts with the prior knowledge that the persistence parameter associated to each stock exceeds a bound $\underline{\rho} > 0$. He observes a long history of stock returns that ends in the distant past. This leads him to the limit posteriors on model parameters derived in Section 4. The investor does not observe recent returns. He assumes, however, that the manager has observed the entire return history, has learned about model parameters in the same way as him, and interprets recent returns in light of these estimates. We denote by $\tilde{E}$ and $\tilde{Var}$, respectively, the manager's expectation and variance operators, as assessed by the investor. The variance $\tilde{Var}_{t-1}(s_{nt})$ is independent of t in steady state, and independent of n because stocks are symmetric. We denote it by $\tilde{Var}_1$.

The investor is aware of the manager's objective and constraints. He assumes that in period t the manager chooses portfolio weights $\{w_{nt}\}_{n=1,\ldots,N}$ to maximize the expected return

$$\tilde{E}_{t-1}\left(\sum_{n=1}^N w_{nt}s_{nt}\right) \tag{27}$$

subject to the tracking error constraint

$$\tilde{Var}_{t-1}\left(\sum_{n=1}^N w_{nt}s_{nt} - s_t\right) \leq TE^2, \tag{28}$$

and the constraint that weights sum to 1, $\sum_{n=1}^N w_{nt} = 1$.

Lemma 4. *The manager's maximum expected return, as assessed by the investor, is*

$$\tilde{E}_{t-1}(s_t) + \frac{TE}{\sqrt{\tilde{Var}_1}}\sqrt{\sum_{n=1}^N \left[\tilde{E}_{t-1}(s_{nt}) - \tilde{E}_{t-1}(s_t)\right]^2}. \tag{29}$$

According to the investor, the manager's expected return exceeds that of a passive fund holding the equally weighted portfolio of all stocks. The manager adds value because of her information, as reflected in the cross-sectional dispersion of return forecasts. When the dispersion is large, the manager can achieve high expected return because her portfolio differs significantly from equal weights. The investor believes that the manager's forecasts should exhibit dispersion because the return on each stock is predictable on the basis of the stock's past history.

6.2. *Fund flows*

Closely related to the active-fund puzzle is a puzzle concerning fund flows. Flows into mutual funds are strongly positively correlated with the funds' lagged returns (e.g. Chevalier and Ellison, 1997; Sirri and Tufano, 1998), and yet lagged returns do not appear to be strong predictors of future returns (e.g. Carhart, 1997).[28] Ideally, the fund-flow and active-fund puzzles should be addressed together within a unified setting: explaining flows into active funds raises the question why people invest in these funds in the first place.

To study fund flows, we extend the asset-market model of Section 6.1, allowing the investor to allocate wealth over an active and a passive fund. The passive fund holds the equally weighted portfolio of all stocks and generates return s_t. The active fund's return is $s_t + \gamma_t$. In a first step, we model the excess return γ_t in reduced form, without reference to the analysis of Section 6.1: we assume that γ_t is generated by equations (1) and (2) for values of $(\sigma_\eta^2, \sigma_\epsilon^2)$ denoted by $(\sigma_{\eta\gamma}^2, \sigma_{\epsilon\gamma}^2)$, γ_t is i.i.d. because $\sigma_{\eta\gamma}^2 = 0$, and γ_t is independent of s_t. Under these assumptions, the investor's beliefs about γ_t can be derived as in Section 4. In a second step, we use the analysis of Section 6.1 to justify some of the assumed properties of γ_t, and more importantly to derive additional properties.

We modify slightly the model of Section 6.1 by replacing the single infinitely lived investor with a sequence of generations. Each generation consists of one investor, who invests over one period and observes all past information. This simplifies the portfolio problem, rendering it myopic, while the inference problem remains as with a single investor.

Investors observe the history of returns of the active and the passive fund. Their portfolio allocation decision, derived below, depends on their beliefs about the excess return γ_t of the active fund. Investors believe that γ_t is generated by equations (1) and (2), is independent of s_t, and its persistence parameter exceeds a bound $\underline{\rho}_\gamma > 0$. After observing a long history of fund returns, investors converge to the limit posteriors on model parameters derived in Section 4.

The investor entering the market in period t starts with wealth $\tilde{W}$, and derives utility over final wealth $\tilde{W}_t$. To keep with the normal-linear structure, we take utility to be exponential with coefficient of absolute risk aversion a. The investor chooses portfolio weight $\tilde{w}_t$ in the active fund to maximize the expected utility

$$-\tilde{E}_{t-1} \exp(-a\tilde{W}_t)$$

subject to the budget constraint

$$\begin{aligned}\tilde{W}_t &= \tilde{W}\left[(1-\tilde{w}_t)(1+s_t) + \tilde{w}_t(1+s_t+\gamma_t)\right]\\ &= \tilde{W}\left[(1+s_t) + \tilde{w}_t\gamma_t\right].\end{aligned}$$

Because of normality and exponential utility, the problem is mean-variance. Independence between s_t and γ_t implies that the solution is

$$\tilde{w}_t = \frac{\tilde{E}_{t-1}(\gamma_t)}{a\tilde{W}\tilde{Var}_{1\gamma}},$$

where $\tilde{Var}_{1\gamma} \equiv \tilde{Var}_{t-1}(\gamma_t)$ and $\tilde{Var}$ denotes the variance assessed by the investor. (The variance $\tilde{Var}_{t-1}(\gamma_t)$ is independent of t in steady state.) The optimal investment in the active fund is a linear increasing function of the investor's expectation of the fund's excess return γ_t. The net

28. Similar findings exist for hedge funds. See, for example, Ding *et al.* (2008) and Fung *et al.* (2008). Berk and Green (2004) propose a rational explanation for the fund-flow puzzle. They assume that ability differs across managers and can be inferred from fund returns. Able managers perform well and receive inflows, but because of decreasing returns to managing a large fund, their performance drops to that of average managers. Their model does not explain the active-fund puzzle, nor does it fully address the origin of decreasing returns.

flow into the active fund in period $t+1$ is the change in investment between periods t and $t+1$:

$$\tilde{F}_{t+1} \equiv \tilde{W}\tilde{w}_{t+1} - \tilde{W}\tilde{w}_t = \frac{\tilde{E}_t(\gamma_{t+1}) - \tilde{E}_{t-1}(\gamma_t)}{a\tilde{V}ar_{1\gamma}}. \tag{30}$$

Lemma 5 determines the covariance between returns and subsequent flows.

Lemma 5. *The covariance between the excess return of the active fund in period t and the net flow into the fund during periods $t+1$ through $t+k$ for $k \geq 1$ is*

$$Cov\left(\gamma_t, \sum_{k'=1}^{k} \tilde{F}_{t+k'}\right) = \frac{\sigma^2_{\epsilon\gamma}\tilde{N}_k}{a\tilde{V}ar_{1\gamma}}. \tag{31}$$

where $\tilde{N}_k$ is defined by equation (22), with γ_t replacing s_t.

Recall from Section 4 (e.g. Figure 1) that $\tilde{N}_k$ is negative for small k and becomes positive as k increases. Hence, returns are negatively correlated with subsequent flows when flows are measured over a short horizon, and are positively correlated over a long horizon. The negative correlation is inconsistent with the evidence on the fund-flow puzzle. It arises because investors attribute high returns partly to luck, and expecting luck to reverse, they reduce their investment in the fund. Investors also develop a fallacious belief that high returns indicate high managerial ability, and this tends to generate positive performance–flow correlation, consistent with the evidence. But the correlation cannot be positive over all horizons because the belief in ability arises to offset but not to overtake the gambler's fallacy.

We next use the analysis of Section 6.1 to derive properties of the active fund's excess return γ_t. Consider first properties under the true model. Since the returns on all stocks are identical and constant over time, all active portfolios achieve zero excess return. The manager is thus indifferent over portfolios, and we assume that he selects one meeting the tracking-error constraint [equation (28)] with equality.[29] Therefore, conditional on past history, γ_t is normal with mean zero and variance $\sigma^2_{\epsilon\gamma} = TE^2$. Since the mean and variance do not depend on past history, γ_t is i.i.d. Moreover, γ_t is uncorrelated with s_t because

$$Cov_{t-1}(\gamma_t, s_t) = Cov_{t-1}\left(\sum_{n=1}^{N} w_{nt}s_{nt} - s_t, s_t\right) = \sigma^2_{\epsilon}\left(\frac{\sum_{n=1}^{N} w_{n,t}}{N} - \frac{1}{N}\right) = 0,$$

where the second step follows because stocks are symmetric and returns are uncorrelated.

Consider next properties of γ_t as assessed by the investor. Since the investor believes that the active manager adds value, γ_t has positive mean. Moreover, this mean varies over time and is serially correlated. Indeed, recall from Lemma 4 that the active fund's expected excess return, as assessed by the investor, is increasing in the cross-sectional dispersion of the manager's return forecasts. This cross-sectional dispersion varies over time and is serially correlated: it is small, for example, when stocks have similar return histories since these histories are used to generate the forecasts.

29. We assume that the manager is rational and operates under the true model. The manager has a strict preference for a portfolio meeting equation (28) with equality if returns are arbitrarily close to i.i.d.

Time variation in the mean of γ_t implies that the value added by the manager is time-varying, i.e. $\sigma^2_{\eta\gamma} > 0$. Section 4 derives belief in time-varying managerial ability as a posterior to which the investor converges after observing a long history of fund returns. This section shows instead that such a belief can arise without observation of fund returns, and as a consequence of the investor's belief in the power of active management. This belief can act as a prior when the investor observes fund returns, and can lead to inferences different than in Section 4. Indeed, suppose that the investor believes on prior grounds that $\sigma^2_{\eta\gamma}$ exceeds a bound $\underline{\sigma}^2_{\eta\gamma} > 0$. If the convergent posterior in the absence of this prior knowledge is smaller than $\underline{\sigma}^2_{\eta\gamma}$, then prior knowledge implies convergence to $\underline{\sigma}^2_{\eta\gamma}$ instead. As a consequence, the fallacious belief that high returns indicate high managerial ability becomes stronger and could generate positive performance–flow correlation over any horizon, consistent with the fund-flow puzzle.

A problem with our analysis is that when stock returns are normal (as in Section 6.1), fund returns, as assessed by the investor, are non-normal. This is because the mean of γ_t depends on the cross-sectional dispersion of stock return forecasts, which is always positive and thus non-normal. Despite this limitation, we believe that the main point is robust: a fallacious belief that active managers add value gives rise naturally to a belief that the added value varies over time, in a way that can help explain the fund-flow puzzle. Belief in time variation could arise through the cross-sectional dispersion of the manager's return forecasts. Alternatively, it could arise because of an assumed dispersion in ability across managers, e.g. only some managers collect information on past returns and make good forecasts according to the investor. Neither of the two mechanisms would generate time variation under the true model because past returns do not forecast future returns. Put differently, a rational investor who knows that active managers do not add value would also rule out the possibility that the added value varies over time.

We illustrate our analysis with a numerical example. We assume that the investor observes fund returns once a year. We set $\underline{\rho}_\gamma = 0.6$, i.e. at least 60% of managerial ability carries over to the next year. We set the tracking error $TE = 5\%$. We follow the calibration of Section 2 and set the parameters of the gambler's fallacy $\alpha = 0.2$ and $\delta = 0.7$. Section 4 then implies that in the absence of prior knowledge on $\sigma_{\eta\gamma}$, the investor converges to $\tilde{\sigma}_{\eta\gamma} = 3.66\%$. He thus believes that managers performing one standard deviation above average beat those performing one standard deviation below by 7.32% per year.

Can prior knowledge lead to the belief that $\sigma_{\eta\gamma}$ exceeds $\tilde{\sigma}_{\eta\gamma}$? Suppose, for example, that there are two types of managers: those collecting information on past returns and those who do not. A manager collecting no information achieves zero expected excess return. To derive the expected excess return, as assessed by the investor, of a manager collecting information, we make the following assumptions. There are $N = 50$ stocks, and the manager observes stock returns and makes trades once a month. We set $\underline{\rho} = 0.8$, i.e. at least 80% of a shock to a stock's expected return carries over to the next month. For large N, the expected excess return in Lemma 4 becomes approximately

$$\frac{TE}{\sqrt{\tilde{Var}_1}}\sqrt{N\tilde{Var}\left[\tilde{E}_{t-1}(s_{nt})\right]}. \tag{32}$$

The ratio $\sqrt{\tilde{Var}[\tilde{E}_{t-1}(s_{nt})]/\tilde{Var}_1}$ is 9.97%, meaning that the cross-sectional dispersion of the manager's return forecasts is 9.97% of her estimate of the standard deviation of each stock. Since we are evaluating returns at a monthly frequency, we set $TE = 5\%/\sqrt{12}$. Substituting into equation (32), we find that the manager's expected excess return is 0.77% per month, i.e. 9.23% per year. Assuming that managers of the two types are equally likely, this translates to a standard deviation in managerial ability of 4.62% per year. This is larger than $\tilde{\sigma}_{\eta\gamma} = 3.66\%$.

Adding the prior knowledge that $\sigma_{\eta\gamma} \geq 4.62\%$ to the analysis of Section 4 yields a positive performance–flow correlation over any horizon, consistent with the fund-flow puzzle.

While flows in the numerical example respond positively to returns over any horizon, the response is strongest over intermediate horizons. This reflects the pattern derived in Section 4 (e.g. Figure 1), with the difference that the $\tilde{N}_k$-curve is positive even for small k. Of course, the delayed reaction of flows to returns could arise for reasons other than the gambler's fallacy, e.g. costs of adjustment.

6.3. *Equilibrium*

Sections 6.1 and 6.2 take stock returns as given. A comprehensive analysis of the general equilibrium effects that would arise if investors prone to the gambler's fallacy constitute a large fraction of the market is beyond the scope of this paper. Nonetheless, some implications can be sketched.

Suppose that investors observe a stock's i.i.d. normal dividends. Suppose, as in Section 6.2, that investors form a sequence of generations, each investing over one period and maximizing exponential utility. Suppose finally that investors are uncertain about whether expected dividends are constant, but are confident that if expected dividends do vary, they are serially correlated ($\rho > 0$). If investors are rational, they would learn that dividends are i.i.d., and equilibrium returns would also be i.i.d.[30] If, instead, investors are prone to the gambler's fallacy, returns would exhibit short-run momentum and long-run reversal. Intuitively, since investors expect a short streak of high dividends to reverse, the stock price under-reacts to the streak. Therefore, a high return is, on average, followed by a high return, implying short-run momentum. Since, instead, investors expect a long streak of high dividends to continue, the stock price over-reacts to the streak. Therefore, a sequence of high returns is, on average, followed by a low return, implying long-run reversal and a value effect. These results are similar to Barberis, Shleifer and Vishny (1998), although the mechanism is different.

