LEDa WORKING PAPERS

Navigating Speculative Bubbles: A Higher-Moment Approach to Optimal Asset Allocation

Gilles de Truchis, Elena-Ivona Dumitrescu, Sébastien Fries, Arthur Thomas

WP 2026-05

Septembre 2026

Navigating Speculative Bubbles: A Higher-Moment Approach to Optimal Asset Allocation

Gilles de Truchis[1,**], Elena-Ivona Dumitrescu[2], Sébastien Fries[3], Arthur Thomas[4,*]

Abstract

We study portfolio allocation in the presence of speculative bubble dynamics using a mixed causal-noncausal (MAR) framework with α-stable innovations. This approach captures forward-looking behavior, heavy tails, and asymmetric return distributions. We derive closed-form expressions for the first four conditional moments of returns and embed them into a portfolio optimization problem with endogenous investment horizons. The resulting framework generates dynamic allocation rules, multiple locally optimal strategies, and sparse rebalancing patterns. Both simulation and empirical results on commodity markets show that the proposed portfolio strategies outperform standard benchmarks, particularly during bubble periods. These gains are primarily driven by the ability of the allocation rule to adjust exposure as the balance between bubble continuation and crash risk evolves. A sensitivity analysis further reveals that skewness and kurtosis mainly improve the robustness, risk-adjusted performance and implementability of portfolio strategies in speculative markets.

Keywords: noncausal process, α-stable, asset allocation, utility function, out-of-sample performance
JEL: C51, C22, G12

1. Introduction

Financial markets frequently exhibit episodes of speculative excess during which asset prices deviate markedly from fundamental values. These episodes, often referred to as bubbles, are characterized by prolonged phases of rapid price appreciation followed by abrupt and sizable corrections. Empirical evidence across a wide range of asset classes, particularly commodities,

*Corresponding author
**We are grateful to Jean-Michel Zakoian, Christophe Pérignon, Serge Darolles, Gaëlle le Fol. We also thank the seminar participants at Dauphine University, University of Paris Nanterre, as well as the participants at the 8th Quantitative Finance and financial Econometrics (Aix-Marseille University), 16th International Conference on Computational and Financial Econometrics (King's College London) and the 21st Conference "Développements Récents de l'Econométrie Appliquée à la Finance" (University of Paris Nanterre) for helpful comments and discussions.
Email addresses: gilles.detruchis@univ-orleans.fr (Gilles de Truchis), elena.dumitrescu@parisnanterre.fr (Elena-Ivona Dumitrescu), sebastien.fries@gmail.com (Sébastien Fries), arthur.thomas@dauphine.psl.eu (Arthur Thomas)

[1]LEO, University of Orléans, Rue de Blois BP 6739, 45067 Orléans Cedex, France.
[2]EconomiX-CNRS, University of Paris Nanterre, 200 Avenue de la République, 92000 Nanterre, France.
[3]Department of Econometrics and Data Science, Vrije Universiteit Amsterdam, Amsterdam, Netherlands
[4]Université Paris-Dauphine, Université PSL, LEDA, CNRS, IRD, [CGEMP], 75016 PARIS, FRANCE

energy markets, and cryptocurrencies, documents such patterns, which give rise to return distributions that are highly asymmetric, nonlinear, and heavy-tailed. These features pose a fundamental challenge for portfolio allocation, as standard asset pricing and portfolio selection frameworks are ill-equipped to accommodate state-dependent risks associated with bubble dynamics.

Traditional portfolio choice models, rooted in mean–variance analysis [Markowitz, 1952] or intertemporal expected utility maximization [Merton, 1969, 1971], are not designed to accommodate the asymmetric and rapidly changing risk profile associated with bubble episodes. Even extensions that incorporate higher moments typically rely on static distributions or exogenous regime changes, leaving little room for risks that emerge endogenously from the asset price dynamics themselves.

To address this issue, we develop a portfolio allocation framework that explicitly accounts for the nonlinear and forward-looking nature of speculative bubbles. Our approach builds on mixed causal–noncausal autoregressive (MAR) processes with α-stable innovations. The MAR specification combines backward-looking and forward-looking dynamics and has proved particularly well suited for describing speculative episodes. Combined with α-stable shocks, it generates asymmetric and heavy-tailed return distributions whose properties vary with the current level of the asset price.

Leveraging recent results on conditional moments of noncausal processes, we derive closed-form expressions for the first four conditional moments of returns in a MAR(1,1) setting. These moments are then incorporated into a portfolio optimization problem in which investors jointly choose the portfolio weight allocated to the speculative asset and the horizon over which the position is held. The endogenous timing dimension is particularly relevant in bubble environments, where expected gains and crash risk evolve rapidly. This feature distinguishes our approach from standard portfolio frameworks, which typically rely on a fixed investment horizon.

The framework yields several implications. First, optimal allocations vary substantially across market states. Investors increase exposure during intermediate stages of a bubble but reduce or reverse their positions as crash risk rises. Second, the optimization problem is non-convex and may generate multiple locally optimal strategies that differ in both exposure and timing. Higher-order moments provide an additional refinement by improving the assessment of asymmetric and tail risks.

We evaluate these mechanisms through Monte Carlo simulations and an empirical application to commodity markets. The simulation analysis shows that parameter estimation uncertainty has only a limited impact on the mapping from the state variable to optimal strategies. Out-of-sample, the proposed bubble portfolio (BP) strategies outperform equally weighted and mean–variance benchmarks, particularly during bubble episodes. A sensitivity analysis further shows that most of these gains originate from the state-dependent allocation mechanism, while higher-order moments mainly improve risk-adjusted performance and the stability of portfolio decisions.

Our paper contributes to several strands of the literature. First, it complements the literature on bubble dynamics and explosive behavior in asset prices [Phillips et al., 2011, 2015] by providing a structural link between price levels and conditional return distributions. Second, it contributes to the literature on higher-moment portfolio choice [Jondeau and Rockinger, 2006, 2012] by showing how skewness and kurtosis arise naturally from speculative price dynamics and affect portfolio decisions. Third, it extends recent work on forward-looking noncausal time series models [Gourieroux and Zakoian, 2017, Fries and Zakoian, 2019] to the domain of portfolio optimization.

The remainder of the paper is organized as follows. Section 2 presents the portfolio allocation framework and characterizes the conditional moments of returns. Section 3 reports Monte Carlo simulations and examines the impact of parameter estimation. Section 4 evaluates the economic value of the proposed strategies. Section 5 provides an empirical application to commodity markets. Section 6 concludes. Appendix A and Appendix B provide the theoretical proofs and detail the utility specifications. A Web Appendix reports additional Monte Carlo and empirical results.

2. Bubble-riding allocation problem

This section introduces the portfolio allocation framework in the presence of a speculative asset exhibiting bubble dynamics. Our approach combines a structural model of bubble behavior with a state-dependent portfolio choice problem based on higher-order conditional moments. We first define the investor's optimization problem and the associated wealth dynamics. We then show how the allocation problem can be approximated using a Taylor expansion of expected utility, which highlights the role of higher-order moments.

2.1. *Optimal portfolio allocation*

We consider an investor allocating his wealth between a speculative asset with price X_t and a bubble-free asset with price S_t. The initial wealth is innocuous to the optimization problem and arbitrarily set to one. Consistent with the discussion in the introduction, the key feature of this framework is that the conditional distribution of returns depends on the current state of the speculative asset, inducing state-dependent higher-order moments that drive portfolio allocation.

The investor has an investment horizon H: at time t, he chooses (i) the portfolio weight ω invested in the speculative asset and (ii) an intermediate horizon $h \leq H$ at which the position in the speculative asset is liquidated and reinvested into the safe asset until the terminal date $t + H$. Short selling is allowed. The resulting terminal wealth is given by

$$W_{t+H} = \frac{S_{t+H}}{S_{t+h}} \left(\omega \frac{X_{t+h}}{X_t} + (1-\omega) \frac{S_{t+h}}{S_t} \right). \tag{1}$$

This formulation introduces an endogenous investment horizon h, which allows the investor to optimally time exit from the speculative asset. This dimension is central in environments characterized by bubble dynamics, where expected returns and downside risks vary sharply with the current state of the asset.

The speculative asset is assumed to follow a MAR(1,1) process with α-stable innovations, capturing locally explosive behavior, heavy tails, and asymmetric risk profiles. The safe asset follows a geometric Brownian motion with drift v and volatility ς. The two processes are assumed independent, which provides a natural hedging benchmark.[5]

The investor maximizes the expected utility of terminal wealth:

$$\max_{(\omega,h)} \mathbb{E}[U(W_{t+H}) \mid X_t, S_t].$$

Therefore, the portfolio allocation problem can be interpreted as mapping the current state variable into optimal decisions through its impact on conditional higher-order moments. We consider

[5]In this framework, the independence hypothesis does not appear as a strong assumption. As we focus on investment during periods where an asset price exhibits a bubble behaviour, its dynamics cannot be correlated over this time-interval with that of a safe(r) asset. In practice it is reasonable to think of the second *asset* as a well-diversified portfolio whose constituents do not exhibit any bubble behaviour. This makes the second asset an attractive hedge against the risk of bubble collapse in the first asset.

both CRRA and LIRA preferences. While CRRA utility is standard in portfolio choice, it may generate degenerate demands in the presence of infrequent extreme risks [Gourieroux and Monfort, 2004]. LIRA preferences address this limitation by ensuring bounded marginal utility for large losses (see Appendix B).

Following Jondeau and Rockinger [2006, 2012], we approximate the expected utility using a fourth-order Taylor expansion around

$$\overline{W} := \mathbb{E}[W_{t+H} \mid X_t, S_t].$$

This yields

$$\mathbb{E}[U(W_{t+H}) \mid X_t, S_t] \approx \sum_{p=0}^{4} \frac{U^{(p)}(\overline{W})}{p!} \mathbb{E}\left[(W_{t+H} - \overline{W})^p \mid X_t, S_t\right]. \tag{2}$$

The optimization problem therefore depends on the first four conditional moments of terminal wealth. Denoting raw moments by

$$M_p := \mathbb{E}[W_{t+H}^p \mid X_t, S_t], \quad p = 1, \ldots, 4,$$

the corresponding central moments entering (2) can be obtained by standard transformations. The Taylor expansion highlights that portfolio choice depends not only on mean and variance, but also on conditional skewness and kurtosis. In particular, skewness enters positively while kurtosis enters negatively, in line with standard preference ordering [Scott and Horvath, 1980].