A second application concerns trading volume. Suppose that all investors are subject to the gambler's fallacy, but observe different subsets of the return history. Then, they would form different forecasts for future dividends. If, in addition, prices are not fully revealing (e.g. because of noise trading), then investors would trade because of the different forecasts. No such trading would occur under rational updating, and therefore the gambler's fallacy would generate excessive trading.

APPENDIX A. ALTERNATIVE SPECIFICATIONS FOR $(\alpha_\rho, \delta_\rho)$

Consider first the case where signals are i.i.d. because $\sigma_\eta^2 = 0$, and the agent initially entertains all parameter values. Under the constant specification, where $(\alpha_\rho, \delta_\rho)$ are independent of ρ and equal to $(\alpha, \delta) \equiv (\alpha_1, \delta_1)$, Proposition 4 is replaced by

Proposition A1. *Suppose that $\alpha > 0$ and $\sigma_\eta^2 = 0$. Then $e(P_0) = 0$, and $m(P_0)$ consists of the two elements*

$$\tilde{p}_1 \equiv \left(\frac{\alpha(1 - \delta^2 + \delta\alpha)}{\delta(1 - \delta + \alpha)^2} \sigma_\omega^2, \delta - \alpha, \frac{\delta - \alpha}{\delta} \sigma_\omega^2, \mu \right)$$

and

$$\tilde{p}_2 \equiv (\sigma_\omega^2, 0, 0, \mu).$$

30. Proofs of all results in this section are available upon request.

As in Proposition 4, the agent ends up predicting the signals correctly as i.i.d., despite the gambler's fallacy. Correct predictions can be made using either of two models. Under model $\tilde{p}_1$, the agent believes that the state θ_t varies over time and is persistent. The state's persistence parameter $\tilde{\rho}$ is equal to the decay rate $\delta - \alpha > 0$ of the gambler's fallacy effect. (Decay rates are derived in (23).) This ensures that the agent's updating on the state can exactly offset his belief that luck should reverse. Under model $\tilde{p}_2$, the agent attributes all variation in the signal to the state, and since the gambler's fallacy applies only to the sequence $\{\epsilon_t\}_{t\geq 1}$, it has no effect.

Proposition 4 can be extended to a broader class of specifications. Suppose that the function $\rho \in [0, 1] \rightarrow \delta_\rho - \alpha_\rho$ crosses the 45-degree line at a single point $\hat{\rho} \in [0, 1)$. Then, the agent ends up predicting the signals correctly as i.i.d., i.e. $e(P_0) = 0$, and $m(P_0)$ consists of the element $\tilde{p}_2$ of Proposition A1 and an element $\tilde{p}_1$ for which $\tilde{\rho} = \hat{\rho}$. Under the linear specification, $\hat{\rho} = 0$, and under the constant specification, $\hat{\rho} = \delta - \alpha$.

Consider next the case where the agent is confident that ρ is bounded away from zero, and the closed support of his priors is $P = P_{\underline{\rho}}$. Under the constant specification, Proposition 6 is replaced by:

Proposition A2. *Suppose that α is small, $\sigma_\eta^2 = 0$, $\rho \geq \underline{\rho} > \delta$, and the agent considers parameter values in the set $P_{\underline{\rho}}$. Then, in steady state $\tilde{\Delta}_k$ is negative for $k = 1$ and becomes positive as k increases.*

As in Proposition 6, the hot-hand fallacy arises after long streaks of signals while the gambler's fallacy arises after short streaks. Note that Proposition A2 requires the additional condition $\underline{\rho} > \delta$. If instead $\underline{\rho} < \delta$, then the agent predicts the signals correctly as i.i.d. Indeed, because $\underline{\rho} < \delta - \alpha$ for small α, model $\tilde{p}_1$ of Proposition A1 belongs to $P_{\underline{\rho}}$ (provided that $\overline{\rho}$ is close to one so that $\overline{\rho} > \delta$), and therefore, $e(P_{\underline{\rho}}) = 0$ and $m(P_{\underline{\rho}}) = \{\tilde{p}_1\}$. Proposition A2 can be extended to the broader class of specifications defined above, with the condition $\underline{\rho} > \delta$ replaced by $\underline{\rho} > \hat{\rho}$.

Consider finally the case where signals are serially correlated. Under the constant specification, Proposition 8 is replaced by:

Proposition A3. *Suppose that α and σ_η^2 are small, $\rho \notin \{0, \delta\}$, the agent has no prior knowledge ($P = P_0$), and one of the following is met*

$$\frac{2\rho}{1+\rho^2} > \delta, \tag{A1}$$

$$\frac{\nu^2\rho^4(1-\rho)(1-\delta^2)}{(1+\rho)^3} > 1. \tag{A2}$$

Then, in steady state $\tilde{\Delta}_k - \Delta_k$ is negative for $k = 1$, becomes positive as k increases, and then becomes negative again.

As in Proposition 8, the agent under-reacts to short streaks, over-reacts to longer streaks, and under-reacts to very long streaks. Proposition A3 is, in a sense, stronger than Proposition 8 because it derives the under/over/under-reaction pattern regardless of whether ρ is larger or smaller than δ. (The relevant comparison in Proposition 8 is between ρ and zero.) The agent converges to a value $\tilde{\rho}$ between ρ and δ. When $\rho > \delta$, the under-reaction to very long streaks is because the agent underestimates the state's persistence. When instead $\rho < \delta$, the agent overestimates the state's persistence, and the under-reaction to very long streaks is due to the large memory parameter δ of the gambler's fallacy. Conditions (A1) and (A2) ensure that when α and σ_η^2 are small, in which case the agent converges to a model predicting signals close to i.i.d., it is $\tilde{\sigma}_\eta^2$ rather than $\tilde{\rho}$ that is small. A small-$\tilde{\rho}$ model predicts signals close to i.i.d. for the same reason as model $\tilde{p}_2$ of Proposition A1: because variation in the signal is mostly attributed to the state. Appendix B shows that when the gambler's fallacy applies to both $\{\epsilon_t\}_{t\geq 1}$ and $\{\eta_t\}_{t\geq 1}$, a model similar to $\tilde{p}_1$ becomes the unique element of $m(P_0)$ in the setting of Proposition A1, and conditions (A1) and (A2) are not needed in the setting of Proposition A3.

APPENDIX B. GAMBLER'S FALLACY FOR THE STATE

Our model and solution method can be extended to the case where the gambler's fallacy applies to both the sequence $\{\epsilon_t\}_{t\geq 1}$ that generates the signals given the state, and the sequence $\{\eta_t\}_{t\geq 1}$ that generates the state. Suppose that according to the agent,

$$\eta_t = \zeta_t - \alpha \sum_{k=0}^{\infty} \delta^k \eta_{t-1-k}, \tag{B1}$$

where the sequence $\{\zeta_t\}_{t\geq 1}$ is i.i.d. normal with mean zero and variance σ_ζ^2. To formulate the recursive-filtering problem, we expand the state vector to

$$x_t \equiv \left[\theta_t - \tilde{\mu}, \epsilon_t^\delta, \eta_t^\delta\right]',$$

where

$$\eta_t^\delta \equiv \sum_{k=0}^{\infty} \delta^k \eta_{t-k}.$$

We also set

$$\tilde{A} \equiv \begin{bmatrix} \tilde{\rho} & 0 & -(1-\tilde{\rho})\alpha \\ 0 & \delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}} & 0 \\ 0 & 0 & \delta - \alpha \end{bmatrix},$$

$$w_t \equiv [(1-\tilde{\rho})\zeta_t, \omega_t, \zeta_t]',$$

$$\tilde{C} \equiv [\tilde{\rho}, -\alpha_{\tilde{\rho}}, -(1-\tilde{\rho})\alpha],$$

$v_t \equiv (1-\tilde{\rho})\zeta_t + \omega_t$, and $\tilde{p} \equiv (\tilde{\sigma}_\zeta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu})$. Under these definitions, the analysis of the agent's inference in Section 3 carries through identical.

When signals are i.i.d. because $\sigma_\eta^2 = 0$, and the agent initially entertains all parameter values, his limit posteriors are as follows:

Proposition B1. *Suppose that $\alpha > 0$ and $\sigma_\eta^2 = 0$. If $(\alpha_\rho, \delta_\rho) = (\alpha\rho, \delta\rho)$, then $e(P_0) = 0$ and*

$$m(P_0) = \{(0, 0, \sigma_\omega^2, \mu)\}.$$

If $(\alpha_\rho, \delta_\rho) = (\alpha, \delta)$, then $e(P_0) = 0$ and

$$m(P_0) = \left\{\left(\frac{\alpha(1-\delta^2+\delta\alpha)}{\delta(1-\delta)^2}\sigma_\omega^2, \delta, \frac{\delta-\alpha}{\delta}\sigma_\omega^2, \mu\right)\right\}.$$

Under both the linear and the constant specification, the agent ends up predicting the signals correctly as i.i.d. Therefore, the main result of Propositions 4 and A1 extends to the case where the gambler's fallacy applies to both $\{\epsilon_t\}_{t\geq 1}$ and $\{\eta_t\}_{t\geq 1}$. The set of models that yield correct predictions, however, becomes a singleton. The only model under the linear specification is one in which the state is constant, and the only model under the constant specification is one similar to model $\tilde{p}_1$ of Proposition A1.

When the agent is confident that ρ is bounded away from zero, he ends up believing that $\tilde{\sigma}_\eta^2 > 0$. When, in addition, α is small, so is $\tilde{\sigma}_\eta^2$. Therefore, the effects of the gambler's fallacy for $\{\eta_t\}_{t\geq 1}$ are small relative to $\{\epsilon_t\}_{t\geq 1}$, and the analysis is the same as when the gambler's fallacy applies only to $\{\epsilon_t\}_{t\geq 1}$. In particular, the agent's predictions after streaks are as in Propositions 6 and A2.

Consider finally the case where signals are serially correlated. When α and σ_η^2 are small, the agent converges to a model where $\tilde{\sigma}_\eta^2$ is small. (Unlike with Proposition A3, Conditions (A1) and (A2) are not needed to rule out the small-$\tilde{\rho}$ model.) Therefore, the effects of the gambler's fallacy for $\{\eta_t\}_{t\geq 1}$ are small relative to $\{\epsilon_t\}_{t\geq 1}$, and the analysis is the same as when the gambler's fallacy applies only to $\{\epsilon_t\}_{t\geq 1}$. In particular, the agent's predictions after streaks are as in Propositions 8 and A3.

APPENDIX C. PROOFS

Proof of Proposition 1. Our formulation of the recursive-filtering problem is as in standard textbooks. For example, equations (8) and (9) follow from (4.1.1) and (4.1.4) in Balakrishnan (1987) if x_{n+1} is replaced by x_t, x_n by x_{t-1}, A_n by $\tilde{A}$, U_n by 0, N_n^s by w_t, v_n by $s_t - \tilde{\mu}$, C_n by $\tilde{C}$, and N_n^0 by v_t. Equation (10) follows from (4.6.14), if the latter is written for $n+1$ instead of n, and $\overline{x}_{n+1}$ is replaced by $\overline{x}_t$, $\overline{x}_n$ by $\overline{x}_{t-1}$, and $AK_n + Q_n$ by $\tilde{G}_t$. That $\tilde{G}_t$ so defined is given by equation (12) follows from (4.1.29) and (4.6.12) if H_{n-1} is replaced by $\tilde{\Sigma}_{t-1}$, $G_n G_n'$ by $\tilde{V}$, and J_n by $\tilde{U}$. Equation (11) follows from (4.6.18) if the latter is written for $n+1$ instead of n, P_n is substituted from (4.1.30), and $F_n F_n'$ is replaced by $\tilde{W}$. ‖

Proof of Proposition 2. It suffices to show (Balakrishnan, pp. 182–184) that the eigenvalues of $\tilde{A} - \tilde{U}\tilde{V}^{-1}\tilde{C}$ have modulus smaller than 1. This matrix is

$$\begin{bmatrix} \frac{\tilde{\sigma}_\omega^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\tilde{\rho} & \frac{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\alpha_{\tilde{\rho}} \\ -\frac{\tilde{\sigma}_\omega^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\tilde{\rho} & \delta_{\tilde{\rho}} - \frac{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\alpha_{\tilde{\rho}} \end{bmatrix}.$$

The characteristic polynomial is

$$\begin{aligned} &\lambda^2 - \lambda\left[\frac{\tilde{\sigma}_\omega^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\tilde{\rho} + \delta_{\tilde{\rho}} - \frac{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\alpha_{\tilde{\rho}}\right] + \frac{\tilde{\sigma}_\omega^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\tilde{\rho}\delta_{\tilde{\rho}} \\ \equiv\ &\lambda^2 - \lambda b + c \end{aligned}$$

Suppose that the roots λ_1, λ_2 of this polynomial are real, in which case $\lambda_1 + \lambda_2 = b$ and $\lambda_1\lambda_2 = c$. Since $c > 0$, λ_1 and λ_2 have the same sign. If λ_1 and λ_2 are negative, they are both greater than -1, since $b > -1$ from $\alpha_{\tilde{\rho}} < 1$ and $\tilde{\rho}, \delta_{\tilde{\rho}} \geq 0$. If λ_1 and λ_2 are positive, then at least one is smaller than 1, since $b < 2$ from $\tilde{\rho}, \delta_{\tilde{\rho}} < 1$ and $\alpha_{\tilde{\rho}} \geq 0$. But since the characteristic polynomial for $\lambda = 1$ takes the value

$$(1-\delta_{\tilde{\rho}})\left[1 - \frac{\tilde{\sigma}_\omega^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\tilde{\rho}\right] + \frac{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2}{(1-\tilde{\rho})^2\tilde{\sigma}_\eta^2+\tilde{\sigma}_\omega^2}\alpha_{\tilde{\rho}} > 0,$$

both λ_1 and λ_2 are smaller than 1. Suppose instead that λ_1, λ_2 are complex. In that case, they are conjugates and the modulus of each is $\sqrt{c} < 1$. ‖

Lemma C1 determines $\overline{s}_t$, the true mean of s_t conditional on $\mathcal{H}_{t-1}$, and $\overline{s}_t(\tilde{p})$, the mean that the agent computes under the parameter vector $\tilde{p}$. To state the lemma, we set $\zeta_t \equiv s_t - \overline{s}_t$, $\tilde{D}_t \equiv \tilde{A} - \tilde{G}_t\tilde{C}$, $\tilde{D} \equiv \tilde{A} - \tilde{G}\tilde{C}$, and

$$\tilde{J}_{t,t'} \equiv \begin{cases} \prod_{k=t'}^{t} \tilde{D}_k & \text{for } t' = 1, \ldots, t, \\ I & \text{for } t' > t. \end{cases}$$

For simplicity, we set the initial condition $\overline{x}_0 = 0$.