In the present setting, these higher-order moments are driven by the nonlinear and noncausal dynamics of X_t, which generate asymmetric return distributions and state-dependent crash risk. Accurately characterizing these moments is therefore essential for determining optimal portfolio strategies.

2.2. Conditional moments of MAR(1,1) α-stable processes

This subsection characterizes the conditional moments of the speculative asset price, which are key inputs of the portfolio optimization problem described in Section 2.1. In our framework, optimal allocation depends on the conditional distribution of future returns, and therefore on the first four conditional moments of X_{t+h}. To obtain them, we model the speculative asset price (X_t) as a mixed causal-noncausal autoregressive process with α-stable innovations. This class of models has been shown to capture salient features of bubble dynamics, including locally explosive behavior, heavy tails, and asymmetric risk profiles [Gourieroux and Zakoian, 2017, Fries and Zakoian, 2019, Hecq and Voisin, 2021].

Let (X_t) be the α-stable solution of the MAR(1,1) process

$$\left(1 - \varphi^{\circ} L^{-1}\right)\left(1 - \varphi^{\bullet} L\right) X_t = \varepsilon_t,$$

where $\varepsilon_t \overset{i.i.d.}{\sim} S(\alpha, \beta, \sigma, \mu)$ with $\alpha \in (0, 2)$, $\alpha \neq 1$, $\beta \in [-1, 1]$, and $\sigma > 0$, and L (resp. L^{-1}) denotes the lag (resp. lead) operator. The process is well defined and strictly stationary for $|\varphi^{\circ}| < 1$ and $|\varphi^{\bullet}| < 1$. It admits a two-sided $MA(\infty)$ representation

$$X_t = \sum_{k \in \mathbb{Z}} a_k \varepsilon_{t+k}.$$

Despite the fact that α-stable distributions may have infinite variance, conditional moments of X_{t+h} given X_t can still exist under suitable conditions. Specifically, results from Fries [2022] ensure the existence of moments up to order p provided that the coefficients (a_k) satisfy appropriate summability conditions. An important implication is that the more noncausal (i.e., forward-looking) the process, the stronger the existence results. In particular, when the MA coefficients decay sufficiently fast, all conditional moments of order strictly below $2\alpha + 1$ exist, which is crucial for implementing higher-moment portfolio allocation. When they exist, the conditional moments $\mathbb{E}[X_{t+h}^p \mid X_t]$, for $p \leq 4$, admit closed-form expressions as functions of four quantities $(\sigma_1^{\alpha}, \beta_1, \kappa_p, \lambda_p)$ and a family of functions $\mathcal{H}$ derived in Fries [2022]. The following proposition provides their expressions in the MAR(1,1) case.

Proposition 1. *For $\alpha \neq 1$, the conditional moments $\mathbb{E}[X_{t+h}^p \mid X_t]$, $p \leq 4$, exist whenever the summability condition of Fries [2022, Th. 2.1–2.2] holds. In the MAR(1,1) case, this condition is satisfied for $\alpha > 1/2$ when $p = 2$, $\alpha > 1$ when $p = 3$, and $\alpha > 3/2$ when $p = 4$. When they exist, these moments can be written as functions of $(\sigma_1^{\alpha}, \beta_1, \kappa_p, \lambda_p)$ and $\mathcal{H}$. Closed-form expressions for the corresponding constants are reported in Appendix A.2.*

This representation is particularly useful for the BP allocation, as it provides tractable expressions for the conditional mean, variance, skewness, and kurtosis of returns.

The structure of the conditional moments becomes especially transparent in the tails of the distribution. As shown by Fries [2022], for large values of $|X_t|$, and provided $|\beta_1| \neq 1$, the conditional moments satisfy

$$x^{-p}\mathbb{E}[X_{t+h}^p \mid X_t = x] \underset{x \to \pm\infty}{\to} \frac{\kappa_p \pm \lambda_p}{1 \pm \beta_1},$$

highlighting the asymmetry between positive and negative bubble states. In fact, a key implication of the MAR specification is that the conditional distribution of returns becomes increasingly asymmetric as the process enters a bubble regime. Notably, recent results by Gouriéroux et al. [2025] show that, for large values of X_t, the conditional distribution of the realized one-period gross return $r_{t+1} := X_{t+1}/X_t$ converges to a two-point distribution, reflecting i) a continuation regime, in which the bubble keeps expanding, and ii) a crash regime, characterized by abrupt mean reversion. As $x \to \infty$ and for $p < \alpha$, the corresponding conditional moments can be approximated as

$$\mathbb{E}\left[\left(\frac{X_{t+1}}{X_t}\right)^p \middle| X_t = x, r_t = r\right] \approx (1-(\varphi^\circ)^\alpha)(\varphi^\bullet)^p + (\varphi^\circ)^\alpha \left(\varphi^\bullet + \frac{1}{\varphi^\circ}\left(1 - \frac{\varphi^\bullet}{r}\right)\right)^p, \tag{3}$$

where $(\varphi^\circ)^\alpha$ can be interpreted as the probability of bubble continuation, while $1-(\varphi^\circ)^\alpha$ is the probability of the bubble bursting or reverting to causal dynamics, and the term $\left(1 - \frac{\varphi^\bullet}{r}\right)$ represents the correction of the bubble trend based on the realized past return ratio r_t. This decomposition provides a direct economic interpretation of the role of higher-order moments in portfolio allocation. In particular, conditional skewness and kurtosis reflect the coexistence of large continuation gains and crash risk, and therefore vary endogenously with the state of the system.

3. Monte Carlo experiments

Section 2 shows that optimal BP decisions depend on the conditional moments of returns, which in turn are driven by the bubble asset prices. The purpose of this simulation analysis is to assess how the mapping from the state variable to optimal strategy pairs (ω^*, h^*), through the conditional higher-order moments, is affected by parameter estimation.

We simulate $M = 500$ trajectories of $N = \{500, 1000, 5000\}$ observations from the MAR(1,1) process $(1-0.9F)(1-0.1B)X_t = \varepsilon_t$ where $\varepsilon_t \overset{i.i.d.}{\sim} S(1.8, 0.3, 1.5, 15)$. We then estimate the conditional power moments by replacing the theoretical constants σ_1^α, β_1, κ_p, λ_p in Proposition 1 by their empirical counterparts computed by plugging-in the MAR(1,1) parameter estimates obtained by Maximum Likelihood.[6]

Estimating mixed causal–noncausal processes remains challenging, and only a limited number of estimators are currently available. The main alternatives are the semi-parametric Generalized Covariance (GCoV) estimator of Gourieroux and Jasiak [2023] and the (approximate) maximum likelihood (ML) estimator of Lanne and Saikkonen [2011]. We rely on the latter throughout the paper. While GCoV has the advantage of not requiring a parametric specification for the innovation distribution, it is generally less statistically efficient than ML for estimating the autoregressive parameters $(\varphi^\circ, \varphi^\bullet)$, owing to its semi-parametric nature. Since these parameters enter all the quantities $(\sigma_1^\alpha, \beta_1, \kappa_p, \lambda_p)$ appearing in Proposition 1, their estimation uncertainty directly propagates to the conditional moments used in the portfolio allocation problem. We therefore favor ML estimation, as our approach requires the additional assumption that ε_t follows an α-stable distribution. Standard errors are computed from the numerical (finite-difference) Hessian of the log-likelihood evaluated at $\hat{\theta}$ and should be interpreted with caution. Although consistency of the ML estimator has been established for α-stable autoregressive models [Andrews et al., 2009], its asymptotic distribution is non-standard and its rate of convergence depends on the tail index α. Consequently, the inverse Hessian only provides an approximate measure of estimation uncertainty. Bootstrap-based inference has been proposed in the literature [Andrews et al., 2009, Cavaliere et al., 2020], but existing procedures are developed for purely noncausal specifications and are not directly applicable to the mixed MAR(1,1) framework considered here. Finally, as shown by the Monte Carlo evidence reported below, the optimal allocation rules are relatively insensitive to estimation uncertainty, suggesting that the economic conclusions are considerably more robust than asymptotic inference on individual parameters.

The results are displayed in Figure 1 for prediction horizons $h = 1, 3, 5, 10$ and conditioning values that correspond to the 0.0005% and 99.9995% quantiles of the marginal distribution of X_t. These results take the form of a pointwise 5% - 95% interquantile interval of the conditional moment estimators for each sample size N. Notice that the theoretical conditional moments,

[6]To facilitate the estimation, we initialize the parameters of the α stable distribution by relying on the approach of McCulloch [1986]. Provided the ML estimator is consistent, which is the case for the one used here [see Andrews et al., 2009], the plug-in estimators of the conditional moments will also be consistent.

based on the true values of the parameters and represented by a black line, always belong to the empirical interquantile range. More precisely, the interquantile intervals are quite narrow around most of the true conditional moments curves, even for small sample sizes. They are larger for higher-order moments and large horizons when $N = 500$ but narrow down as the sample size increases. Overall, the plug-in method appears to be a good way to estimate the conditional moments even when the conditioning values $X_t = x$ are in the right tail of the marginal distribution of the process, which is the case of interest here.

In the second step we hence investigate the impact of parameter estimation on the selected portfolios. The simulated conditional moments of returns obtained from the MLE estimates of the MAR(1,1) process are plugged in the CRRA portfolio optimization program to get the optimal portfolio strategies in the form of couples (ω^*, h^*), which define the part of the wealth to invest in the bubble asset and the horizon of this investment given that the overall investment horizon is fixed to $H = 26$ periods, i.e., six-months of weekly trading activity. To be more precise, we search for optima (ω^*, h^*) in the set $[-1, 1] \times [0, 26]$, thus allowing short strategies. As the optimization program is likely non-convex, several strategies may lead to the same terminal wealth, and in this case they are all labeled as optimal strategies. We round ω^* to the closest percentage point and report h^* in weeks. Besides, by convention, if either $\omega^* = 0$ or $h^* = 0$, we report $(\omega^*, h^*) = (0, 0)$. For the bubble-free asset, we set υ and ς so that the annual return and volatility equal 2%.

Figures 2 and 3 display the optimal investment strategies under correct specification of the DGP and the impact of parameter estimation for investors with CRRA ($\gamma = 5$) and power-three LIRA preferences. Results are generally robust to the preference specification, see the Web appendix for the other four cases. For each initial value $X_t = x$, defined by a quantile of its marginal distribution, and each sample size N, we display the empirical distribution of optimal strategies across simulations in the share–horizon space. Colored dots represent the frequency of selected strategies, while black target symbols denote the true optima.[7]

While the starting value of S_t is immaterial, the initial level of the speculative asset X_t plays a central role in shaping the investment landscape, giving rise to strongly state-dependent allocation patterns. We first discuss the theoretical optimal strategies and then turn to their simulation-

[7]We do not report the precise values as they vary across subplots due to the multiple equilibrium issue discussed earlier and the plot would become too dense.

Figure 1: Conditional moments of return distribution

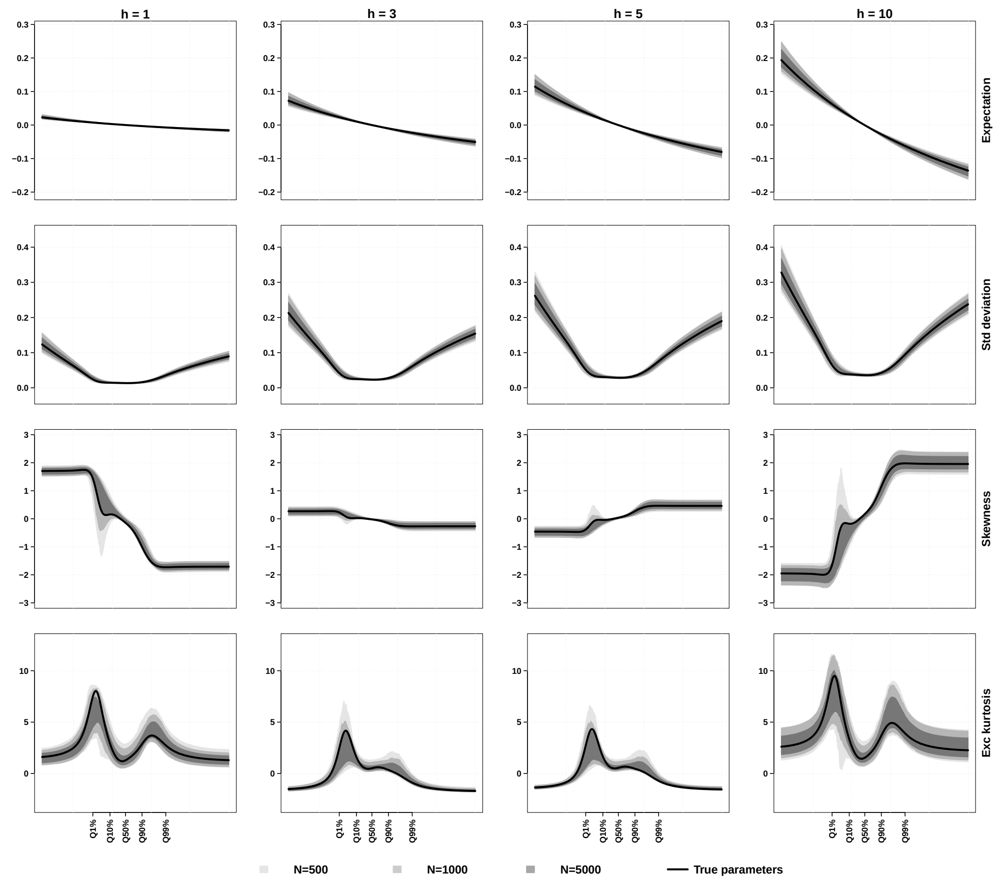

Notes: First four conditional moments of returns when the price series follows a MAR(1,1) process $(1-0.9F)(1-0.1B)X_t = \varepsilon_t$ with $\varepsilon_t \overset{i.i.d.}{\sim} S(1.8, 0.3, 1.5, 15)$. The conditional moments are obtained for conditioning values $X_t = x \in (7, 391)$, i.e. 99.999% of the probability mass of the marginal distribution of X_t is supported on this interval. The black line represents the theoretical moments, whereas the gray shaded areas correspond to the simulated conditional moments based on 500 simulated trajectories of the MAR(1,1) process for the three sample sizes. In the latter case the MAR(1,1) parameters are estimated by MLE and plugged in the formulae of Proposition 1. The results are displayed for four horizons, $h = 1, 3, 5, 10$ and three sample sizes, $N = 500, 1000, 5000$.

based counterparts.

At low quantiles (left tail), optimal strategies consist of long positions with relatively large horizons (around 12 to 20 periods), reflecting expectations of mean reversion and significant upside potential. As the conditioning value moves toward the lower-middle part of the distribution (approximately $Q_{0.2}$–$Q_{0.3}$), both the optimal share and the investment horizon decrease, indicating a more cautious stance as the persistence of favorable dynamics weakens.

A clear regime shift occurs around the center of the distribution. For quantiles at and above the median, optimal strategies switch to short positions held over relatively long horizons, reflecting increasing overvaluation concerns. As the quantile rises further, short exposure intensifies while the investment horizon shortens, consistent with a growing likelihood of a near-term correction. In the upper tail, this pattern becomes non-monotonic: while strong short positions dominate intermediate high quantiles (e.g., $Q_{0.7}$–$Q_{0.95}$), exposure gradually declines as one moves to the extreme right tail, with optimal shares reverting toward zero. This reflects a trade-off between rising crash risk and the possibility of continued bubble expansion, which increases uncertainty about timing.

Turning to the simulation-based strategies, estimated optima are asymptotically tightly clustered around the theoretical ones, indicating that parameter uncertainty has a limited impact on portfolio allocation. However, the empirical distribution exhibits some polarization, particularly in the upper-middle of the distribution, with a non-negligible mass around corner or near-corner allocations at zero and H. This reflects heterogeneous beliefs about bubble dynamics, leading to either very short horizons or longer horizons combined with moderate exposures. Importantly, dispersion is asymmetric across decision variables. While the investment share (ω^*) remains relatively well identified, uncertainty primarily affects the horizon choice (h^*), especially when the state variable lies in the upper tail of the distribution. This reflects the difficulty of timing exit in environments where continuation and crash regimes coexist.

As the sample size increases, dispersion shrinks and the mass concentrates around the true optima, supporting the consistency of the plug-in approach. This contraction remains more pronounced in the share dimension, while residual uncertainty persists in the horizon choice near the bubble peak. Importantly, these results are robust across the utility specifications considered and to alternative parameterizations of risk aversion, as well as to changes in the underlying data-generating process.[8] Overall, the evidence highlights the strongly state-dependent nature of optimal strategies and the presence of sharp regime shifts in both exposure and investment horizons, features that cannot be captured by standard allocation frameworks.

4. Economic Value

Sections 2–3 establish that the BP framework generates state-dependent portfolio strategies through conditional higher-order moments. In this section, we evaluate the economic value of this mechanism by comparing our BP strategies to standard benchmarks that rely on unconditional moments and therefore do not adjust to changes in the distribution of returns.

To assess the performance gains from our approach, we generate 500 MAR(1,1) trajectories of length $N = 1000$ in the same setup as in Section 3. For each trajectory, the first half of the sample is used to estimate the conditional distribution of returns and to derive optimal strategies (ω^*, h^*) over a grid of quantiles of the state variable X_t. These strategies are then retained and used for the out-of-sample evaluation. For each out-of-sample observation, the realized value of the state variable X_t is mapped to the closest in-sample quantile below it, and the corresponding optimal strategy is implemented. The portfolio is subsequently held for the prescribed horizon h, implying at most one reallocation of wealth from the speculative asset to the safe asset before the terminal date H. This procedure yields out-of-sample allocations that remain explicitly state-dependent without implementing dynamic trading rules.

Both benchmark portfolios are implemented as buy-and-hold strategies over the H-period investment horizon, with no interim rebalancing. The EW benchmark allocates a fixed 50%/50% share to the speculative and safe assets, respectively, while the MV benchmark uses the portfolio weight obtained in-sample by minimizing the variance of the non-overlapping H-period return. This weight is estimated once and then applied unchanged to all out-of-sample investment dates. This setup differs from the empirical application of Section 5.2, where all portfolios are rebalanced weekly.

Tables 1 and 2 summarize the results for the same utility functions as in Section 3. Several important findings emerge. First, BP strategies dominate benchmark portfolios in terms of terminal

[8] The results are qualitatively similar when using a non-causal $AR(1)$ DGP, although this specification imposes an abrupt crash dynamic that is less flexible than the $MAR(1,1)$ framework.

Figure 2: Optimal portfolio strategies (CRRA $\gamma = 5$ utility)