Lemma C1. *The true mean $\overline{s}_t$ is given by*

$$\overline{s}_t = \mu + \sum_{t'=1}^{t-1} CA^{t-t'-1}G_{t'}\zeta_{t'} \tag{C1}$$

and the agent's mean $\overline{s}_t(\tilde{p})$ by

$$\overline{s}_t(\tilde{p}) = \tilde{\mu} + \sum_{t'=1}^{t-1} \tilde{C}\tilde{M}_{t,t'}\zeta_{t'} + \tilde{C}\tilde{M}_t^\mu(\mu - \tilde{\mu}), \tag{C2}$$

where

$$\begin{aligned} \tilde{M}_{t,t'} &\equiv \tilde{J}_{t-1,t'+1}\tilde{G}_{t'} + \sum_{k=t'+1}^{t-1} \tilde{J}_{t-1,k+1}\tilde{G}_k CA^{k-t'-1}G_{t'}, \\ \tilde{M}_t^\mu &\equiv \sum_{t'=1}^{t-1} \tilde{J}_{t-1,t'+1}\tilde{G}_{t'}. \end{aligned}$$

Proof. Consider the recursive-filtering problem under the true model, and denote by $\overline{x}_t$ the true mean of x_t. Equation (9) implies that

$$\overline{s}_t = \mu + C\overline{x}_{t-1}. \tag{C3}$$

Equation (10) then implies that

$$\overline{x}_t = A\overline{x}_{t-1} + G_t(s_t - \overline{s}_t) = A\overline{x}_{t-1} + G_t\zeta_t.$$

Iterating between $t-1$ and zero, we find

$$\overline{x}_{t-1} = \sum_{t'=1}^{t-1} A^{t-t'-1} G_{t'}\zeta_{t'}. \tag{C4}$$

Plugging into equation (C3), we find equation (C1).

Consider next the agent's recursive-filtering problem under $\tilde{p}$. Equation (10) implies that

$$\overline{x}_t(\tilde{p}) = (\tilde{A} - \tilde{G}_t\tilde{C})\overline{x}_{t-1}(\tilde{p}) + \tilde{G}_t(s_t - \tilde{\mu}).$$

Iterating between $t-1$ and zero, we find

$$\begin{aligned} \overline{x}_{t-1}(\tilde{p}) &= \sum_{t'=1}^{t-1} \tilde{J}_{t-1,t'+1}\tilde{G}_{t'}(s_{t'} - \tilde{\mu}) \qquad \text{(C5)} \\ &= \sum_{t'=1}^{t-1} \tilde{J}_{t-1,t'+1}\tilde{G}_{t'}(\zeta_{t'} + \mu - \tilde{\mu} + C\overline{x}_{t'-1}), \end{aligned}$$

where the second step follows from $s_{t'} = \zeta_{t'} + \overline{s}_{t'}$ and equation (C3). Substituting $\overline{x}_{t'-1}$ from equation (C4), and grouping terms, we find

$$\overline{x}_{t-1}(\tilde{p}) = \sum_{t'=1}^{t-1} \tilde{M}_{t,t'}\zeta_{t'} + \tilde{M}_t^{\mu}(\mu - \tilde{\mu}). \tag{C6}$$

Combining this with

$$\overline{s}_t(\tilde{p}) = \tilde{\mu} + \tilde{C}\overline{x}_{t-1}(\tilde{p}) \tag{C7}$$

(which follows from equation (9)), we find equation (C2). ‖

We next prove Lemma 3. While this Lemma is stated after Theorem 1 and Lemmas 1 and 2, its proof does not rely on these results.

Proof of Lemma 3. Lemma C1 implies that

$$\overline{s}_t(\tilde{p}) - \overline{s}_t = \sum_{t'=1}^{t-1} e_{t,t'}\zeta_{t'} + N_t^{\mu}(\tilde{\mu} - \mu), \tag{C8}$$

where

$$e_{t,t'} \equiv \tilde{C}\tilde{M}_{t,t'} - CA^{t-t'-1}G_{t'},$$

$$N_t^{\mu} \equiv 1 - \tilde{C}\tilde{M}_t^{\mu}.$$

Therefore,

$$[\overline{s}_t(p) - \overline{s}_t]^2 = \sum_{t',t''=1}^{t-1} e_{t,t'}e_{t,t''}\zeta_{t'}\zeta_{t''} + (N_t^{\mu})^2(\tilde{\mu} - \mu)^2 + 2\sum_{t'=1}^{t-1} e_{t,t'}N_t^{\mu}\zeta_{t'}(\tilde{\mu} - \mu). \tag{C9}$$

Since the sequence $\{\zeta_{t'}\}_{t'=1,\dots,t-1}$ is independent under the true measure and mean-zero, we have

$$E\left[\bar{s}_t(p)-\bar{s}_t\right]^2=\sum_{t'=1}^{t-1}e_{t,t'}^2\sigma_{s,t'}^2+(N_t^{\mu})^2(\tilde{\mu}-\mu)^2. \tag{C10}$$

We first determine the limit of $\sum_{t'=1}^{t-1}e_{t,t'}^2\sigma_{s,t'}^2$ when t goes to ∞. Defining the double sequence $\{\phi_{k,t}\}_{k,t\geq 1}$ by

$$\phi_{k,t}\equiv\begin{cases} e_{t,t-k}^2\sigma_{s,t-k}^2 & \text{for } k=1,\dots,t-1,\\ 0 & \text{for } k>t-1,\end{cases},$$

we have

$$\sum_{t'=1}^{t-1}e_{t,t'}^2\sigma_{s,t'}^2=\sum_{k=1}^{t-1}e_{t,t-k}^2\sigma_{s,t-k}^2=\sum_{k=1}^{\infty}\phi_{k,t}.$$

The definitions of $e_{t,t'}$ and $\tilde{M}_{t,t'}$ imply that

$$e_{t,t-k}=\tilde{C}\tilde{J}_{t-1,t-k+1}\tilde{G}_{t-k}+\sum_{k'=1}^{k-1}\tilde{C}\tilde{J}_{t-1,t-k+k'+1}\tilde{G}_{t-k+k'}CA^{k'-1}G_{t-k}-CA^{k-1}G_{t-k}. \tag{C11}$$

Equation (9) applied to the recursive-filtering problem under the true model implies that

$$\sigma_{s,t}^2=C\Sigma_{t-1}C'+V.$$

When t goes to ∞, G_t goes to G, $\tilde{G}_t$ to $\tilde{G}$, Σ_t to Σ, and $\tilde{J}_{t,t-k}$ to $\tilde{D}^{k+1}$. Therefore,

$$\lim_{t\to\infty}e_{t,t-k}=\tilde{C}\tilde{D}^{k-1}\tilde{G}+\sum_{k'=1}^{k-1}\tilde{C}\tilde{D}^{k-1-k'}\tilde{G}CA^{k'-1}G-CA^{k-1}G=\tilde{N}_k-N_k\equiv e_k,$$

$$\lim_{t\to\infty}\sigma_{s,t-k}^2=C\Sigma C'+V=C\Sigma C'+(1-\rho)^2\sigma_{\eta}^2+\sigma_{\omega}^2=\sigma_s^2, \tag{C12}$$

implying that

$$\lim_{t\to\infty}\phi_{k,t}=e_k^2\sigma_s^2.$$

The dominated convergence theorem will imply that

$$\lim_{t\to\infty}\sum_{k=1}^{\infty}\phi_{k,t}=\sum_{k=1}^{\infty}\lim_{t\to\infty}\phi_{k,t}=\sigma_s^2\sum_{k=1}^{\infty}e_k^2, \tag{C13}$$

if there exists a sequence $\{\overline{\phi}_k\}_{k\geq 1}$ such that $\sum_{k=1}^{\infty}\overline{\phi}_k<\infty$ and $|\phi_{k,t}|\leq\overline{\phi}_k$ for all $k,t\geq 1$. To construct such a sequence, we note that the eigenvalues of A have modulus smaller than 1, and so do the eigenvalues of $\tilde{D}\equiv\tilde{A}-\tilde{G}\tilde{C}$ (Balakrishnan, Theorem 4.2.3, p. 111). Denoting by $a<1$ a number exceeding the maximum of the moduli, we can construct a dominating sequence $\{\overline{\phi}_k\}_{k\geq 1}$ decaying geometrically at the rate a^{2k}.

We next determine the limit of N_t^{μ}. Defining the double sequence $\{\chi_{k,t}\}_{k,t\geq 1}$ by

$$\chi_{k,t}\equiv\begin{cases} \tilde{J}_{t-1,t-k+1}\tilde{G}_{t-k} & \text{for } k=1,\dots,t-1,\\ 0 & \text{for } k>t-1,\end{cases}$$

we have

$$N_t^{\mu}=1-\tilde{C}\sum_{k=1}^{t-1}\tilde{J}_{t-1,t-k+1}\tilde{G}_{t-k}=1-\tilde{C}\sum_{k=1}^{\infty}\chi_{k,t}.$$

It is easy to check that the dominated convergence theorem applies to $\{\chi_{k,t}\}_{k,t\geq 1}$, and thus

$$\lim_{t\to\infty} N_t^{\mu} = 1 - \tilde{C} \lim_{t\to\infty}\left[\sum_{k=1}^{\infty}\chi_{k,t}\right] = 1 - \tilde{C}\sum_{k=1}^{\infty}\lim_{t\to\infty}\chi_{k,t} = 1 - \tilde{C}\sum_{k=1}^{\infty}\tilde{D}^{k-1}\tilde{G} = N^{\mu}. \tag{C14}$$

The lemma follows by combining equations (C10), (C13), and (C14). ∥

Proof of Theorem 1. Equation (15) implies that

$$\frac{2\log L_t(\mathcal{H}_t|\tilde{p})}{t} = -\frac{\sum_{t'=1}^{t}\log\left[2\pi\sigma_{s,t'}^2(\tilde{p})\right]}{t} - \frac{1}{t}\sum_{t'=1}^{t}\frac{[s_{t'} - \overline{s}_{t'}(\tilde{p})]^2}{\sigma_{s,t'}^2(\tilde{p})}. \tag{C15}$$

To determine the limit of the first term, we note that equation (9) applied to the agent's recursive-filtering problem under $\tilde{p}$ implies that

$$\sigma_{s,t}^2(\tilde{p}) = \tilde{C}\tilde{\Sigma}_{t-1}\tilde{C}' + \tilde{V}.$$

Therefore,

$$\lim_{t\to\infty}\sigma_{s,t}^2(\tilde{p}) = \tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V} = \sigma_s^2(\tilde{p}), \tag{C16}$$

$$\lim_{t\to\infty}\frac{\sum_{t'=1}^{t}\log\sigma_{s,t'}^2(\tilde{p})}{t} = \lim_{t\to\infty}\log\sigma_{s,t}^2(\tilde{p}) = \log\sigma_s^2(\tilde{p}). \tag{C17}$$

We next fix $k \geq 0$ and determine the limit of the sequence

$$S_{k,t} \equiv \frac{1}{t}\sum_{t'=1}^{t}\zeta_{t'}\zeta_{t'+k}$$

when t goes to ∞. This sequence involves averages of random variables that are non-independent and non-identically distributed. An appropriate law of large numbers (LLN) for such sequences is that of McLeish (1975). Consider a probability space $(\Omega, \mathcal{F}, P)$, a sequence $\{\mathcal{F}_t\}_{t\in\mathbb{Z}}$ of σ-algebras, and a sequence $\{U_t\}_{t\geq 1}$ of random variables. The pair $\left(\{\mathcal{F}_t\}_{t\in\mathbb{Z}}, \{U_t\}_{t\geq 1}\right)$ is a mixingale (McLeish, 1975, Definition 1.2, p. 830) if and only if there exist sequences $\{c_t\}_{t\geq 1}$ and $\{\psi_m\}_{m\geq 0}$ of non-negative constants, with $\lim_{m\to\infty}\psi_m = 0$, such that for all $t \geq 1$ and $m \geq 0$:

$$\|E_{t-m}U_t\|_2 \leq \psi_m c_t, \tag{C18}$$

$$\|U_t - E_{t+m}U_t\|_2 \leq \psi_{m+1}c_t, \tag{C19}$$

where $\|.\|_2$ denotes the L_2 norm, and $E_{t'}U_t$ the expectation of U_t conditional on $\mathcal{F}_{t'}$. McLeish's LLN (Corollary 1.9, p. 832) states that if $\left(\{\mathcal{F}_t\}_{t\in\mathbb{Z}}, \{U_t\}_{t\geq 1}\right)$ is a mixingale, then

$$\lim_{t\to\infty}\frac{1}{t}\sum_{t'=1}^{t}U_{t'} = 0$$

almost surely, provided that $\sum_{t=1}^{\infty}c_t^2/t^2 < \infty$ and $\sum_{m=1}^{\infty}\psi_m < \infty$. In our model, we take the probability measure to be the true measure, and define the sequence $\{\mathcal{F}_t\}_{t\in\mathbb{Z}}$ as follows: $\mathcal{F}_t = \{\Omega, \emptyset\}$ for $t \leq 0$, and $\mathcal{F}_t$ is the σ-algebra generated by $\{\zeta_{t'}\}_{t'=1,\ldots,t}$ for $t \geq 1$. Moreover, we set $U_t \equiv \zeta_t^2 - \sigma_{s,t}^2$ when $k = 0$, and $U_t \equiv \zeta_t\zeta_{t+k}$ when $k \geq 1$. Since the sequence $\{\zeta_t\}_{t\geq 1}$ is independent, we have $E_{t-m}U_t = 0$ for $m \geq 1$. We also trivially have $E_{t+m}U_t = U_t$ for $m \geq k$. Therefore, when $k = 0$, equations (C18) and (C19) hold with $\psi_0 = 1$, $\psi_m = 0$ for $m \geq 1$, and $c_t = \sup_{t\geq 1}\|\zeta_t^2 - \sigma_{s,t}^2\|_2$ for $t \geq 1$. McLeish's LLN implies that

$$\lim_{t\to\infty}S_{0,t} = \lim_{t\to\infty}\frac{1}{t}\sum_{t'=1}^{t}\left(U_{t'} + \sigma_{s,t'}^2\right) = \lim_{t\to\infty}\frac{\sum_{t'=1}^{t}\sigma_{s,t'}^2}{t} = \lim_{t\to\infty}\sigma_{s,t}^2 = \sigma_s^2 \tag{C20}$$

almost surely. When $k \geq 1$, equations (C18) and (C19) hold with $\psi_m = 1$ for $m = 0, \ldots, k-1$, $\psi_m = 0$ for $m \geq k$, and $c_t = \sup_{t \geq 1} \|\zeta_t\|_2^2$ for $t \geq 1$. McLeish's LLN implies that

$$\lim_{t\to\infty} S_{k,t} = \lim_{t\to\infty} \frac{1}{t} \sum_{t'=1}^{t} U_{t'} = 0 \tag{C21}$$

almost surely. Finally, a straightforward application of McLeish's LLN to the sequence $U_t \equiv \zeta_t$ implies that

$$\lim_{t\to\infty} \frac{1}{t} \sum_{t'=1}^{t} \zeta_{t'} = 0 \tag{C22}$$

almost surely. Since $\mathbb{N}$ is countable, we can assume that equations (C20), (C21) for all $k \geq 1$, and equation (C22), hold in the same measure-one set. In what follows, we consider histories in that set.
To determine the limit of the second term in equation (C15), we write it as

$$\underbrace{\frac{1}{t}\sum_{t'=1}^{t} \frac{\zeta_{t'}^2}{\sigma_{s,t'}^2(\tilde{p})}}_{X_t} - \underbrace{\frac{2}{t}\sum_{t'=1}^{t} \frac{\zeta_{t'}[\bar{s}_{t'}(\tilde{p}) - \bar{s}_{t'}]}{\sigma_{s,t'}^2(\tilde{p})}}_{Y_t} + \underbrace{\frac{1}{t}\sum_{t'=1}^{t} \frac{[\bar{s}_{t'}(\tilde{p}) - \bar{s}_{t'}]^2}{\sigma_{s,t'}^2(\tilde{p})}}_{Z_t}.$$