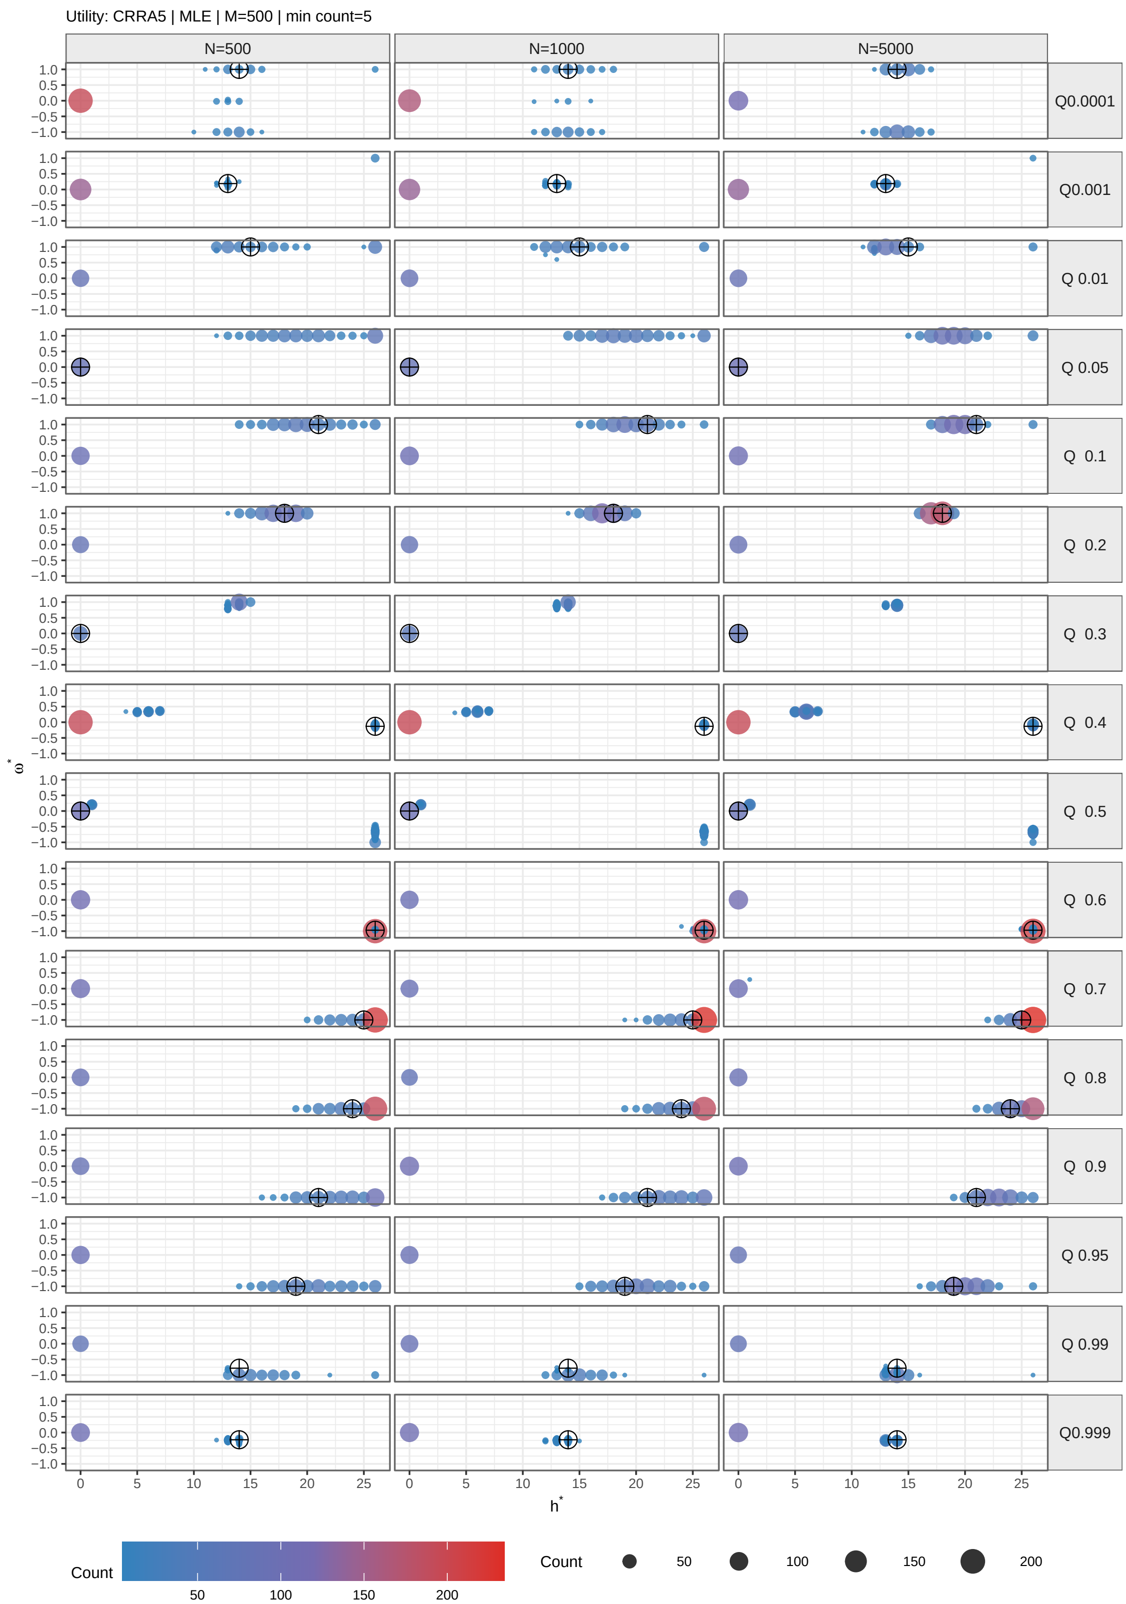

Notes: Mass repartition of the optimal portfolio strategies when the speculative asset's parameters are estimated by MLE across 500 simulated trajectories of length $N = 500, 1000, 5000$ trading days and for several starting values defined by the quantiles of the true marginal distribution of X_t. The DGP for the speculative asset price is a MAR(1,1) process $(1 - 0.9F)(1 - 0.1B)X_t = \varepsilon_t$ with $\varepsilon_t \overset{i.i.d.}{\sim} S(1.8, 0.3, 1.5, 15)$. The results are displayed in the share (vertical axis) - horizon (horizontal axis) space. The larger and redder the dots, the bigger the proportion of selected portfolios falling in that area across the simulations. A black target symbol indicates a *true* optimal portfolio, i.e. obtained for the true values of the parameters.

Figure 3: Optimal portfolio strategies (LIRA Power 3 utility)

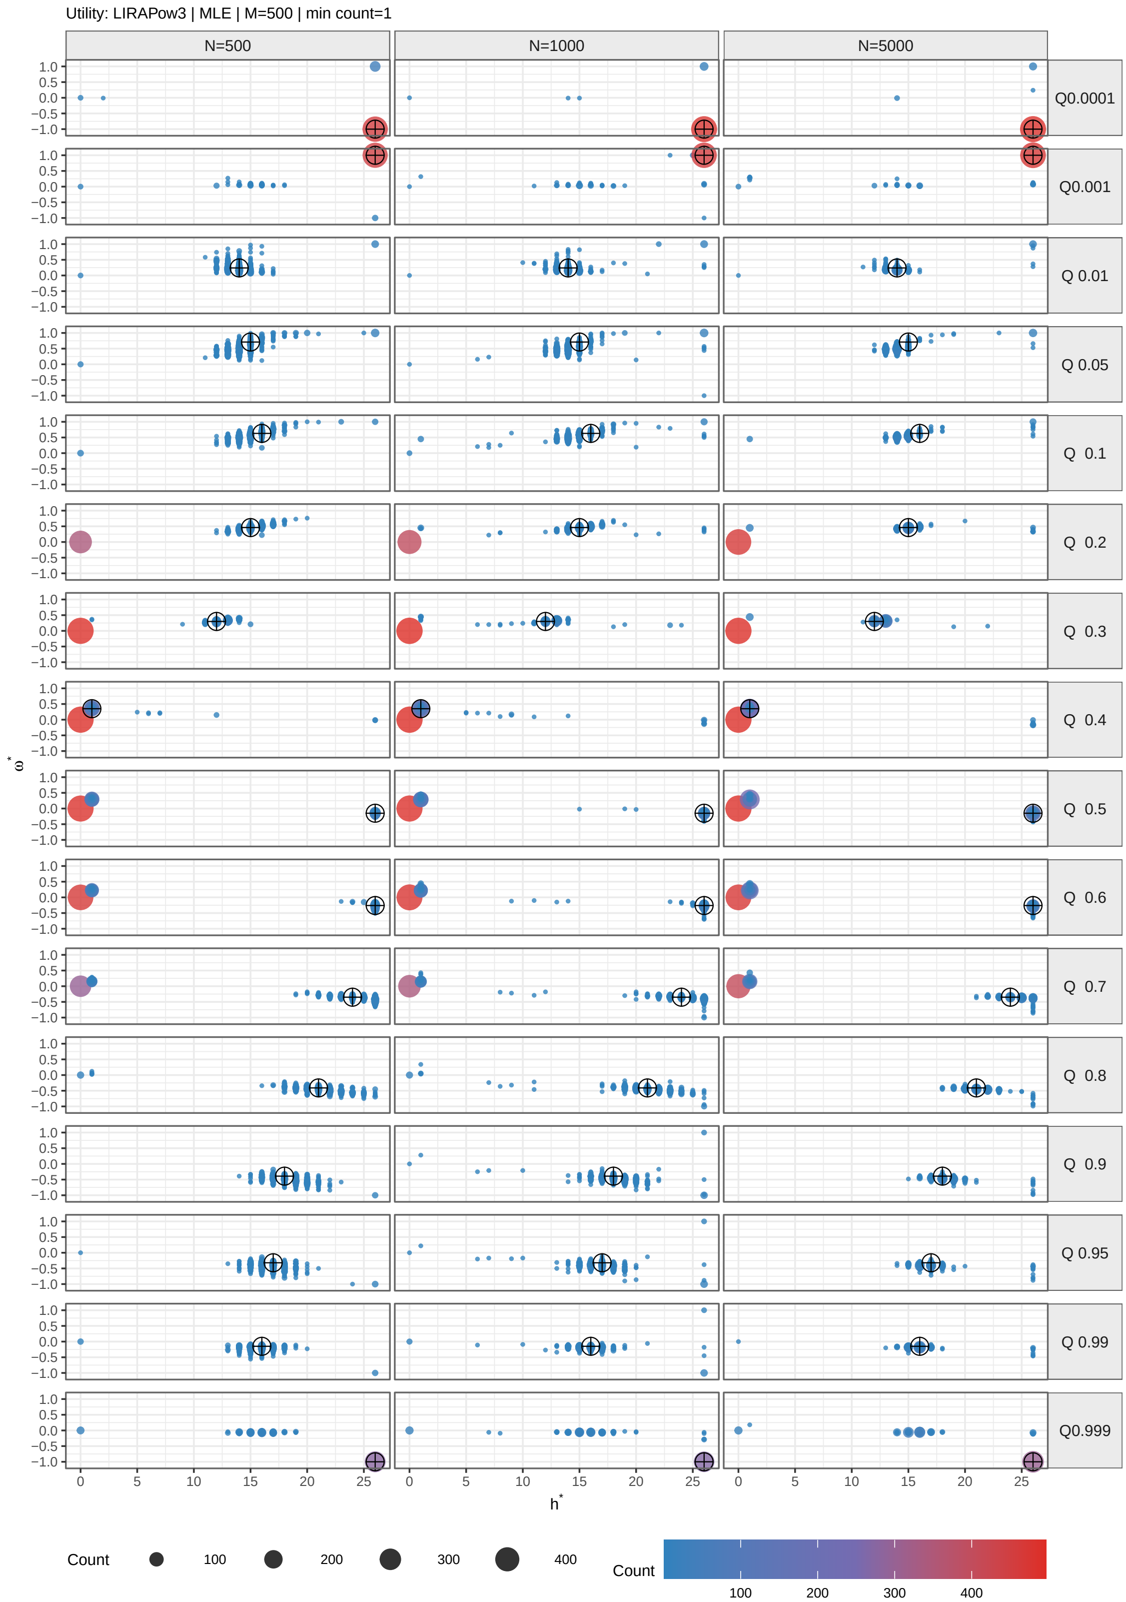

Notes: See note to figure 2.