Since $\lim_{t\to\infty} \sigma_{s,t}^2(\tilde{p}) = \sigma_s^2(\tilde{p})$, we have

$$\lim_{t\to\infty} X_t = \frac{1}{\sigma_s^2(\tilde{p})} \lim_{t\to\infty} \frac{1}{t} \sum_{t'=1}^{t} \zeta_{t'}^2 = \frac{1}{\sigma_s^2(\tilde{p})} \lim_{t\to\infty} S_{0,t} = \frac{\sigma_s^2}{\sigma_s^2(\tilde{p})}. \tag{C23}$$

Using equation (C8), we can write Y_t as

$$Y_t = 2\sum_{k=1}^{\infty} \psi_{k,t} + \frac{2}{t}\sum_{t'=1}^{t} \frac{N_{t'}^{\mu}}{\sigma_{s,t'}^2(\tilde{p})} \zeta_{t'}(\tilde{\mu} - \mu),$$

where the double sequence $\{\psi_{k,t}\}_{k,t\geq 1}$ is defined by

$$\psi_{k,t} \equiv \begin{cases} \frac{1}{t}\sum_{t'=k+1}^{t} \frac{e_{t',t'-k}}{\sigma_{s,t'}^2(\tilde{p})} \zeta_{t'}\zeta_{t'-k} & \text{for } k = 1, \ldots, t-1, \\ 0 & \text{for } k > t-1. \end{cases}$$

Since $\lim_{t\to\infty} e_{t,t-k} = e_k$ and $\lim_{t\to\infty} N_t^{\mu} = N^{\mu}$, we have

$$\lim_{t\to\infty} \psi_{k,t} = \frac{e_k}{\sigma_s^2(\tilde{p})} \lim_{t\to\infty} \frac{1}{t} \sum_{t'=k+1}^{t} \zeta_{t'}\zeta_{t'-k} = \frac{e_k}{\sigma_s^2(\tilde{p})} \lim_{t\to\infty} S_{k,t} = 0,$$

$$\lim_{t\to\infty} \frac{1}{t}\sum_{t'=1}^{t} \frac{N_{t'}^{\mu}}{\sigma_{s,t'}^2(\tilde{p})} \zeta_{t'}(\tilde{\mu} - \mu) = \frac{N^{\mu}}{\sigma_s^2(\tilde{p})}(\tilde{\mu} - \mu) \lim_{t\to\infty} \frac{1}{t}\sum_{t'=1}^{t} \zeta_{t'} = 0.$$

The dominated convergence theorem will imply that

$$\lim_{t\to\infty} Y_t = 2\lim_{t\to\infty} \sum_{k=1}^{\infty} \psi_{k,t} = 2\sum_{k=1}^{\infty} \lim_{t\to\infty} \psi_{k,t} = 0 \tag{C24}$$

if there exists a sequence $\{\overline{\psi}_k\}_{k\geq 1}$ such that $\sum_{k=1}^{\infty} \overline{\psi}_k < \infty$ and $|\psi_{k,t}| \leq \overline{\psi}_k$ for all $k, t \geq 1$. Such a sequence can be constructed by the same argument as for $\phi_{k,t}$ (Lemma 3) since

$$\left| \frac{1}{t} \sum_{t'=k+1}^{t} \zeta_{t'}\zeta_{t'-k} \right| \leq \frac{1}{t} \sum_{t'=k+1}^{t} |\zeta_{t'}|\,|\zeta_{t'-k}| \leq \sqrt{\frac{1}{t}\sum_{t'=k+1}^{t} \zeta_{t'}^2} \sqrt{\frac{1}{t}\sum_{t'=k+1}^{t} \zeta_{t'-k}^2} \leq \sup_{t\geq 1} S_{0,t} < \infty,$$

where the last inequality holds because the sequence $S_{0,t}$ is convergent.

Using similar arguments as for Y_t, we find

$$\lim_{t\to\infty} Z_t = \frac{\sigma_s^2}{\sigma_s^2(\tilde{p})}\sum_{k=1}^{\infty} e_k^2 + \frac{(N^\mu)^2}{\sigma_s^2(\tilde{p})}(\tilde{\mu}-\mu)^2 = \frac{e(\tilde{p})}{\sigma_s^2(\tilde{p})}. \tag{C25}$$

The theorem follows from equations (C15), (C17), (C23), (C24), and (C25). ‖

Proof of Lemma 1. We first show that the set $m(P)$ is non-empty. Equation (C16) implies that when $\tilde{\sigma}_\eta^2$ or $\tilde{\sigma}_\omega^2$ go to ∞, $\sigma_s^2(\tilde{p})$ goes to ∞ and $F(\tilde{p})$ goes to $-\infty$. Therefore, we can restrict the maximization of $F(\tilde{p})$ to bounded values of $(\tilde{\sigma}_\eta^2, \tilde{\sigma}_\omega^2)$. Equation (20) implies that

$$N^\mu = 1 - \tilde{C}(I-\tilde{D})^{-1}\tilde{G} = 1 - \tilde{C}(I-\tilde{A}+\tilde{G}\tilde{C})^{-1}\tilde{G}.$$

Replacing $\tilde{A}$ and $\tilde{C}$ by their values, and denoting the components of $\tilde{G}$ by $\tilde{G}_1$ and $\tilde{G}_2$, we find

$$N^\mu = 1 - \tilde{C}(I-\tilde{A}+\tilde{G}\tilde{C})^{-1}\tilde{G} = \frac{(1-\tilde{\rho})(1-\delta_{\tilde{\rho}}+\alpha_{\tilde{\rho}})}{[1-\tilde{\rho}(1-\tilde{G}_1)](1-\delta_{\tilde{\rho}}+\alpha_{\tilde{\rho}}) - \alpha_{\tilde{\rho}}(1-\tilde{\rho})\tilde{G}_2}. \tag{C26}$$

Since $\alpha_{\tilde{\rho}}, \delta_{\tilde{\rho}} \in [0,1)$ and $\tilde{\rho} \in [0,\overline{\rho}]$, $N^\mu \neq 0$. Therefore, when $|\tilde{\mu}|$ goes to ∞, Lemma 3 implies that $e(\tilde{p})$ goes to ∞ and $F(\tilde{p})$ goes to $-\infty$. This means that we can restrict the maximization of $F(\tilde{p})$ to bounded values of $\tilde{\mu}$, and thus to a compact subset of P. We can also assume that $F(\tilde{p})$ is continuous in that subset since the only point of discontinuity is for $\sigma_s^2(\tilde{p}) = 0$, in which case $F(\tilde{p}) = -\infty$. Therefore, $F(\tilde{p})$ has a maximum and the set $m(P)$ is non-empty.

To show that the measure π_t converges weakly to a measure giving weight only to $m(P)$, it suffices to show that for all closed sets S having zero intersection with $m(P)$, $\pi_t(S)$ goes to zero (Billingsley, 1986, Theorem 29.1, p. 390). The maximum F_S of $F(\tilde{p})$ over S is the same as in a compact subset of S and is thus smaller than the value of $F(\tilde{p})$ in $m(P)$. Consider a compact neighbourhood B of a point in $m(P)$ such that the minimum F_B of $F(\tilde{p})$ over B exceeds F_S. Consider also two constants (F_1, F_2) such that $F_B > F_2 > F_1 > F_S$. For large enough t,

$$\min_{\tilde{p}\in B} \frac{\log L_t(\mathcal{H}_t|\tilde{p})}{t} > F_2. \tag{C27}$$

Indeed, if equation (C27) does not hold, there exists a convergent sequence $\{\tilde{p}_t\}_{t\geq 1}$ in B such that

$$\frac{\log L_t(\mathcal{H}_t|\tilde{p}_t)}{t} \leq F_2.$$

Denoting the limit of this sequence by $\tilde{p} \in B$, Theorem 1 implies that $F(\tilde{p}) \leq F_2$, a contradiction. (Theorem 1 concerns the convergence of the likelihood for a given $\tilde{p}$, but extending the argument to a convergent sequence $\{\tilde{p}_t\}_{t\geq 1}$ is straightforward.) Likewise, we can show that for large enough t,

$$\max_{\tilde{p}\in S} \frac{\log L_t(\mathcal{H}_t|\tilde{p})}{t} < F_1. \tag{C28}$$

Bayes' law, equation (C27), and equation (C28) imply that for large enough t,

$$\pi_t(S) = \frac{E_{\pi_0}\left[L_t(\mathcal{H}_t|\tilde{p})1_{\{\tilde{p}\in S\}}\right]}{E_{\pi_0}\left[L_t(\mathcal{H}_t|\tilde{p})\right]} < \frac{E_{\pi_0}\left[L_t(\mathcal{H}_t|\tilde{p})1_{\{\tilde{p}\in S\}}\right]}{E_{\pi_0}\left[L_t(\mathcal{H}_t|\tilde{p})1_{\{\tilde{p}\in B\}}\right]} < \frac{\exp(tF_1)\pi_0(S)}{\exp(tF_2)\pi_0(B)}.$$

Since $F_2 > F_1$, $\pi_t(S)$ goes to zero when t goes to ∞. ‖

Proof of Lemma 2. Consider $\tilde{p} \in P$ such that $e(\tilde{p}) = e(P)$ and $\sigma_s^2(\tilde{p}) = \sigma_s^2 + e(\tilde{p})$. We will show that $F(\tilde{p}) \geq F(\hat{p})$ for any $\hat{p} = (\hat{\sigma}_\eta^2, \hat{\rho}, \hat{\sigma}_\omega^2, \hat{\mu}) \in P$. Denote by $\hat{\Sigma}$ and $\hat{G}$ the steady-state variance and regression coefficient for the recursive-filtering problem under $\hat{p}$, and by $\hat{\Sigma}_\lambda$ and $\hat{G}_\lambda$ those under $\hat{p}_\lambda \equiv (\lambda\hat{\sigma}_\eta^2, \hat{\rho}, \lambda\hat{\sigma}_\omega^2, \hat{\mu})$ for $\lambda > 0$. It is easy to check that $\lambda\hat{\Sigma}$ solves equation (13) for $\hat{p}_\lambda$. Since this equation has a unique solution, $\hat{\Sigma}_\lambda = \lambda\hat{\Sigma}$. Equation (12) then implies that $\hat{G}_\lambda = \hat{G}$, and equations (17) and (C16) imply that $e(\hat{p}_\lambda) = e(\hat{p})$ and $\sigma_s^2(\hat{p}_\lambda) = \lambda\sigma_s^2(\hat{p})$. Therefore,

$$F(\hat{p}_\lambda) = -\frac{1}{2}\left[\log\left[2\pi\lambda\sigma_s^2(\hat{p})\right] + \frac{\sigma_s^2 + e(\hat{p})}{\lambda\sigma_s^2(\hat{p})}\right]. \tag{C29}$$

Since this function is maximized for

$$\lambda^* = \frac{\sigma_s^2 + e(\hat{p})}{\sigma_s^2(\hat{p})},$$

we have

$$F(\hat{p}) \leq F(\hat{p}_{\lambda^*}) = -\frac{1}{2}\left[\log\left[2\pi\left[\sigma_s^2 + e(\hat{p})\right]\right] + 1\right] \leq -\frac{1}{2}\left[\log\left[2\pi\left[\sigma_s^2 + e(\tilde{p})\right]\right] + 1\right] = F(\tilde{p}).$$

The proof of the converse is along the same lines. ‖

Lemma C2 determines when a model can predict the signals equally well as the true model.

Lemma C2. *The error $e(\tilde{p})$ is zero if and only if*

- $\tilde{C}\tilde{A}^{k-1}\tilde{G} = CA^{k-1}G$ *for all* $k \geq 1$
- $\tilde{\mu} = \mu$.

Proof. From Lemma 3 and $N^{\mu} \neq 0$, it suffices to show that $\{e_k\}_{k\geq 1} = 0$ is equivalent to $\tilde{C}\tilde{A}^{k-1}\tilde{G} = CA^{k-1}G$ for all $k \geq 1$. Setting $a_k \equiv \tilde{C}\tilde{A}^{k-1}\tilde{G} - CA^{k-1}G$ and

$$b_k \equiv \tilde{D}^{k-1}\tilde{G} + \sum_{k'=1}^{k-1} \tilde{D}^{k-1-k'}\tilde{G}CA^{k'-1}G - \tilde{A}^{k-1}\tilde{G},$$

we have $e_k = \tilde{C}b_k + a_k$ for $k \geq 1$. Simple algebra shows that

$$b_k = \tilde{D}b_{k-1} - \tilde{G}a_{k-1}.$$

Iterating between k and 1, and using the initial condition $b_1 = 0$, we find

$$b_k = -\sum_{k'=1}^{k-1} \tilde{D}^{k-1-k'}\tilde{G}a_{k'}.$$

Therefore,

$$e_k = -\sum_{k'=1}^{k-1} \tilde{C}\tilde{D}^{k-1-k'}\tilde{G}a_{k'} + a_k. \tag{C30}$$

Equation (C30) implies that $\{e_k\}_{k\geq 1} = 0$ if and only if $\{a_k\}_{k\geq 1} = 0$. ‖

Proof of Proposition 3. Under rational updating, it is possible to achieve minimum error $e(P_0) = 0$ by using the vector of true parameters p. Since $e(P_0) = 0$, Lemmas 2 and C2 imply that $\tilde{p} \in m(P_0)$ if and only if (i) $\tilde{C}\tilde{A}^{k-1}\tilde{G} = CA^{k-1}G$ for all $k \geq 1$, (ii) $\tilde{\mu} = \mu$, and (iii) $\sigma_s^2(\tilde{p}) = \sigma_s^2$. Since $\alpha = 0$, we can write Condition (i) as

$$\tilde{\rho}^k\tilde{G}_1 = \rho^k G_1. \tag{C31}$$

We can also write element (1,1) of equation (13) as

$$\tilde{\Sigma}_{11} = \frac{\left[\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2\right]\tilde{\sigma}_\omega^2}{\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2}, \tag{C32}$$

$$\Sigma_{11} = \frac{\left[\rho^2\Sigma_{11} + (1-\rho)^2\sigma_\eta^2\right]\sigma_\omega^2}{\rho^2\Sigma_{11} + (1-\rho)^2\sigma_\eta^2 + \sigma_\omega^2}, \tag{C33}$$

and the first element of equation (14) as

$$\tilde{G}_1 = \frac{\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2}{\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2}, \tag{C34}$$

$$G_1 = \frac{\rho^2\Sigma_{11} + (1-\rho)^2\sigma_\eta^2}{\rho^2\Sigma_{11} + (1-\rho)^2\sigma_\eta^2 + \sigma_\omega^2}, \tag{C35}$$

where the first equation in each case is for $\tilde{p}$ and the second for p. Using equations (C12) and (C16), we can write Condition (iii) as

$$\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \rho^2\Sigma_{11} + (1-\rho)^2\sigma_\eta^2 + \sigma_\omega^2. \tag{C36}$$

Suppose that $\rho\sigma_\eta^2 > 0$, and consider $\tilde{p}$ that satisfies Conditions (i)–(iii). Equation (C35) implies that $G_1 > 0$. Since equation (C31) must hold for all $k \geq 1$, we have $\tilde{\rho} = \rho$ and $\tilde{G}_1 = G_1$. We next write equations (C32)–(C35) in terms of the normalized variables $\tilde{s}_\eta^2 \equiv \tilde{\sigma}_\eta^2/\tilde{\sigma}_\omega^2$, $\tilde{S}_{11} \equiv \tilde{\Sigma}_{11}/\tilde{\sigma}_\omega^2$, $s_\eta^2 \equiv \sigma_\eta^2/\sigma_\omega^2$, and $S_{11} \equiv \Sigma_{11}/\sigma_\omega^2$. Equations (C32) and (C33) imply that $\tilde{S}_{11} = g(\tilde{s}_\eta^2)$ and $S_{11} = g(s_\eta^2)$ for the same function g. Equations (C34), (C35), and $\tilde{G}_1 = G_1$ then imply that $\tilde{s}_\eta^2 = s_\eta^2$, and equation (C36) implies that $\tilde{\sigma}_\omega^2 = \sigma_\omega^2$. Thus, $\tilde{p} = p$.

Suppose next that $\rho\sigma_\eta^2 = 0$, and consider $\tilde{p}$ that satisfies Conditions (i)–(iii). If $\rho = 0$, equation (C31) implies that $\tilde{\rho}^k\tilde{G}_1 = 0$, and equation (C36) that $\tilde{\rho}^2\tilde{\Sigma}_{11} + (1-\tilde{\rho})^2\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \sigma_\eta^2 + \sigma_\omega^2$. If $\sigma_\eta^2 = 0$, the same implications follow because $\Sigma = 0$ and $G = [0, 1]'$ from equations (13) and (14). Equation $\tilde{\rho}^k\tilde{G}_1 = 0$ implies that either $\tilde{\rho} = 0$, or $\tilde{G}_1 = 0$ in which case $\tilde{\sigma}_\eta^2 = 0$. If $\tilde{\rho} = 0$, then $\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \sigma_\eta^2 + \sigma_\omega^2$. If $\tilde{\sigma}_\eta^2 = 0$, then $\tilde{\sigma}_\omega^2 = \sigma_\eta^2 + \sigma_\omega^2$. Therefore, $\tilde{p}$ is as in the proposition. Showing that all $\tilde{p}$ in the proposition satisfy Conditions (i)–(iii) is obvious. ||

Proof of Proposition 4. We determine the parameter vectors $\tilde{p}$ that belong to $m(P_0)$ and satisfy $e(\tilde{p}) = 0$. From Lemmas 2 and C2, these must satisfy (i) $\tilde{C}\tilde{A}^{k-1}\tilde{G} = CA^{k-1}G$ for all $k \geq 1$, (ii) $\tilde{\mu} = \mu$, and (iii) $\sigma_s^2(\tilde{p}) = \sigma_s^2$. Since $\sigma_\eta^2 = 0$, we can write Condition (i) as

$$\tilde{\rho}^k\tilde{G}_1 - \alpha_{\tilde{\rho}}(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^{k-1}\tilde{G}_2 = 0, \tag{C37}$$

and Condition (iii) as

$$\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V} = \sigma_\omega^2. \tag{C38}$$

Using $(\alpha_{\tilde{\rho}}, \delta_{\tilde{\rho}}) = (\alpha\tilde{\rho}, \delta\tilde{\rho})$, we can write equation (C37) as

$$\tilde{\rho}^k\left[\tilde{G}_1 - \alpha(\delta-\alpha)^{k-1}\tilde{G}_2\right] = 0. \tag{C39}$$

If $\tilde{\rho} \neq 0$, equation (C39) implies that $\tilde{G}_1 - \alpha(\delta-\alpha)^{k-1}\tilde{G}_2 = 0$ for all $k \geq 1$, which in turn implies that $\tilde{G} = 0$. For $\tilde{G} = 0$, equation (13) becomes $\tilde{\Sigma} = \tilde{A}\tilde{\Sigma}\tilde{A}' + \tilde{W}$. Solving for $\tilde{\Sigma}$, we find

$$\begin{aligned}
\tilde{\Sigma}_{11} &= \frac{(1-\tilde{\rho})\tilde{\sigma}_\eta^2}{1+\tilde{\rho}}, \\
\tilde{\Sigma}_{12} &= 0, \\
\tilde{\Sigma}_{22} &= \frac{\tilde{\sigma}_\omega^2}{1-(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^2}.
\end{aligned}$$

Substituting $\tilde{\Sigma}$ into equation (14), we find $(\tilde{\sigma}_\eta^2, \tilde{\sigma}_\omega^2) = 0$. But then, $\tilde{\Sigma} = 0$, which contradicts equation (C38) since $\sigma_\omega^2 = \sigma_\epsilon^2 > 0$. Therefore, $\tilde{\rho} = 0$. Equation (C38) then implies that $\tilde{\sigma}_\eta^2 + \tilde{\sigma}_\omega^2 = \sigma_\omega^2$. Showing that all $\tilde{p}$ in the proposition satisfy Conditions (i)–(iii) is obvious. ||

Proof of Proposition 5. Consider a sequence $\{\alpha_n\}_{n\in\mathbb{N}}$ converging to zero, and an element $\tilde{p}_n \equiv ((\tilde{\sigma}_\eta^2)_n, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ from the set $m(P_{\underline{\rho}})$ corresponding to α_n. The proposition will follow if we show that $(\frac{(\tilde{s}_\eta^2)_n}{\alpha_n}, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ converges to $(z, \underline{\rho}, \sigma_\omega^2, \mu)$, where $(\tilde{s}_\eta^2)_n \equiv (\tilde{\sigma}_\eta^2)_n/(\tilde{\sigma}_\omega^2)_n$. Denoting the limits of $(\frac{(\tilde{s}_\eta^2)_n}{\alpha_n}, (\tilde{\sigma}_\eta^2)_n, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ by $(\ell_s, \ell_\eta, \ell_\rho, \ell_\omega, \ell_\mu)$, the point $(\ell_\eta, \ell_\rho, \ell_\omega, \ell_\mu)$ belongs to the set $m(P_{\underline{\rho}})$ derived for $\alpha = \sigma_\eta^2 = 0$. (If the sequences do not converge, we extract converging subsequences.) All elements in that set satisfy $\tilde{\mu} = \mu$. Proposition 3 implies that they also satisfy $\tilde{\sigma}_\eta^2 = 0$ and $\tilde{\sigma}_\omega^2 = \sigma_\omega^2$ since $\tilde{\rho} \geq \underline{\rho} > 0$. Therefore, $(\ell_\mu, \ell_\eta, \ell_\omega) = (\mu, 0, \sigma_\omega^2)$.

When $v \equiv (\alpha, \frac{\tilde{s}_\eta^2}{\alpha}, \tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2)$ converges to $\ell_v \equiv (0, \ell_s, 0, \ell_\rho, \sigma_\omega^2)$, $\tilde{C}\tilde{D}^k\tilde{G}$ converges to zero, $\frac{\tilde{G}_1}{\tilde{s}_\eta^2}$ to $\frac{1-\ell_\rho}{1+\ell_\rho}$, and $\tilde{G}_2$ to 1. These limits follow by continuity if we show that when $\alpha = \tilde{\sigma}_\eta^2 = 0$, $\tilde{C}\tilde{D}^k\tilde{G} = 0$ and $\tilde{G}_2 = 1$, and when $\alpha = 0$, $\lim_{\tilde{\sigma}_\eta^2\to 0}\frac{\tilde{G}_1}{\tilde{s}_\eta^2} = \frac{1-\tilde{\rho}}{1+\tilde{\rho}}$. When $\tilde{\sigma}_\eta^2 = 0$, the unique solution of equation (13) is $\tilde{\Sigma} = 0$, and equation (14) implies that $\tilde{G} = [0, 1]'$. Therefore, when $\alpha = \tilde{\sigma}_\eta^2 = 0$, we have $\tilde{C}\tilde{G} = 0$,

$$\tilde{C}\tilde{D} = \tilde{C}(\tilde{A} - \tilde{G}\tilde{C}) = \tilde{C}\tilde{A} = \tilde{\rho}\tilde{C},$$

and $\tilde{C}\tilde{D}^k\tilde{G} = \tilde{\rho}^k\tilde{C}\tilde{G} = 0$. Moreover, if we divide both sides of equations (C32) and (C34) (derived from equations (13) and (14) when $\alpha = 0$) by $\tilde{\sigma}_\omega^2$, we find

$$\tilde{S}_{11} = \frac{\tilde{\rho}^2\tilde{S}_{11} + (1-\tilde{\rho})^2\tilde{s}_\eta^2}{\tilde{\rho}^2\tilde{S}_{11} + (1-\tilde{\rho})^2\tilde{s}_\eta^2 + 1}, \tag{C40}$$

$$\tilde{G}_1 = \frac{\tilde{\rho}^2\tilde{S}_{11} + (1-\tilde{\rho})^2\tilde{s}_\eta^2}{\tilde{\rho}^2\tilde{S}_{11} + (1-\tilde{\rho})^2\tilde{s}_\eta^2 + 1}. \tag{C41}$$

When $\tilde{\sigma}_\eta^2$ converges to zero, $\tilde{s}_\eta^2$ and $\tilde{S}_{11}$ converge to zero. Equations (C40) and (C41) then imply that $\frac{\tilde{S}_{11}}{\tilde{s}_\eta^2}$ and $\frac{\tilde{G}_1}{\tilde{s}_\eta^2}$ converge to $\frac{1-\tilde{\rho}}{1+\tilde{\rho}}$.

Using the above limits, we find

$$\lim_{v\to\ell_v} \frac{\tilde{C}\tilde{A}^{k-1}\tilde{G}}{\alpha} = \lim_{v\to\ell_v} \frac{\tilde{\rho}^k\tilde{G}_1 - \alpha_{\tilde{\rho}}(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^{k-1}\tilde{G}_2}{\alpha}$$
$$= \lim_{v\to\ell_v} \tilde{\rho}^k \left[\frac{\tilde{s}_\eta^2}{\alpha}\frac{\tilde{G}_1}{\tilde{s}_\eta^2} - (\delta-\alpha)^{k-1}\tilde{G}_2\right] = \ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right].$$

Equations (18), (C30), $\lim_{v\to\ell_v}\tilde{C}\tilde{D}^k\tilde{G} = 0$, and $N_k = CA^kG = 0$, imply that

$$\lim_{v\to\ell_v}\frac{e_k}{\alpha} = \lim_{v\to\ell_v}\frac{\tilde{N}_k}{\alpha} = \lim_{v\to\ell_v}\frac{\tilde{C}\tilde{D}^{k-1}\tilde{G}}{\alpha} = \lim_{v\to\ell_v}\frac{a_k}{\alpha} = \lim_{v\to\ell_v}\frac{\tilde{C}\tilde{A}^{k-1}\tilde{G}}{\alpha} = \ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right]. \tag{C42}$$

Since $\tilde{p}_n$ minimizes $e(\tilde{p}_n)$, equation (17) implies that $((\sigma_\eta^2)_n, \tilde{\rho}_n, (\sigma_\omega^2)_n)$ minimizes $\sigma_s^2(\tilde{p})\sum_{k=0}^\infty e_k^2$. Since from equation (C42),

$$\lim_{v\to\ell_v}\frac{\sigma_s^2(\tilde{p})\sum_{k=1}^\infty e_k^2}{\alpha^2} = \sigma_\omega^2\sum_{k=1}^\infty \ell_\rho^{2k}\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right]^2 \equiv \sigma_\omega^2 F(\ell_s,\ell_\rho),$$

(ℓ_s,ℓ_ρ) must minimize F. Treating F as a function of $\left(\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}, \ell_\rho\right)$, the minimizing value of the second argument is clearly $\ell_\rho = \underline{\rho}$. The first-order condition with respect to the first argument is

$$\sum_{k=1}^\infty \ell_\rho^{2k}\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right] = 0. \tag{C43}$$

Substituting $\ell_\rho = \underline{\rho}$ into equation (C43), we find $\ell_s = z$. ‖

Proof of Proposition 6. Equations (C5) and (C7) imply that in steady state

$$\tilde{E}_{t-1}(s_t) \equiv \overline{s}_t(\tilde{p}) = \tilde{\mu} + \sum_{k=1}^\infty \tilde{C}\tilde{D}^{k-1}\tilde{G}(s_{t-k} - \tilde{\mu}). \tag{C44}$$

Therefore,

$$\tilde{\Delta}_k = \sum_{k'=1}^k \tilde{C}\tilde{D}^{k'-1}\tilde{G}.$$

Equation (C42) implies that

$$\tilde{\Delta}_k = \alpha g_k + o(\alpha),$$

where

$$g_k \equiv \sum_{k'=1}^k f_{k'},$$
$$f_k \equiv \underline{\rho}^k\left[\frac{z(1-\underline{\rho})}{1+\underline{\rho}} - \delta^{k-1}\right].$$