Table 1: Relative performance of portfolio strategies — CRRA ($\gamma = 5$)

Sample	Metric		Strategy				
			BP_{med}	BP_{inf}	BP_{sup}	EW	MV
Full sample	Wealth	μ	$1.031^{***/***}$	$1.030^{***/***}$	$1.032^{***/***}$	1.005	1.010
		σ	0.008	0.008	0.007	0.003	0.003
	OC vs EW	μ	0.027	0.026	0.028		
		σ	0.009	0.010	0.009		
	OC vs MV	μ	0.024	0.023	0.025		
		σ	0.008	0.008	0.008		
Bubble periods	Wealth	μ	$1.087^{***/***}$	$1.087^{***/***}$	$1.087^{***/***}$	0.962	1.005
		σ	0.036	0.036	0.036	0.016	0.011
	OC vs EW	μ	0.119	0.118	0.120		
		σ	0.047	0.047	0.046		
	OC vs MV	μ	0.076	0.075	0.076		
		σ	0.039	0.039	0.039		

Notes: MAR(1,1)-based strategies are compared with equally weighted (EW) and mean-variance (MV) benchmarks. Results are reported for the full sample and for bubble periods ($X_t \geq Q_{95\%}$). Terminal wealth and opportunity cost (OC) are summarized by their mean (μ) and standard deviation (σ). Asterisks $^{*/**/***}$ indicate rejection of one-sided Wilcoxon tests (BP $>$ benchmark) at the 90%, 95%, and 99% levels versus EW/MV, respectively.

wealth across all specifications. In the full sample, the average terminal wealth of BP strategies is between 1.017 and 1.035, compared to 1.01 for the benchmarks. While these differences may appear moderate, they are highly statistically significant and economically meaningful given the short investment horizon.

Second, the performance gains become substantially larger in bubble regimes. When conditioning on the upper tail of the state distribution ($X_t \geq Q_{95\%}$), BP strategies achieve average terminal wealth levels between 1.041 and 1.099, while the EW strategy performs poorly (below 1), reflecting its inability to adjust to looming crash risk. This highlights the importance of state-dependent allocation when the return distribution becomes highly asymmetric.

Third, the opportunity cost results confirm these findings from a utility perspective. The OC is positive in all cases, indicating that investors would be willing to pay to switch from benchmark strategies to BP allocations. Moreover, the magnitude of the OC increases sharply in bubble regimes (e.g., above 9% relative to EW), underscoring the substantial welfare gains from incorporating conditional information. These results stem from the ability of BP strategies to adapt to the changing balance between bubble continuation and crash risk. In contrast, MV and EW portfolios rely on unconditional moments and cannot adjust to shifts in the conditional

Table 2: Relative performance of portfolio strategies — LIRA Power 3 utility

Sample	Metric		Strategy				
			BP_{med}	BP_{inf}	BP_{sup}	EW	MV
Full sample	Wealth	μ	$1.019^{***/***}$	$1.017^{***/***}$	$1.023^{***/***}$	1.005	1.010
		σ	0.005	0.005	0.005	0.003	0.003
	OC vs EW	μ	0.018	0.015	0.021		
		σ	0.006	0.006	0.006		
	OC vs MV	μ	0.014	0.012	0.018		
		σ	0.008	0.009	0.008		
Bubble periods	Wealth	μ	$1.049^{***/***}$	$1.048^{***/***}$	$1.051^{***/***}$	0.962	1.005
		σ	0.023	0.025	0.023	0.016	0.011
	OC vs EW	μ	0.090	0.088	0.091		
		σ	0.029	0.031	0.029		
	OC vs MV	μ	0.047	0.045	0.048		
		σ	0.023	0.025	0.023		

Notes: See note to Table 1.

distribution of returns.

Overall, the Monte Carlo evidence highlights the economic value of state-dependent portfolio allocation in environments characterized by nonlinear and heavy-tailed dynamics. These findings are robust across CRRA risk aversion levels and alternative LIRA parameterizations (see the Web Appendix).

5. Empirical application

5.1. *Data and in-sample analysis*

We now assess whether the state-dependent allocation mechanism identified in the previous sections is relevant in real data. We consider four commodity price series, *Brent crude oil, rice, soybean-oil, and soybeans*, as candidates for the speculative asset. Commodity markets provide a particularly suitable environment for studying bubble dynamics, as prior evidence documents the presence of nonlinear, forward-looking, and noncausal features in their price behavior [Gourieroux et al., 2021, Cubadda et al., 2023, Hecq and Velasquez-Gaviria, 2025, Blasques et al., 2023]. As a benchmark bubble-free asset, we use the settlement price of the 30-year U.S. Treasury bond. Weekly data covering January 2000 to March 2025 are obtained from Refinitiv Eikon. The sample is split into an in-sample period (2000–2012) and an out-of-sample period (2012–2025).

A central issue in empirical asset pricing is whether the apparent nonstationarity of price series reflects genuine unit-root behavior or instead arises from nonlinear dynamics associated with speculative episodes. A large literature shows that bubble processes with periodically collapsing dynamics can closely mimic unit-root behavior, making standard linear tests unreliable [Evans, 1991]. More recent approaches interpret financial time series as sequences of locally explosive regimes embedded within persistent dynamics [Phillips et al., 2011, 2015]. In this context, failure to reject a unit root does not provide evidence against bubble behavior, but rather reflects the low power of linear tests in the presence of nonlinearities.

Given this ambiguity, we do not rely on unit-root testing to characterize the stochastic properties of commodity prices. Instead, we adopt a structural decomposition consistent with the MAR framework introduced in Section 2.2. Specifically, we decompose observed prices as

$$Y_t = T(t) + X_t,$$

where $T(t)$ captures low-frequency deterministic movements driven by fundamentals, and (X_t) is a stationary component governed by nonlinear and potentially noncausal dynamics. This decomposition separates long-run trends from short-run speculative fluctuations, the latter being the focus of the portfolio allocation problem. Importantly, it avoids confounding nonlinear dynamics with stochastic trends, which could arise when applying standard differencing or filtering techniques.

As emphasized by Hencic and Gouriéroux [2015], using deterministic detrending methods is preferable to filtering-based approaches, as the latter can induce spurious evidence of noncausality. We therefore estimate and remove a deterministic trend prior to modeling the stochastic component. Among the available approaches, and by following Hecq and Voisin [2019], we adopt a second-order polynomial specification, which strikes a balance between flexibility and robustness in out-of-sample applications.[9] The resulting series exhibit pronounced nonlinear and bubble-like features, consistent with the MAR representation (see Figure 4).

The detrended component (X_t) is modeled using mixed causal-noncausal autoregressive specifications. More specifically, a MAR(1,1) structure is selected for all assets. This choice is primarily

[9] A first-order polynomial is retained if the second-order coefficient is not statistically significant.

motivated by tractability: it allows us to exploit the closed-form expressions for conditional moments derived in Section 2.2, while preserving the key features of bubble dynamics. Even if possibly slightly misspecified for some time-series, this specification captures the essential nonlinearities and asymmetries that drive state-dependent portfolio decisions which is of utmost importance in a forecasting exercise. Throughout the analysis, we consider a fixed investment horizon of $H = 26$ weeks, corresponding to a six-month investment period. For each asset, the estimated MAR model and the associated conditional moments serve as inputs to the portfolio optimization problem, yielding state-dependent allocation strategies (ω^*, h^*) that vary with the level of the detrended price process.

Figure 4: Detrended price series for each asset

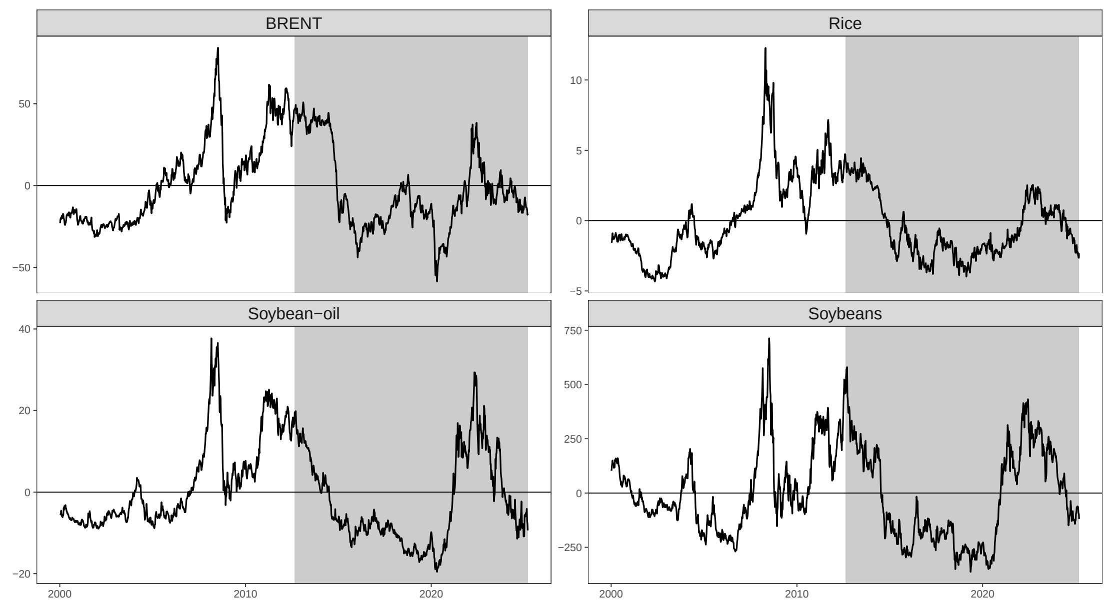

Notes: The shaded area corresponds to the out-of-sample data.

Table 3 reports the estimated trend and MAR(1,1) parameters for the commodity series. The polynomial trend coefficients are statistically significant, validating the detrending strategy. Importantly, the noncausal component is always significant and dominates the causal part, reflecting the asymmetric dynamics of periods of gradual price increases followed by abrupt corrections, consistent with speculative bubble behavior. This suggests that price dynamics are largely driven by forward-looking expectations rather than backward-looking adjustments, which implies that the current price level contains information about future returns beyond what is captured by standard autoregressive models. Besides, in all cases, the estimated tail indices are significantly below two, indicating substantial departures from Gaussianity and the presence of heavy-tailed risks. In the case of *Soybean-oil*, α is slightly inferior to the 1.5 threshold of existence of the conditional kurtosis, but we retain the series and treat it as a slightly misspecified case.

Table 3: Estimates of the polynomial trend and α-stable MAR(1,1) models

	Trend			α-stable MAR(1,1)				
	(Intercept)	τ_1	τ_2	$\varphi^{\bullet}$	φ°	α	β	σ
Brent								
	17.5^{***}	0.09842^{***}	$7.208e-05^{***}$	0.01341^{***}	0.9725^{***}	1.532^{***}	0.099^{***}	1.490^{***}
	(1.79E+00)	(1.26E-02)	(1.85E-05)	(8.70E-07)	(1.45E-07)	(7.0001E-07)	(1.3002E-06)	(1.4006E-06)
Rice								
	3.794^{***}	0.01956^{***}		0.1052^{***}	0.9742^{***}	1.510^{***}	0.010^{***}	0.183^{***}
	(1.64E-01)	(4.31E-04)		(2.15E-07)	(4.35E-08)	(1.9002E-07)	(2.4006E-07)	(3.4008E-08)
Soybean oil								
	14.49^{***}	0.03057^{***}	$5.401e-05^{***}$	0.09669^{***}	0.9568^{***}	1.466^{***}	-0.018^{***}	0.554^{***}
	(8.12E-01)	(5.69E-03)	(8.36E-06)	(2.58E-06)	(2.73E-07)	(9.2003E-07)	(4.7009E-06)	(8.4001E-07)
Soybeans								
	461.4^{***}	0.3387^{***}	0.001705^{***}	0.09382^{***}	0.9424^{***}	1.571^{***}	0.159^{***}	15.930^{***}
	(1.72E+01)	(1.21E-01)	(1.78E-04)	(1.99E-07)	(1.57E-07)	(2.9000E-06)	(5.6001E-07)	(1.9003E-05)

Notes: The table reports estimates of the polynomial trend and α-stable MAR(1,1) models. Standard errors are reported in parentheses and computed from the numerical (finite-difference) Hessian of the log-likelihood at the estimated parameters. Asterisks *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table 4: Moments of the US 30-year Treasury bond returns

	μ	σ
Returns	0.0007	0.0145

Notes: The table reports the mean and the standard deviation of the US 30-year Treasury bond returns over the period January 2000 to August 2012.

Table 4 reports the mean and standard deviation of the safe asset, modeled as a geometric Brownian motion, in line with standard assumptions in the portfolio optimization literature. Given its relatively low volatility and stable dynamics, the safe asset mainly acts as a hedging instrument, portfolio performance being primarily driven by the interaction between bubble dynamics and the endogenous investment horizon.

Table 5: Optimal portfolio strategies (Brent)

	$q_{0.500}$	$q_{0.600}$	$q_{0.700}$	$q_{0.800}$	$q_{0.900}$	$q_{0.950}$	$q_{0.990}$	$q_{0.999}$	$q_{0..9999}$
CRRA									
$\gamma = 5$	(0.390,26)	(0.480,1) (0.310,26)	(0.470,1)	(0.450,1)	(0.410,1)	(0.360,1) (-0.150,26)	(-0.120,26)	(-0.080,26)	(-0.070,26)
$\gamma = 10$	(0.480,1)	(0.470,1)	(0.460,1)	(0.450,1)	(0.420,1)	(0.370,1)	(0.210,1)	(-0.040,26)	(-0.030,26)
LIRA									
LIRAGauss	(1,26)	(1,26)	(1,26) (0.480,1) (0,0) (-1,26)	(-1,26) (0.460,1) (1,26) (0,0)	(-1,26)	(-1,26)	(-1,8)	(-1,5)	(-1,4)
LIRADExp	(0.320,26) (0.480,1) (0,0)	(0.480,1) (0,0)	(0.470,1) (0,0)	(0.450,1) (0,0)	(0.420,1) (-0.010,26) (0,0)	(0.370,1) (-0.070,26) (0,0)	(0.210,1) (-0.070,26) (0,0)	(-0.050,26) (0.090,1) (0,0)	(-0.040,26) (0.060,1) (0,0)
LIRAUnif	(0.430,26)	(0.480,1) (0.340,26)	(0.470,1)	(0.460,1)	(0.430,1)	(0.390,1)	(-0.100,26) (0.240,1)	(-1,26)	(-1,26)
LIRAPow3	(0.480,1)	(0.470,1)	(0.470,1)	(0.460,1)	(0.430,1)	(0.380,1)	(0.220,1)	(-1,26)	(-1,26)

Notes: Optimal portfolio strategies (ω^*, h^*) conditional on the quantile of the detrended in-sample distribution in which the bubble asset is at the time of the investment. ω^* is reported in percentages of the investment and h^* in weeks.

Using the estimated MAR(1,1) parameters, we compute the conditional moments of the detrended price series and, subsequently, those of the observed returns by reintroducing the trend dynamics at the time of investment. These conditional moments of returns serve as inputs to the portfolio optimization programs, which deliver optimal strategies in terms of couples (ω^*, h^*) for each quantile of the detrended state variable. Multiple optimal strategies may arise, as all allocations within one percent of the maximal expected utility are considered equivalent and retained. This reflects the non-convexity induced by higher-order moments and nonlinear dynamics. Economically, this reflects the coexistence of heterogeneous beliefs regarding bubble continuation versus collapse, leading to multiple locally optimal but qualitatively distinct allocation strategies. The results are broadly consistent across assets.

For brevity, Table 5 reports the strategies for the *Brent* portfolio over the upper quantiles of the state distribution, while the results for the other assets are reported in the Web Appendix. When the conditioning variable $X_t = x$ moves towards the upper quantiles, the conditional distribution becomes increasingly skewed and heavy-tailed (and asymptotically bimodal), reflecting a rising probability of abrupt corrections. This feature, which is absent from linear Gaussian models, generates a sharp asymmetry between continuation gains and crash risk. As a result, optimal portfolio strategies become inherently state-dependent: investors optimally "ride the bubble" at intermediate quantiles (long positions), but reduce exposure or take short positions as the asset approaches the upper tail of its distribution, where the likelihood of a regime shift increases sharply. Such behavior cannot be replicated by standard allocation rules, such as mean–variance or equally weighted portfolios, which rely on unconditional moments and ignore the current state of the system. This feature will be illustrated in Section 5.4 through a decision-tree representation of optimal investment paths. Note also that the results reported in the Web Appendix indicate that these features of portfolio strategies are stable across assets.

Overall, these findings suggest that accounting for nonlinear and forward-looking dynamics in the risky asset price materially affects optimal portfolio behavior, leading to state-dependent strategies that dynamically balance exposure to bubble continuation against crash risk, a mechanism absent from standard portfolio allocation frameworks.

5.2. Out-of-sample performance analysis

We now evaluate the real-time performance of the proposed portfolio strategies using out-of-sample data. At each investment date, the polynomial trend is recursively re-estimated and removed from prices, yielding a detrended state variable X_{T+k}. This value is mapped to the closest lower empirical quantile of the in-sample distribution, which determines the corresponding optimal allocation (ω^*, h^*). This is a slightly conservative choice, which implies the investor gives more weight to bubble crashes than to its continuation.

To ensure economic realism, portfolios are dynamically rebalanced. Because the optimal holding horizon h is state-dependent, rebalancing occurs at irregular intervals and depends on the realized path of the state variable. At each rebalancing date, the strategy is updated based on the newly observed state. Positions in the speculative asset are liquidated either when the implied holding horizon exceeds the remaining investment period H, or when the optimal allocation collapses to $(\omega^*, h^*) = (0, 0)$. Figure 5 illustrates this rebalancing mechanism.

Because the mapping from the state variable to optimal strategies may yield multiple solutions

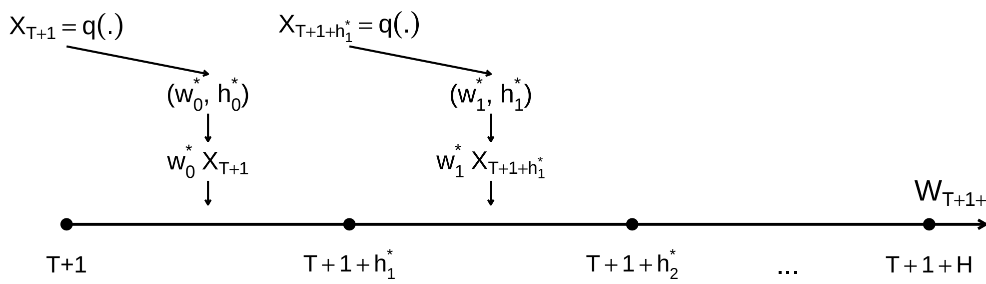

Figure 5: Rebalancing mechanism

and because dynamic rebalancing introduces path dependence, performance is evaluated over the distribution of investment outcomes. Our results are hence reported in terms of quantiles of terminal wealth distribution across paths and out-of-sample period, denoted $BP_{10\%}$ for the worst 10%, $BP_{50\%}$ for the median of the distribution, and $BP_{90\%}$ for the best 10%. [10]

Table 6: Number of rebalancings and strategy paths – $BP_{50\%}$

		LIRA				CRRA	
		LIRAGauss	LIRADExp	LIRAUnif	LIRAPow3	$\gamma = 5$	$\gamma = 10$
Brent	$\mathrm{Med}_{\mathrm{Rebal}}$	0.000	3.000	0.000	17.000	0	3
	σ_{Rebal}	1.535	3.444	4.785	5.659	4.574	—
	$\mathrm{Med}_{\mathrm{Paths}}$	1	27	1	1	1	1
	σ_{Paths}	4.900	12.100	4.360	9.330	6.242	0.000
Rice	$\mathrm{Med}_{\mathrm{Rebal}}$	0.000	4.000	0.000	2.250	0	0
	σ_{Rebal}	0.000	4.499	5.311	4.616	—	4.542
	$\mathrm{Med}_{\mathrm{Paths}}$	1	27	1	1	1	1
	σ_{Paths}	0.310	10.400	4.870	24.790	0.000	1.630
Soybean oil	$\mathrm{Med}_{\mathrm{Rebal}}$	0.000	0.000	0.000	0.000	0	0
	σ_{Rebal}	1.879	5.069	4.847	0.848	2.160	—
	$\mathrm{Med}_{\mathrm{Paths}}$	2	1	1	1	1	1
	σ_{Paths}	9.180	24.980	3.560	0.720	1.546	0.000
Soybeans	$\mathrm{Med}_{\mathrm{Rebal}}$	0.000	0.000	0.000	0.000	0	0
	σ_{Rebal}	2.729	2.563	3.249	4.070	2.272	4.145
	$\mathrm{Med}_{\mathrm{Paths}}$	1	2	1	1	1	1
	σ_{Paths}	8.310	12.230	3.790	1.420	3.754	1.417

Notes: Median and standard deviation of the number of rebalancings and strategy paths for the $BP_{50\%}$ strategies across out-of-sample periods. For the equally-weighted and mean-variance portfolios, there is always one path with 26 rebalancings for each asset.