Using equation (24) to substitute z, we find

$$f_1 = g_1 = \frac{\underline{\rho}^3(\delta - 1)}{1 - \underline{\rho}^2\delta} < 0,$$

$$g_\infty = \underline{\rho}\left[\frac{z}{1+\underline{\rho}} - \frac{1}{1-\underline{\rho}\delta}\right] = \frac{\underline{\rho}^2(1-\delta)}{(1-\underline{\rho}^2\delta)(1-\underline{\rho}\delta)} > 0.$$

The function f_k is negative for $k = 1$ and positive for large k. Since it can change sign only once, it is negative and then positive. The function g_k is negative for $k = 1$, then decreases ($f_k < 0$), then increases ($f_k > 0$), and is eventually positive ($g_\infty > 0$). Therefore, g_k is negative and then positive. ‖

Proof of Proposition 7. Consider a sequence $\{\alpha_n\}_{n\in\mathbb{N}}$ converging to zero, an element $\tilde{p}_n \equiv ((\tilde{\sigma}_\eta^2)_n, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ from the set $m(P_0)$ corresponding to α_n, and set $(\sigma_\eta^2)_n \equiv \nu\alpha_n\sigma_\omega^2$. The proposition will follow if we show that $(\frac{(\tilde{s}_\eta^2)_n}{\alpha_n}, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ converges to $(z, r, \sigma_\omega^2, \mu)$. Denoting the limits of $(\frac{(\tilde{s}_\eta^2)_n}{\alpha_n}, (\tilde{\sigma}_\eta^2)_n, \tilde{\rho}_n, (\tilde{\sigma}_\omega^2)_n, \tilde{\mu}_n)$ by $(\ell_s, \ell_\eta, \ell_\rho, \ell_\omega, \ell_\mu)$, the point $(\ell_\eta, \ell_\rho, \ell_\omega, \ell_\mu)$ belongs to the set $m(P_0)$ derived for $\alpha = \sigma_\eta^2 = 0$. Proposition 3 implies that $\ell_\mu = \mu$. If $\ell_\rho > 0$, then Proposition 3 implies also that $(\ell_\eta, \ell_\omega) = (0, \sigma_\omega^2)$, and same arguments as in the proof of Proposition 5 imply that

$$\lim_{v\to\ell_v} \frac{\sigma_s^2(\tilde{p})\sum_{k=1}^\infty e_k^2}{\alpha^2} = \sigma_\omega^2 \sum_{k=1}^\infty \left(\ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right] - \rho^k\frac{\nu(1-\rho)}{1+\rho}\right)^2 \equiv \sigma_\omega^2 H(\ell_s, \ell_\rho).$$

If $\ell_\rho = 0$, then the limit is instead

$$\lim_{v\to\ell_v} \frac{\sigma_s^2(\tilde{p})\sum_{k=1}^\infty e_k^2}{\alpha^2} = \sigma_\omega^2\left(\left[\ell_\phi - \rho\frac{\nu(1-\rho)}{1+\rho}\right]^2 + \sum_{k=2}^\infty\left[\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2\right) \equiv \sigma_\omega^2 H_0(\ell_\phi),$$

where ℓ_ϕ denotes the limit of $\frac{\tilde{\rho}_n(\tilde{G}_1)_n}{\alpha_n}$. The cases $\ell_\rho > 0$ and $\ell_\rho = 0$ can be nested by noting that $H_0(\ell_\phi)$ is equal to H evaluated for $\ell_s\ell_\rho = \ell_\phi$ and $\ell_\rho = 0$. Treating H as a function of $\left(\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}, \ell_\rho\right)$, the first-order condition with respect to the first argument is

$$\sum_{k=1}^\infty \ell_\rho^k\left(\ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right] - \rho^k\frac{\nu(1-\rho)}{1+\rho}\right) = 0, \tag{C45}$$

and the derivative with respect to the second argument is

$$\sum_{k=1}^\infty k\ell_\rho^{k-1}\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right]\left(\ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1}\right] - \rho^k\frac{\nu(1-\rho)}{1+\rho}\right). \tag{C46}$$

Computing the infinite sums, we can write equation (C45) as

$$\frac{\ell_s\ell_\rho}{(1+\ell_\rho)^2} - \frac{\ell_\rho}{1-\ell_\rho^2\delta} - \frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho\ell_\rho)} = 0 \tag{C47}$$

and equation (C46) as

$$\begin{aligned}&\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}\left[\frac{\ell_s\ell_\rho}{(1+\ell_\rho)^2(1-\ell_\rho^2)} - \frac{\ell_\rho}{(1-\ell_\rho^2\delta)^2} - \frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho\ell_\rho)^2}\right]\\ &-\left[\frac{\ell_s\ell_\rho(1-\ell_\rho)}{(1+\ell_\rho)(1-\ell_\rho^2\delta)^2} - \frac{\ell_\rho}{(1-\ell_\rho^2\delta^2)^2} - \frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho\ell_\rho\delta)^2}\right].\end{aligned} \tag{C48}$$

Substituting ℓ_s from equation (C47), we can write the first square bracket in equation (C48) as

$$-\frac{\nu\rho(1-\rho)\ell_\rho(\rho-\ell_\rho)}{(1+\rho)(1-\rho\ell_\rho)^2(1-\ell_\rho^2)} + \frac{\ell_\rho^3(1-\delta)}{(1-\ell_\rho^2)(1-\ell_\rho^2\delta)}$$

and the second as

$$\frac{\nu\rho(1-\rho)\ell_\rho(\rho-\ell_\rho)\left[1-2\delta+(\rho+\ell_\rho)\ell_\rho\delta^2-\rho\ell_\rho^3\delta^2\right]}{(1+\rho)(1-\rho\ell_\rho)(1-\ell_\rho^2\delta)^2(1-\rho\ell_\rho\delta)^2}-\frac{\ell_\rho^3(1-\delta)\left[1-2\delta+\ell_\rho^2(1+\delta)\delta^2-\ell_\rho^4\delta^3\right]}{(1-\ell_\rho^2\delta)^3(1-\ell_\rho^2\delta^2)^2}.$$

We next substitute ℓ_s from equation (C47) into the term $\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}$ that multiplies the first square bracket in equation (C48). Grouping terms, we can write equation (C48) as

$$-\frac{\nu\rho(1-\rho)(\rho-\ell_\rho)}{(1+\rho)(1-\rho\ell_\rho)^2}H_1(\ell_\rho)+\frac{\ell_\rho^2(1-\delta)}{(1-\ell_\rho^2\delta)^2}H_2(\ell_\rho). \tag{C49}$$

The functions $H_1(\ell_\rho)$ and $H_2(\ell_\rho)$ are positive for $\ell_\rho \in [0,\overline{\rho}]$ because

$$2-\rho\ell_\rho(1+\delta)-\ell_\rho^2\delta+\rho^2\ell_\rho^4\delta^2>0, \tag{C50}$$

$$2-\ell_\rho^2\delta^2-\ell_\rho^4\delta^3>0. \tag{C51}$$

(To show equations (C50) and (C51), we use $\rho,\delta<1$. For equation (C50) we also note that the left-hand side is decreasing in δ and positive for $\delta=1$.) Since $H_1(\ell_\rho), H_2(\ell_\rho)>0$, equation (C49) is negative for $\ell_\rho=0$ and positive for $\ell_\rho\in[\rho,\overline{\rho}]$. Therefore, the value of ℓ_ρ that minimizes H is in $(0,\rho)$ and renders equation (C49) equal to zero. Comparing equations (26) and (C49), we find that this value is equal to r. Substituting $\ell_\rho=r$ into equation (C47), we find $\ell_s=z$. ‖

Proof of Proposition 8. Proceeding as in the proof of Proposition 6, we find

$$\tilde{\Delta}_k-\Delta_k=\alpha g_k+o(\alpha),$$

where

$$g_k \equiv \sum_{k'=1}^{k} f_{k'},$$

$$f_k \equiv r^k\left[\frac{z(1-r)}{1+r}-\delta^{k-1}\right]-\rho^k\frac{\nu(1-\rho)}{1+\rho}. \tag{C52}$$

The proposition will follow if we show that $f_1=g_1<0$ and $g_\infty<0$. Indeed, suppose that $f_1=g_1<0$. Since $r<\rho$, f_k is negative for large k. Moreover, equation (C45) can be written for $(\ell_s,\ell_\rho)=(z,r)$ as

$$\sum_{k=1}^{\infty} r^k f_k=0, \tag{C53}$$

implying that f_k has to be positive for some k. Since f_k can change sign at most twice (because the derivative of f_k/ρ^k can change sign at most once, implying that f_k/ρ^k can change sign at most twice), it is negative, then positive, and then negative again. The function g_k is negative for $k=1$, then decreases ($f_k<0$), then increases ($f_k>0$), and then decreases again ($f_k<0$). If $g_\infty<0$, then g_k can either be (i) always negative or (ii) negative, then positive, and then negative again. To rule out (i), we write equation (C53) as

$$g_1 r+\sum_{k=2}^{\infty}(g_k-g_{k-1})r^k=0 \Leftrightarrow \sum_{k=1}^{\infty} g_k\left(r^k-r^{k+1}\right)=0.$$

We next show that $f_1=g_1<0$ and $g_\infty<0$. Using equation (25) to substitute z, we find

$$f_1=g_1=\frac{\rho r(\rho-r)(1-\rho)(\nu-\nu_1)}{(1+\rho)(1-\rho r)},$$

$$g_\infty=\frac{zr}{1+r}-\frac{r}{1-r\delta}-\frac{\nu\rho}{1+\rho}=\frac{\rho(\rho-r)(\nu_\infty-\nu)}{(1+\rho)(1-\rho r)},$$

where

$$\nu_1 \equiv \frac{(1+\rho)(1-\rho r)r^2(1-\delta)}{\rho(\rho-r)(1-\rho)(1-r^2\delta)}$$

$$\nu_\infty \equiv \frac{(1+\rho)(1-\rho r)r^2(1-\delta)}{\rho(\rho-r)(1-r^2\delta)(1-r\delta)}.$$

We thus need to show that $\nu_1 > \nu > \nu_\infty$. These inequalities will follow if we show that when ν is replaced by ν_1 (resp. ν_∞) in equation (26), the L.H.S becomes larger (resp. smaller) than the R.H.S. To show these inequalities, we make use of $0 < r < \rho < 1$ and $0 \le \delta < 1$. The inequality for ν_1 is

$$\frac{(\rho - r\delta)r^2}{(\rho - r)(1-\rho r)(1-r^2\delta)} + \frac{Y_1}{(1-\rho r)(1-\rho r\delta)^2(1-r^2\delta^2)^2(1-r^2\delta)} > 0, \tag{C54}$$

where

$$Y_1 = (1-r^2\delta^2)^2\left[2-\rho r(1+\delta) - r^2\delta + \rho^2 r^4\delta^2\right] - (1-\rho r)(1-\rho r\delta)^2(2-r^2\delta^2-r^4\delta^3).$$

Since $\rho > r > r\delta$, equation (C54) holds if $Y_1 > 0$. Algebraic manipulations show that $Y_1 = (\rho - r\delta)rZ_1$, where

$$\begin{aligned} Z_1 \equiv\ & (2-r^2\delta^2-r^4\delta^3)\left[\delta(1-r^2\delta)(2-r^2\delta^2-\rho r\delta) + (1-\rho r\delta)^2\right] \\ & -(1-r^2\delta^2)^2\left[1+\delta-(\rho+r\delta)r^3\delta^2\right]. \end{aligned}$$

Since $2 - r^2\delta^2 - r^4\delta^3 > 2(1-r^2\delta^2)$, inequality $Z_1 > 0$ follows from

$$2\left[\delta(1-r^2\delta)(2-r^2\delta^2-\rho r\delta) + (1-\rho r\delta)^2\right] - (1-r^2\delta^2)\left[1+\delta-(\rho+r\delta)r^3\delta^2\right] > 0. \tag{C55}$$

To show equation (C55), we break the L.H.S into

$$2\delta(1-r^2\delta)(1-r^2\delta^2) - (1-r^2\delta^2)(\delta - r^4\delta^3) = \delta(1-r^2\delta^2)(1-r^2\delta)^2 > 0$$

and

$$2\left[\delta(1-r^2\delta)(1-\rho r\delta) + (1-\rho r\delta)^2\right] - (1-r^2\delta^2)(1-\rho r^3\delta^2). \tag{C56}$$

Equation (C56) is positive because of the inequalities

$$\begin{aligned} & 2(1-\rho r\delta) > 1 - r^2\delta^2, \\ & \delta(1-r^2\delta) + 1 - \rho r\delta > 1 - \rho r^3\delta^2. \end{aligned}$$

The inequality for ν_∞ is

$$-\frac{(\rho - r\delta)(1-\rho)(1-r)r}{(\rho - r)(1-\rho r)(1-r^2\delta)(1-r\delta)^2} + \frac{Y_\infty}{(1-\rho r)(1-\rho r\delta)^2(1-r^2\delta)(1-r^2\delta^2)^2(1-r\delta)} < 0, \tag{C57}$$

where

$$Y_\infty = (1-\rho)(1-r^2\delta^2)\left[2-\rho r(1+\delta)-r^2\delta+\rho^2r^4\delta^2\right] - (1-\rho r)(1-\rho r\delta)^2(1-r\delta)(2-r^2\delta^2-r^4\delta^3).$$

Algebraic manipulations show that $Y_\infty = -(\rho - r\delta)Z_\infty$, where

$$\begin{aligned} Z_\infty \equiv\ & (2-r^2\delta^2-r^4\delta^3)\left[(1-\rho r\delta)^2(1-r) - (1-\rho)r\delta(1-r^2\delta)(2-r^2\delta^2-\rho r\delta)\right] \\ & +(1-\rho)r(1-r^2\delta^2)^2\left[1+\delta-(\rho+r\delta)r^3\delta^2\right]. \end{aligned}$$

To show that $Z_\infty > 0$, we break it into

$$\begin{aligned} & (2-r^2\delta^2-r^4\delta^3)\left[(1-\rho r\delta)^2(1-r) - (1-\rho)r\delta(1-r^2\delta)(1-\rho r\delta)\right] \\ > \ & (2-r^2\delta^2-r^4\delta^3)(1-\rho r\delta)\left[(1-\rho r\delta)(1-r) - (1-\rho)(1-r^2\delta)\right] \\ = \ & (2-r^2\delta^2-r^4\delta^3)(1-\rho r\delta)(\rho - r)(1-r\delta) > 0 \end{aligned}$$

and

$$\begin{aligned} & (1-\rho)r(1-r^2\delta^2)^2\left[1+\delta-(\rho+r\delta)r^3\delta^2\right] - (2-r^2\delta^2-r^4\delta^3)(1-\rho)r\delta(1-r^2\delta)(1-r^2\delta^2) \\ & \quad = (1-\rho)r(1-r^2\delta^2)\left[(1-r^2\delta^2)\left[1+\delta-(\rho+r\delta)r^3\delta^2\right] - (2-r^2\delta^2-r^4\delta^3)\delta(1-r^2\delta)\right]. \end{aligned}$$