Table 6 reports summary statistics on rebalancing activity and the number of strategy paths. Two robust patterns emerge. First, BP strategies involve markedly fewer rebalancings than bench-

[10] This notation differs from the BP_{med}, BP_{inf}, and BP_{sup} statistics of Section 4: The latter refer to different locally optimal strategies for a given state, whereas $BP_{10\%}$, $BP_{50\%}$, and $BP_{90\%}$ denote quantiles of the out-of-sample terminal wealth distribution across dates and paths.

mark portfolios, despite their state-dependent nature. Indeed, both benchmark portfolios are rebalanced weekly over the $H = 26$ periods in order to maintain fixed target allocations. The EW benchmark is rebalanced to a constant 50%/50% allocation, while the MV benchmark is rebalanced to the portfolio weight estimated once in-sample from the mean–variance optimization problem. Although these target weights remain unchanged throughout the out-of-sample period, portfolio positions are adjusted at every period to restore the target allocation.

Table 7: Out-of-sample terminal wealth and Sharpe ratios – $BP_{50\%}$

			LIRA				CRRA			
			LIRAGauss	LIRADExp	LIRAUnif	LIRAPow3	$\gamma = 5$	$\gamma = 10$	*EW*	*MV*
Brent	Wealth	μ	$0.970^{***/***}$	1.000	1.000	$1.010^{***/***}$	1.000	$1.000^{\cdot/**}$	1.000	1.000
		σ	0.160	0.050	0.060	0.080	0.060	0.070	0.110	0.050
	Sharpe	μ	$-0.030^{***/**}$	$0.010^{\cdot/*}$	-0.000	$0.040^{*/***}$	-0.010	$0.020^{\cdot/**}$	0.010	-0.000
		σ	0.190	0.170	0.190	0.160	0.180	0.180	0.200	0.170
	mSharpe	μ	$-0.030^{***/*}$	$0.010^{\cdot/*}$	-0.010	$0.040^{**/***}$	$-0.010^{*/\cdot}$	$0.020^{\cdot/**}$	0.010	-0.000
		σ	0.210	0.180	0.220	0.200	0.210	0.210	0.200	0.170
Rice	Wealth	μ	$0.990^{***/*}$	$1.000^{***/**}$	1.000	$1.000^{***/***}$	1.000	$1.000^{*/*}$	1.000	0.990
		σ	0.110	0.040	0.060	0.050	0.060	0.070	0.060	0.060
	Sharpe	μ	$-0.040^{***/**}$	$0.010^{*/**}$	-0.010	$0.010^{*/**}$	-0.010	-0.010	-0.000	-0.010
		σ	0.200	0.150	0.190	0.170	0.190	0.180	0.170	0.190
	mSharpe	μ	$-0.040^{***/**}$	$0.000^{\cdot/*}$	-0.020	$0.000^{\cdot/*}$	-0.020	-0.020	-0.010	-0.020
		σ	0.210	0.170	0.200	0.180	0.200	0.200	0.170	0.190
Soybean oil	Wealth	μ	$0.990^{\cdot/*}$	1.000	$1.000^{**/*}$	0.990	$1.000^{*/\cdot}$	1.000	1.000	0.990
		σ	0.070	0.080	0.060	0.070	0.070	0.080	0.080	0.060
	Sharpe	μ	-0.020	-0.000	$0.010^{**/*}$	-0.020	0.000	-0.010	-0.010	-0.010
		σ	0.180	0.200	0.190	0.200	0.200	0.200	0.210	0.190
	mSharpe	μ	-0.030	-0.000	$0.010^{*/\cdot}$	-0.020	0.000	-0.010	-0.020	-0.010
		σ	0.200	0.210	0.210	0.210	0.210	0.200	0.200	0.190
Soybeans	Wealth	μ	0.990	$1.000^{***/***}$	$1.000^{***/***}$	$1.000^{***/**}$	$1.000^{***/***}$	$1.000^{**/\cdot}$	0.990	0.990
		σ	0.070	0.050	0.060	0.060	0.060	0.060	0.070	0.060
	Sharpe	μ	-0.010	$0.010^{**/**}$	$0.010^{***/***}$	$0.010^{**/**}$	$0.010^{**/***}$	-0.000	-0.010	-0.020
		σ	0.190	0.170	0.190	0.180	0.180	0.190	0.190	0.180
	mSharpe	μ	-0.010	$0.010^{**/**}$	$0.010^{**/***}$	$-0.000^{*/*}$	$0.010^{**/**}$	-0.010	-0.010	-0.020
		σ	0.210	0.180	0.210	0.190	0.200	0.200	0.180	0.190

Notes: Out-of-sample average (μ) and standard deviation (σ) of terminal wealth, Sharpe ratio, and modified Sharpe ratio (mSharpe) for the $BP_{50\%}$ strategies. EW: equal-weight benchmark; MV: mean-variance benchmark. For μ rows, superscripts denote significance of Wilcoxon tests vs. EW/MV: $(\cdot / * / ** / ***)$ at the 90%, 95%, and 99% levels; a dot $(\cdot)$ indicates non-significance.

By contrast, BP strategies trade only when dictated by the state-dependent allocation rule. Second, the number of distinct BP investment paths is typically close to one, indicating that optimal strategies are stable for a given state. However, substantial heterogeneity arises for some assets, most notably Soybean-oil, especially under LIRA preferences, where the number and dispersion of paths are large. These results point to the presence of multiple local optima induced by nonlinear risk-return trade-offs. This stands in contrast with benchmark portfolios, which rebalance mechanically at each period and admit a unique allocation driven solely by unconditional moments.

We next evaluate the economic performance using terminal wealth, Sharpe ratios, modified Sharpe ratios, and opportunity costs. Terminal wealth W_{T+H} is obtained by applying the sequence of optimal allocations to realized (non-detrended) returns.

Table 7 shows that BP strategies perform on par with, and sometimes outperform, the benchmark portfolios for all assets and specifications. While average gains in terminal wealth are moderate, they are statistically significant and economically meaningful. Risk-adjusted performance confirms these findings. Sharpe and modified Sharpe ratios are slightly higher for BP strategies, indicating that gains are not driven by excessive risk-taking but by improved timing and state-contingent allocation.

Table 8: Opportunity cost – $BP_{50\%}$

			LIRA				CRRA	
			LIRAGauss	LIRADExp	LIRAUnif	LIRAPow3	$\gamma = 5$	$\gamma = 10$
Brent	EW vs	μ	-0.030	-0.010	-0.010	0.010	-0.010	0.000
		σ	0.170	0.110	0.120	0.090	0.120	0.110
	MV vs	μ	-0.030	0.000	0.000	0.010	-0.000	0.010
		σ	0.150	0.050	0.050	0.080	0.050	0.070
Rice	EW vs	μ	-0.000	0.000	0.000	0.010	-0.000	0.010
		σ	0.090	0.050	0.060	0.050	0.060	0.070
	MV vs	μ	0.000	0.010	0.000	0.010	0.000	0.010
		σ	0.090	0.040	0.030	0.030	0.030	0.050
Soybean oil	EW vs	μ	-0.010	0.000	0.000	-0.010	0.000	-0.000
		σ	0.070	0.090	0.070	0.100	0.090	0.100
	MV vs	μ	-0.010	0.010	0.010	-0.000	0.010	0.000
		σ	0.070	0.060	0.040	0.030	0.040	0.040
Soybeans	EW vs	μ	-0.000	0.010	0.010	0.010	0.010	0.010
		σ	0.070	0.060	0.070	0.060	0.070	0.060
	MV vs	μ	-0.000	0.010	0.010	0.010	0.010	0.010
		σ	0.050	0.040	0.030	0.040	0.040	0.040

Notes: Opportunity cost (as a fraction of initial wealth) for the $BP_{50\%}$ strategies vs. EW and MV benchmarks. Out-of-sample average (μ) and standard deviation (σ). A dash indicates that both the median and standard deviation of strategy paths equal zero.

A key feature of the BP approach is its heterogeneous performance across the distribution of outcomes. Results in the Web Appendix show that the lower tail ($BP_{10\%}$) is broadly comparable to, or slightly worse than, the benchmark strategies in terms of terminal wealth and risk-adjusted performance. However, performance improves markedly as one moves up the distribution. In particular, the upper quantile ($BP_{90\%}$) exhibits substantial gains in both terminal wealth and Sharpe ratios, with improvements reaching up to three percentage points relative to the median strategy and significantly outperforming both benchmarks.

From a utility perspective, Table 8 and the additional results reported in the Web Appendix provide further support for the economic value of the proposed approach. Opportunity costs are generally positive for $BP_{50\%}$ and increase significantly for $BP_{90\%}$ strategies, indicating that investors would require compensation to switch to benchmark portfolios. In contrast, opportunity costs for $BP_{10\%}$ strategies are often small or negative, reflecting weaker performance in unfavorable scenarios. This pattern is consistent with the distributional results discussed above and reinforces the interpretation of BP strategies as providing asymmetric payoff profiles.

We next focus on periods where the detrended state variable lies in the upper tail of its distribution, corresponding to bubble dynamics, i.e. above the 75% threshold. This threshold, rather than a higher one, is chosen to retain a sufficiently large number of out-of-sample observations to validate our strategy. The results are reported in the Web Appendix. The BP approach exhibits a larger number of rebalancings and strategy paths, reflecting a more complex conditional dependence structure. At the same time, the relative performance of BP strategies becomes more pronounced in these regimes. While downside performance ($BP_{10\%}$) remains close to that of benchmark portfolios, the upper quantiles deliver substantial gains with controlled risk exposure. It is important to note, however, that in this setting the number of out-of-sample observations is limited. As a result, we rely on the median as a robust measure of central tendency, and the associated Wilcoxon tests are expected to exhibit lower statistical power than in the baseline analysis.

Overall, these findings indicate that the proposed state-dependent allocation mechanism effectively captures the time-varying risk-return trade-offs associated with speculative episodes. By conditioning on the current state of the bubble asset and incorporating higher-order moments of returns, BP strategies deliver economically meaningful improvements over standard benchmark approaches, particularly in environments characterized by nonlinear and asymmetric price dynamics.