Since $2 - r^2\delta^2 - r^4\delta^3 < 2 - r^2\delta^2 - r^4\delta^4 = (1 - r^2\delta^2)(2 + r^2\delta^2)$, the last square bracket is greater than

$$(1 - r^2\delta^2)\left[1 + \delta - (\rho + r\delta)r^3\delta^2 - \delta(1 - r^2\delta)(2 + r^2\delta^2)\right]$$
$$= (1 - r^2\delta^2)\left[(1 - \delta)(1 - r^4\delta^3) + r^2\delta^2(2 - \delta - \rho r)\right] > 0. \quad \|$$

Proof of Lemma 4. Because stocks are symmetric and returns are uncorrelated, equation (28) takes the form

$$\sum_{n=1}^{N}\left(w_{nt} - \frac{1}{N}\right)^2 \tilde{V}ar_1 \leq TE^2. \tag{C58}$$

The Lagrangian is

$$L \equiv \sum_{n=1}^{N} w_{nt}\tilde{E}_{t-1}(s_{nt}) + \lambda_1\left[TE^2 - \sum_{n=1}^{N}\left(w_{nt} - \frac{1}{N}\right)^2 \tilde{V}ar_1\right] + \lambda_2\left(1 - \sum_{n=1}^{N} w_{nt}\right).$$

The first-order condition with respect to w_{nt} is

$$\tilde{E}_{t-1}(s_{nt}) - 2\lambda_1\left(w_{nt} - \frac{1}{N}\right)\tilde{V}ar_1 - \lambda_2 = 0. \tag{C59}$$

Summing over n and using $\sum_{n=1}^{N} w_{nt} = 1$, we find

$$\lambda_2 = \tilde{E}_{t-1}(s_t). \tag{C60}$$

Substituting λ_2 from equation (C60) into equation (C59), we find

$$w_{nt} = \frac{1}{N} + \frac{\tilde{E}_{t-1}(s_{nt}) - \tilde{E}_{t-1}(s_t)}{2\lambda_1\tilde{V}ar_1}. \tag{C61}$$

Substituting w_{nt} from equation (C61) into equation (C58), which holds as an equality, we find

$$\lambda_1 = \frac{\sqrt{\sum_{n=1}^{N}\left[\tilde{E}_{t-1}(s_{nt}) - \tilde{E}_{t-1}(s_t)\right]^2}}{2TE\sqrt{\tilde{V}ar_1}}. \tag{C62}$$

Substituting w_{nt} from equation (C61) into equation (27), and using equation (C62), we find equation (29). $\|$

Proof of Lemma 5. Equation (30) implies that

$$\sum_{k'=1}^{k}\tilde{F}_{t+k'} = \frac{\tilde{E}_{t+k-1}(\gamma_{t+k}) - \tilde{E}_{t-1}(\gamma_t)}{a\tilde{V}ar_{1\gamma}}. \tag{C63}$$

Replacing s_t by γ_t in equation (C44), substituting $(\tilde{E}_{t-1}(\gamma_t), \tilde{E}_{t+k-1}(\gamma_{t+k}))$ from equation (C44) into equation (C63), and noting that γ_t is i.i.d. with variance $\sigma^2_{\epsilon\gamma}$, we find

$$Cov\left(\gamma_t, \sum_{k'=1}^{k}\tilde{F}_{t+k'}\right) = \frac{\sigma^2_{\epsilon\gamma}\tilde{C}\tilde{D}^{k-1}\tilde{G}}{a\tilde{V}ar_{1\gamma}}. \tag{C64}$$

Since γ_t is i.i.d., $CA^{k'-1}G = 0$ for all $k' \geq 1$. Equation (18) then implies that $\tilde{N}_k = \tilde{C}\tilde{D}^{k-1}\tilde{G}$, and equation (C64) becomes equation (31). $\|$

Proof of Proposition A1. The parameter vectors $\tilde{p}$ that belong to $m(P_0)$ and satisfy $e(\tilde{p}) = 0$ must satisfy equation (C37) for all $k \geq 1$, $\tilde{\mu} = \mu$, and equation (C38). Using $(\alpha_{\tilde{p}}, \delta_{\tilde{p}}) = (\alpha, \delta)$, we can write equation (C37) as

$$\tilde{\rho}^k\tilde{G}_1 - \alpha(\delta - \alpha)^{k-1}\tilde{G}_2 = 0. \tag{C65}$$

Consider first the case where $\tilde{G}_2 \neq 0$. Equation (C65) implies that $\tilde{\rho} = \delta - \alpha$. Multiplying equation (14) from the left by an arbitrary 1×2 vector v, and noting that $v\tilde{A} = (\delta - \alpha)v$, we find

$$v\tilde{G} = \frac{(\delta - \alpha)v\tilde{\Sigma}\tilde{C}' + v\tilde{U}}{\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V}}. \tag{C66}$$

Writing equation (13) as

$$\tilde{\Sigma} = \tilde{A}\tilde{\Sigma}\tilde{A}' - \tilde{G}\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right)' + \tilde{W},$$

multiplying from the left by v, and noting that $v\tilde{A} = (\delta - \alpha)v$, we find

$$v\tilde{\Sigma} = \left[-v\tilde{G}\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right)' + v\tilde{W}\right][I - (\delta - \alpha)\tilde{A}']^{-1}. \tag{C67}$$

Substituting $v\Sigma$ from equation (C67) into equation (C66), we find

$$v\tilde{G}\left(1 + \frac{(\delta - \alpha)\left(\tilde{A}\tilde{\Sigma}\tilde{C}' + \tilde{U}\right)'[I - (\delta - \alpha)\tilde{A}']^{-1}\tilde{C}'}{\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V}}\right) = \frac{(\delta - \alpha)v\tilde{W}[I - (\delta - \alpha)\tilde{A}']^{-1}\tilde{C}' + v\tilde{U}}{\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{V}}. \tag{C68}$$

If v satisfies $v\tilde{G} = 0$, then equation (C66) implies that

$$(\delta - \alpha)v\tilde{\Sigma}\tilde{C}' + v\tilde{U} = 0, \tag{C69}$$

and equation (C68) implies that

$$(\delta - \alpha)v\tilde{W}[I - (\delta - \alpha)\tilde{A}']^{-1}\tilde{C}' + v\tilde{U} = 0. \tag{C70}$$

Equation (C65) for $k = 1$ implies that $\tilde{C}\tilde{G} = 0$. Therefore, equations (C69) and (C70) hold for $v = \tilde{C}$:

$$(\delta - \alpha)\tilde{C}\tilde{\Sigma}\tilde{C}' + \tilde{C}\tilde{U} = 0, \tag{C71}$$

$$(\delta - \alpha)\tilde{C}\tilde{W}[I - (\delta - \alpha)\tilde{A}']^{-1}\tilde{C}' + \tilde{C}\tilde{U} = 0. \tag{C72}$$

Equations (C38) and (C71) imply that

$$-\frac{\tilde{C}\tilde{U}}{\delta - \alpha} + \tilde{V} = \sigma_\omega^2. \tag{C73}$$

Substituting for $\tilde{\rho} = \delta - a$ and $(\tilde{A}, \tilde{C}, \tilde{V}, \tilde{W}, \tilde{U})$, we find that the solution $(\tilde{\sigma}_\eta^2, \tilde{\sigma}_\omega^2)$ to the system of equations (C72) and (C73) is as in $\tilde{p}_1$.

Consider next the case where $\tilde{G}_2 = 0$. Equation (C70) holds for $v \equiv (0, 1)$ because $v\tilde{A} = (\delta - \alpha)v$ and $v\tilde{G} = 0$. Solving this equation, we find $\tilde{\sigma}_\omega^2 = 0$. The unique solution of equation (13) is $\tilde{\Sigma} = 0$, and equation (C38) implies that $\tilde{\sigma}_\eta^2 = \sigma_\omega^2$. Equation (14) implies that $G_1 = 1$, and $CG = 0$ implies that $\tilde{\rho} = 0$. Therefore, $\tilde{p} = \tilde{p}_2$.

Showing that $\tilde{p}_2$ satisfies equation (C37) for all $k \geq 1$, $\tilde{\mu} = \mu$, and equation (C38) is obvious. Showing the same for $\tilde{p}_1$ follows by retracing the previous steps and noting that the term in brackets in equation (C68) is non-zero. ‖

Proof of Proposition A2. We first show a counterpart of Proposition 5, namely that when α converges to zero, the set

$$\left\{\left(\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) : \left(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu}\right) \in m(P_{\underline{\rho}})\right\}$$

converges to the point $(z, \underline{\rho}, \sigma_\omega^2, \mu)$, where

$$z \equiv \frac{(1 + \underline{\rho})^2}{\underline{\rho}(1 - \underline{\rho}\delta)}. \tag{C74}$$

To show this result, we follow the same steps as in the proof of Proposition 5. The limits of $(\tilde{\sigma}_\omega^2, \tilde{\mu})$ are straightforward. The limits (ℓ_s, ℓ_ρ) of $(\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}, \tilde{\rho})$ minimize the function

$$F(\ell_s, \ell_\rho) \equiv \sum_{k=1}^{\infty} \left[\ell_\rho^k \frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1} \right]^2.$$

Treating F as a function of $\left(\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}, \ell_\rho\right)$, the first-order condition with respect to the first argument is

$$\sum_{k=1}^{\infty} \ell_\rho^k \left[\ell_\rho^k \frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1} \right] = 0. \tag{C75}$$

Since $\ell_\rho \geq \underline{\rho} > \delta$, the function

$$\hat{f}_k \equiv \ell_\rho^k \left[\ell_\rho^k \frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho} - \delta^{k-1} \right]$$

is negative in an interval $k \in \{1, \ldots, k_0 - 1\}$ and becomes positive for $k \in \{k_0, \ldots, \infty\}$. Using this fact and equation (C75), we find

$$\sum_{k_0}^{\infty} \hat{f}_k = -\sum_{k=1}^{k_0-1} \hat{f}_k \Rightarrow \sum_{k_0}^{\infty} \frac{k}{k_0} \hat{f}_k > -\sum_{k=1}^{k_0-1} \frac{k}{k_0} \hat{f}_k \Rightarrow \sum_{k=1}^{\infty} k \hat{f}_k > 0. \tag{C76}$$

The left-hand side of the last inequality in equation (C76) has the same sign as the derivative of F with respect to the second argument. Since that derivative is positive, F is minimized for $\ell_\rho = \underline{\rho}$. Substituting $\ell_\rho = \underline{\rho}$ into equation (C75), we find $\ell_s = z$.

We complete the proof of Proposition A2 by following the same steps as in the proof of Proposition 6. The function g_k is defined as in that proof, while the function f_k is defined as

$$f_k \equiv \underline{\rho}^k \frac{z(1-\underline{\rho})}{1+\underline{\rho}} - \delta^{k-1}.$$

Using equation (C74) to substitute z, we find

$$f_1 = g_1 = \frac{\underline{\rho}(\delta - \underline{\rho})}{1 - \underline{\rho}\delta} < 0,$$

$$g_\infty = \frac{z\underline{\rho}}{1+\underline{\rho}} - \frac{1}{1-\delta} = \frac{\underline{\rho} - \delta}{(1-\underline{\rho}\delta)(1-\delta)} > 0.$$

The remainder of the proof is identical to the last step in the proof of Proposition 6. ||

Proof of Proposition A3. We first show a counterpart of Proposition 7, namely that when α and σ_η^2 converge to zero, holding ν constant, the set

$$\left\{ \left(\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu} \right) : \left(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2, \tilde{\mu} \right) \in m(P_0) \right\}$$

converges to the point $(z, r, \sigma_\omega^2, \mu)$, where

$$z \equiv \frac{\nu\rho(1-\rho)(1+r)^2}{r(1+\rho)(1-\rho r)} + \frac{(1+r)^2}{r(1-r\delta)} \tag{C77}$$

and r solves

$$\frac{\nu(1-\rho)\rho(\rho - r)}{(1+\rho)(1-\rho r)^2} = \frac{r - \delta}{(1-r\delta)^2}. \tag{C78}$$

To show this result, we follow the same steps as in the proof of Proposition 7. The limit of $\tilde{\mu}$ is straightforward. If the limit ℓ_ρ of $\tilde{\rho}$ is positive, then

$$\lim_{v\to\ell_v}\frac{\sigma_s^2(\tilde{p})\sum_{k=1}^{\infty}e_k^2}{\alpha^2}=\sigma_\omega^2\sum_{k=1}^{\infty}\left[\ell_\rho^k\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}-\delta^{k-1}-\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2\equiv\sigma_\omega^2 H(\ell_s,\ell_\rho),$$

where ℓ_s denotes the limit of $\frac{\tilde{\sigma}_\eta^2}{\alpha\tilde{\sigma}_\omega^2}$. If $\ell_\rho=0$, then

$$\lim_{v\to\ell_v}\frac{\sigma_s^2(\tilde{p})\sum_{k=1}^{\infty}e_k^2}{\alpha^2}=\sigma_\omega^2\left(\left[\ell_\phi-\ell_G-\rho\frac{\nu(1-\rho)}{1+\rho}\right]^2+\sum_{k=2}^{\infty}\left[\delta^{k-1}\ell_G+\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2\right)\equiv\sigma_\omega^2 H_0(\ell_\phi),$$

where (ℓ_ϕ,ℓ_G) denote the limits of $\left(\frac{\tilde{\rho}\tilde{G}_1}{\alpha_n},\tilde{G}_2\right)$. Because ℓ_G can differ from 1, the cases $\ell_\rho>0$ and $\ell_\rho=0$ cannot be nested. To show that $\ell_\rho>0$, we compare instead the minima of H and H_0. Since $\ell_G\geq 0$,

$$\min_{\ell_\phi}H_0(\ell_\phi)\geq\sum_{k=2}^{\infty}\left[\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2.$$

Since, in addition, equation (A1) implies that

$$\sum_{k=2}^{\infty}\left[\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2>\min_{\ell_s}H(\ell_s,\delta),$$

and equation (A2) implies that

$$\sum_{k=2}^{\infty}\left[\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]^2>H(\nu,\rho),$$

the minimum of H_0 exceeds that of H, and therefore, $\ell_\rho>0$. Since $\ell_\rho>0$, the limit of $\tilde{\sigma}_\omega^2$ is σ_ω^2.

To minimize H, we treat it as a function of $\left(\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho},\ell_\rho\right)$. The first-order condition with respect to the first argument is

$$\sum_{k=1}^{\infty}\ell_\rho^k\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}\ell_\rho^k-\delta^{k-1}-\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]=0, \tag{C79}$$

and the derivative with respect to the second argument is

$$\sum_{k=1}^{\infty}k\ell_\rho^{k-1}\left[\frac{\ell_s(1-\ell_\rho)}{1+\ell_\rho}\ell_\rho^k-\delta^{k-1}-\rho^k\frac{\nu(1-\rho)}{1+\rho}\right]. \tag{C80}$$

Computing the infinite sums, we can write equation (C79) as

$$\frac{\ell_s\ell_\rho}{(1+\ell_\rho)^2}-\frac{1}{1-\ell_\rho\delta}-\frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho\ell_\rho)}=0, \tag{C81}$$

and equation (C80) as

$$\frac{\ell_s\ell_\rho}{(1+\ell_\rho)^2(1-\ell_\rho^2)}-\frac{1}{(1-\ell_\rho\delta)^2}-\frac{\nu\rho(1-\rho)}{(1+\rho)(1-\rho\ell_\rho)^2}. \tag{C82}$$

Substituting ℓ_s from equation (C81), we can write equation (C82) as

$$-\frac{\nu\rho(1-\rho)\ell_\rho(\rho-\ell_\rho)}{(1+\rho)(1-\rho\ell_\rho)^2(1-\ell_\rho^2)}+\frac{\ell_\rho(\ell_\rho-\delta)}{(1-\ell_\rho\delta)^2(1-\ell_\rho^2)}. \tag{C83}$$

Equation (C49) is negative for $\ell_\rho\in[0,\min\{\delta,\rho\}]$ and positive for $\ell_\rho\in[\max\{\delta,\rho\},\overline{\rho}]$. Therefore, it is equal to zero for a value between δ and ρ. Comparing equations (26) and (C49), we find that this value is equal to r. Substituting $\ell_\rho=r$ into equation (C47), we find $\ell_s=z$.

We complete the proof of Proposition A1 by following the same steps as in the proof of Proposition 8. The function g_k is defined as in that proof, while the function f_k is defined as

$$f_k \equiv r^k \frac{z(1-r)}{1+r} - \delta^{k-1} - \rho^k \frac{\nu(1-\rho)}{1+\rho}.$$

Using equation (C77) to substitute z, and then equation (C78) to substitute ν, we find

$$f_1 = g_1 = \frac{r^2(\delta - r)(\rho - \delta)}{(1-r\delta)^2}.$$

This is negative because equation (C78) implies that r is between δ and ρ. We similarly find

$$g_\infty = \frac{zr}{1+r} - \frac{1}{1-\delta} - \frac{\nu\rho}{1+\rho} = \frac{(1-r)(\delta-r)(\rho-\delta)}{(1-r\delta)^2(1-\rho)(1-\delta)} < 0.$$

The remainder of the proof follows by the same argument as in the proof of Proposition 8. ‖

Proof of Proposition B1. The parameter vectors $\tilde{p}$ that belong to $m(P_0)$ and satisfy $e(\tilde{p}) = 0$ must satisfy (i) $\tilde{C}\tilde{A}^{k-1}\tilde{G} = CA^{k-1}G$ for all $k \geq 1$, (ii) $\tilde{\mu} = \mu$, and (iii) $\sigma_s^2(\tilde{p}) = \sigma_s^2$. Since $\sigma_\eta^2 = 0$ and

$$\tilde{A}^k = \begin{bmatrix} \tilde{\rho}^k & 0 & -(1-\tilde{\rho})\alpha \frac{\tilde{\rho}^k - (\delta-\alpha)^k}{\tilde{\rho} - (\delta-\alpha)} \\ 0 & (\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^k & 0 \\ 0 & 0 & (\delta-\alpha)^k \end{bmatrix},$$

we can write Condition (i) as

$$\tilde{\rho}^{k+1}\tilde{G}_1 - \alpha_{\tilde{\rho}}(\delta_{\tilde{\rho}} - \alpha_{\tilde{\rho}})^k \tilde{G}_2 - (1-\tilde{\rho})\alpha \frac{\tilde{\rho}^{k+1} - (\delta-\alpha)^{k+1}}{\tilde{\rho} - (\delta-\alpha)} \tilde{G}_3 = 0, \tag{C84}$$

and Condition (iii) as equation (C38).

Consider first the linear specification $(\alpha_\rho, \delta_\rho) = (\alpha\rho, \delta\rho)$. If $\tilde{G}_3 \neq 0$, then $\tilde{\rho} \neq \delta - \alpha$; otherwise, equation (C84) would include a term linear in k. Since $\delta - \alpha \notin \{\tilde{\rho}, (\delta-\alpha)\tilde{\rho}\}$ (the latter because $\tilde{\rho} < 1$), equation (C84) implies that $\tilde{G}_3 = 0$, which is a contradiction. Therefore, $\tilde{G}_3 = 0$, and equation (C84) becomes equation (C37). Using the same argument as in the proof of Proposition 4, we find $\tilde{\rho} = 0$. Since $v\tilde{A} = (\delta - \alpha)v$ and $v\tilde{G} = 0$ for $v \equiv (0, 0, 1)$, v satisfies equation (C70). Solving equation (C70), we find $\tilde{\sigma}_\zeta^2 = 0$. The unique solution of equation (13) is $\tilde{\Sigma} = 0$, and equation (C38) implies that $\tilde{\sigma}_\omega^2 = \sigma_\omega^2$. Therefore, $\tilde{p}$ is as in the proposition.

Consider next the constant specification $(\alpha_\rho, \delta_\rho) = (\alpha, \delta)$. If $\tilde{G}_3 = 0$, then $v \equiv (0, 0, 1)$ satisfies equation (C70), and therefore, $\tilde{\sigma}_\zeta^2 = 0$ and $\tilde{\Sigma} = 0$. Equation (14) then implies that $\tilde{G}_1 = 0$, and equation (C84) implies that $\tilde{G}_2 = 0$. Since $v\tilde{A} = (\delta - \alpha)v$ and $v\tilde{G} = 0$ for $v \equiv (0, 1, 0)$, v satisfies equation (C70). Solving equation (C70), we find $\tilde{\sigma}_\omega^2 = 0$. Equations $\tilde{\sigma}_\zeta^2 = \tilde{\sigma}_\omega^2 = 0$ and $\tilde{\Sigma} = 0$ contradict equation (C38) since $\sigma_\omega^2 = \sigma_\epsilon^2 > 0$. Therefore, $\tilde{G}_3 \neq 0$, which implies $\tilde{\rho} \neq \delta - \alpha$. Identifying terms in $\tilde{\rho}^k$ and $(\delta - \alpha)^k$ in equation (C84), we find

$$G = G_3 \left(\frac{(1-\tilde{\rho})\alpha}{\tilde{\rho} - (\delta-\alpha)}, \frac{(1-\tilde{\rho})(\delta-\alpha)}{\tilde{\rho} - (\delta-\alpha)}, 1 \right)'.$$

We set

$$v_1 \equiv \left(0, -1, \frac{(1-\tilde{\rho})(\delta-\alpha)}{\rho - (\delta-\alpha)} \right),$$

$$v_2 \equiv \left(1, 0, -\frac{(1-\tilde{\rho})\alpha}{\rho - (\delta-\alpha)} \right),$$

and $(\lambda_1, \lambda_2) \equiv (\delta - \alpha, \tilde{\rho})$. Since $v_i\tilde{A} = \lambda_i v_i$ and $v_i\tilde{G} = 0$ for $i = 1, 2$, equation (C69) generalizes to

$$\lambda_i v_i \tilde{\Sigma}\tilde{C}' + v_i \tilde{U} = 0, \tag{C85}$$

and equation (C70) to

$$\lambda_i v_i \tilde{W} \left(I - \lambda_i \tilde{A}' \right)^{-1} \tilde{C}' + v_i \tilde{U} = 0. \tag{C86}$$

Noting that $\tilde{C} = \sum_{i=1,2} \mu_i v_i$, where $(\mu_1, \mu_2) \equiv (\alpha, \tilde{\rho})$, and using equation (C85), we find

$$\tilde{C} \tilde{\Sigma} \tilde{C}' = - \sum_{i=1,2} \frac{\mu_i}{\lambda_i} v_i \tilde{U}. \tag{C87}$$

Equations (C38) and (C87) imply that

$$- \sum_{i=1,2} \frac{\mu_i}{\lambda_i} v_i \tilde{U} + \tilde{V} = \sigma_\omega^2. \tag{C88}$$

Substituting for $(\tilde{A}, \tilde{C}, \tilde{V}, \tilde{W}, \tilde{U})$, we find that the solution $(\tilde{\sigma}_\eta^2, \tilde{\rho}, \tilde{\sigma}_\omega^2)$ to the system of equations (C86) and (C88) is as in the proposition. Showing that the parameter vectors in the proposition satisfy Conditions (i)–(iii) follows by the same argument as for $\tilde{p}_1$ in Proposition A1. ‖

Acknowledgements. This paper was circulated previously under the title "The Gambler's and Hot-Hand Fallacies in a Dynamic-Inference Model". We thank Anat Admati, Nick Barberis, Bruno Biais, Markus Brunnermeier, Eddie Dekel, Peter DeMarzo, Xavier Gabaix, Vassilis Hajivassiliou, Michael Jansson, Leonid Kogan, Jon Lewellen, Andrew Patton, Anna Pavlova, Francisco Penaranda, Paul Pfleiderer, Jean Tirole, the referees, and seminar participants at Bocconi, LSE, MIT, UCL, and Yale for useful comments. Jeff Holman provided excellent research assistance. Rabin acknowledges financial support from the NSF.

REFERENCES

ANDERSON, B. and MOORE, J. (1979), *Optimal Filtering* (New York: Prentice Hall).

ASPAROUHOVA, E., HERTZEL, M. and LEMMON, M. (2008), "Inference from Streaks in Random Outcomes: Experimental Evidence on Beliefs in Regime-Shifting and the Law of Small Numbers" (Working Paper, University of Utah).

BALAKRISHNAN, A. (1987), *Kalman Filtering Theory* (New York: Optimization Software Inc., Publications Division).

BAR-HILLEL, M. and WAGENAAR, W. (1991), "Perceptions of Randomness", *Advances of Applied Mathematics,* **12**, 428–454.

BARBERIS, N., SHLEIFER, A. and VISHNY, R. (1998), "A Model of Investor Sentiment", *Journal of Financial Economics,* **49**, 307–343.

BERK, J. and GREEN, R. (2004), "Mutual Fund Flows and Performance in Rational Markets", *Journal of Political Economy,* **112**, 1269–1295.

BILLINGSLEY, P. (1986), *Probability and Measure* (New York: John Wiley).

BLOOMFIELD, R. and HALES, J. (2002), "Predicting the Next Step of a Random Walk: Experimental Evidence of Regime-Shifting Beliefs", *Journal of Financial Economics,* **65**, 397–414.

BOGLE, J. (2005), "The Mutual Fund Industry 60 Years Later: For Better or Worse?" *Financial Analysts Journal,* **61**, 15–24.

BUDESCU, D. and RAPOPORT, A. (1994), "Subjective Randomization in One- and Two-Person Games", *Journal of Behavioral Decision Making,* **7**, 261–278.

CAMERER, C. (1989), "Does the Basketball Market Believe in the 'Hot Hand'?" *American Economic Review,* **79**, 1257–1261.

CARHART, M. (1997), "On Persistence in Mutual Fund Performance", *Journal of Finance,* **52**, 57–82.

CHEVALIER, J. and ELLISON, G. (1997), "Risk Taking by Mutual Funds as a Response to Incentives", *Journal of Political Economy,* **105**, 1167–1200.

CLOTFELTER, C. and COOK, P. (1993), "The 'Gambler's Fallacy in Lottery Play", *Management Science,* **39**, 1521–1525.

CROSON, R. and SUNDALI, J. (2005), "The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos", *Journal of Risk and Uncertainty,* **30**, 195–209.

DING, B., GETMANSKY, M., LIANG, B. and WERMERS, R. (2008), "Investor Flows and Share Restrictions in the Hedge Fund Industry" (Working Paper, University of Massachusetts at Amherst).

EDWARDS, W. (1961), "Probability Learning in 1000 Trials", *Journal of Experimental Psychology,* **62**, 385–394.

FRENCH, K. (2008), "The Cost of Active Investing", *Journal of Finance,* **63**, 1537–1573.

FUNG, W., HSIEH, D., NAIK, N. and RAMADORAI, T. (2008), "Hedge Funds: Performance, Risk, and Capital Formation", *Journal of Finance,* **58**, 1777–1803.

GILOVICH, T., VALLONE, R. and TVERSKY, A. (1985), "The Hot Hand in Basketball: On the Misperception of Random Sequences", *Cognitive Psychology,* **17**, 295–314.

GOLD, E. and HESTER, G. (2008), "The Gambler's Fallacy and the Coin's Memory", in J. Krueger (ed.) *Rationality and Social Responsibility: Essays in Honor of Robyn Mason Dawes* (New York: Psychology Press).

GREENE, W. (2008), *Econometric Analysis*, 6th edn (New York: Prentice Hall).

GRUBER, M. (1996), "Another Puzzle: The Growth in Actively Managed Mutual Funds", *Journal of Finance,* **51**, 783–810.

McLEISH, D. (1975), "A Maximal Inequality and Dependent Strong Laws", *The Annals of Probability,* **3**, 829–839.

METZGER, M. (1984), "Biases in Betting: An Application of Laboratory Findings", *Psychological Reports,* **56**, 883–888.

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

RAPOPORT, A. and BUDESCU, D. (1992), "Generation of Random Binary Series in Strictly Competitive Games", *Journal of Experimental Psychology: General,* **121**, 352–364.

RAPOPORT, A. and BUDESCU, D. (1997), "Randomization in Individual Choice Behavior", *Psychological Review,* **104**, 603–617.

ROLL, R. (1992), "A Mean/Variance Analysis of Tracking Error", *Journal of Portfolio Management,* **18**, 15–24.

SIRRI, E. and TUFANO, P. (1998), "Costly Search and Mutual Fund Flows", *Journal of Finance,* **53**, 1655–1694.

TERRELL, D. (1994), "A Test of the Gambler's Fallacy—Evidence from Pari-Mutuel Games", *Journal of Risk and Uncertainty,* **8**, 309–317.

TERRELL, D. (1998), "Biases in Assessments of Probabilities: New Evidence from Greyhound Races", *Journal of Risk and Uncertainty,* **17**, 151–166.

TERRELL, D. and FARMER A. (1996), "Optimal Betting and Efficiency in Parimutuel Betting Markets with Information Costs", *The Economic Journal,* **106**, 846–868.

TVERSKY, A. and GILOVICH T. (1989*a*), "The Cold Facts about the 'Hot Hand' in Basketball", *Chance,* **2**, 16–21.

TVERSKY, A. and GILOVICH T. (1989*b*), "The Hot Hand: Statistical Reality or Cognitive Illusion?" *Chance,* **2**, 31–34.