5.3. Sensitivity analysis

We assess the sensitivity of our results by isolating the role of higher-order conditional moments in the portfolio optimization problem. In particular, we examine whether the performance gains documented in the previous section are driven by the inclusion of skewness and kurtosis, or whether they primarily reflect the state-dependent nature of the allocation rule.

To this end, we consider a restricted version of the BP framework in which the speculative asset continues to follow a MAR(1,1) process, but the portfolio optimization program relies exclusively on the first two conditional moments of terminal wealth. In this setup, both CRRA and LIRA utility functions are approximated using a second-order Taylor expansion, effectively reducing the allocation problem to a mean–variance framework. The resulting strategies remain state-dependent through the MAR dynamics, but abstract from higher-order risks.

This exercise directly assesses the incremental contribution of higher-order moments beyond the state-dependent allocation mechanism. The corresponding results, reported in Section OA.3 of the Web Appendix, indicate that the mean–variance BP strategies deliver performance levels that are qualitatively similar to those obtained under the full four-moment specification in terms of terminal wealth and opportunity costs, both over the full distribution and when focusing on investments initiated during bubble periods. Moreover, the overall patterns of optimal allocation across states remain broadly unchanged.

However, several important differences emerge. First, the two-moment specification exhibits a larger number of strategy paths and greater dispersion across paths, reflecting a less parsimonious mapping from market states to optimal strategies, and leading to less stable allocation patterns across states. In contrast, incorporating higher-order moments tends to simplify the structure of optimal strategies, resulting in fewer distinct paths and more stable rebalancing behavior, thereby reducing turnover and improving the implementability of the strategy in the presence of transaction costs.

Second, while average performance levels are broadly comparable, the inclusion of higher-order moments improves risk-adjusted outcomes in several cases, as reflected in higher Sharpe and modified Sharpe ratios. This highlights the role of skewness and kurtosis in capturing the asymmetric trade-off between continuation gains and crash risk, which is only partially accounted for in a mean–variance framework.

Overall, these findings emphasize that the primary driver of performance in our framework is the state-dependent allocation mechanism induced by the MAR dynamics. Accounting for higher-order moments provides an additional refinement by improving risk management and stabilizing allocation decisions, but it does not, by itself, fully explain the performance gains. In other words, state dependence is the primary source of economic value, while skewness and kurtosis improve the efficiency, robustness, and implementability of the resulting allocation rules.

5.4. Tree path analysis

To illustrate the economic mechanism underlying our framework, we represent the set of optimal strategies for Brent as a decision tree for the investment date 20 May 2022. This date lies in the out-of-sample period and corresponds to a state in which the detrended price falls in the 99.9% quantile of its in-sample distribution, i.e., an extreme right-tail realization associated with elevated crash risk. The tree provides a compact visualization of the mapping from the state variable to optimal allocation rules and highlights how state-dependent moments translate into discrete allocation regimes. Each node represents a rebalancing date, and strategies are summarized in the form "sign$(\omega) \times h.\omega$," where the sign indicates the position (long/short), ω the portfolio weight, and h the investment horizon in weeks. The root node ("0") corresponds to the initial decision, and "26" denotes the terminal horizon H.

Figure 6 reports the resulting strategy trees for both CRRA and LIRA preferences. A striking feature is that, at such extreme states, most specifications lead to a nearly degenerate solution characterized by a single dominant trajectory: a small short position held over the entire horizon. This reflects the predominance of downside risk when the asset price lies deep in the upper tail of its distribution.

Differences across preferences arise primarily in the timing dimension. Under Gaussian LIRA preferences, the strategy involves infrequent monitoring, with re-optimization occurring every five to eight periods. By contrast, the double-exponential LIRA specification produces a sequential decision structure in which the investor repeatedly trades off continuation gains and crash risks: at each date, the optimal policy alternates between a short-term long position (21% over one period) if the continuation regime is deemed more likely, and a small short position held to maturity otherwise.

Overall, the tree representation provides a transparent visualization of the decision-making process generated by the BP framework. Rather than prescribing a single portfolio allocation, the optimization problem delivers a set of feasible state-contingent strategies that differ jointly in exposure and investment horizon. The resulting paths highlight the importance of timing decisions in speculative environments and illustrate how small changes in the perceived balance between continuation and crash risk can lead to distinct, yet locally optimal, allocation rules. More generally, the tree representation makes clear that the economic value of the BP approach arises from its ability to adapt portfolio decisions to evolving market conditions rather than from a static buy-and-hold allocation.

(a) CRRA

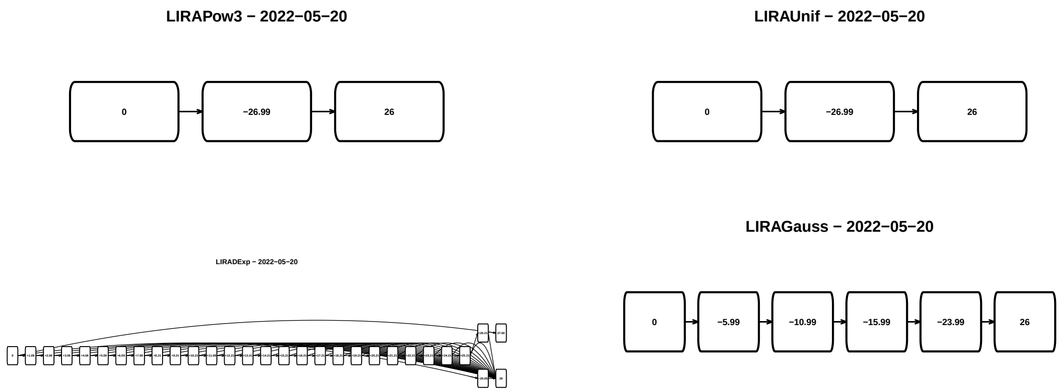

(b) LIRA

Figure 6: Decision-making trees
Reading key: $\text{sign}(\omega) * h.\omega$, with ω the investment share in the bubble asset and h the horizon of this investment.

6. Conclusion

This paper studies portfolio allocation under speculative dynamics using a mixed causal–noncausal (MAR) framework with α-stable innovations. By linking the conditional distribution of returns to the level of the detrended price,the model generates allocation rules that adapt to the changing balance between bubble continuation and crash risk.

Both simulation and empirical results show that the proposed bubble portfolio strategies outperform standard benchmarks, particularly during bubble episodes. These gains primarily arise from the dynamic mapping between market states and portfolio decisions, which allows investors to adjust both exposure and timing as market conditions evolve.

The sensitivity analysis further shows that this state-dependence is the main driver of economic value. While a mean–variance restriction preserves the broad allocation patterns and delivers comparable performance in terms of terminal wealth and opportunity costs, incorporating higher-order moments improves risk-adjusted outcomes and yields more stable and parsimonious strategies, with reduced dispersion in allocation paths and rebalancing activity.

Overall, the results highlight the importance of accounting for state-dependent dynamics when allocating capital in speculative markets. Higher-order moments refine portfolio decisions, but the primary source of performance gains lies in exploiting the state-dependent nature of speculative dynamics.

References

B. Andrews, M. Calder, and R. A. Davis. Maximum likelihood estimation for α-stable autoregressive processes. *The Annals of Statistics*, 37(4):1946–1982, 2009.

F. Blasques, S. J. Koopman, and G. Mingoli. Observation-driven filters for time-series with stochastic trends and mixed causal non-causal dynamics. Technical report, Tinbergen Institute Discussion Paper, 2023.

G. Cavaliere, H. B. Nielsen, and A. Rahbek. Bootstrapping noncausal autoregressions: with applications to explosive bubble modeling. *Journal of Business & Economic Statistics*, 38(1):55–67, 2020.

G. Cubadda, A. Hecq, and E. Voisin. Detecting common bubbles in multivariate mixed causal–noncausal models. *Econometrics*, 11(1):9, 2023.

G. W. Evans. Pitfalls in testing for explosive bubbles in asset prices. *American Economic Review*, 81(4):922–930, 1991.

S. Fries. Conditional moments of noncausal alpha-stable processes and the prediction of bubble crash odds. *Journal of Business & Economic Statistics*, 40(4):1596–1616, 2022.

S. Fries and J.-M. Zakoian. Mixed causal-noncausal ar processes and the modelling of explosive bubbles. *Econometric Theory*, 35(6):1234–1270, 2019.

C. Gourieroux and J. Jasiak. Generalized covariance estimator. *Journal of Business & Economic Statistics*, 41(4):1315–1327, 2023.

C. Gourieroux and A. Monfort. Infrequent extreme risks. *The Geneva Papers on Risk and Insurance Theory*, 29(1):5–22, 2004.

C. Gourieroux and J.-M. Zakoian. Local explosion modelling by non-causal process. *Journal of the Royal Statistical Society Series B*, 79(3):737–756, 2017.

C. Gourieroux, J. Jasiak, and M. Tong. Convolution-based filtering and forecasting: An application to wti crude oil prices. *Journal of Forecasting*, 40(7):1230–1244, 2021.

C. Gouriéroux, Y. Lu, and C.-Y. Robert. The causal-noncausal tail processes: An introduction. *arXiv preprint arXiv:2506.04046*, 2025.

A. Hecq and D. Velasquez-Gaviria. Spectral estimation for mixed causal-noncausal autoregressive models. *Econometric Reviews*, 44(7):939–962, 2025.

A. Hecq and E. Voisin. Predicting bubble bursts in oil prices during the COVID-19 pandemic with mixed causal-noncausal models. *arXiv preprint arXiv:1911.10916*, 2019.

A. Hecq and E. Voisin. Forecasting bubbles with mixed causal-noncausal autoregressive models. *Econometrics and Statistics*, 20:29–45, 2021.

A. Hencic and C. Gouriéroux. Noncausal autoregressive model in application to bitcoin/usd exchange rates. In *Econometrics of risk*, pages 17–40. Springer, 2015.

E. Jondeau and M. Rockinger. Optimal portfolio allocation under higher moments. *European Financial Management*, 12:29 – 55, 01 2006.

E. Jondeau and M. Rockinger. On the importance of time variability in higher moments for asset allocation. *Journal of Financial Econometrics*, 10:84–123, 12 2012. doi: 10.1093/jjfinec/nbr006.

M. Lanne and P. Saikkonen. Noncausal autoregressions for economic time series. *Journal of Time Series Econometrics*, 3(3), 2011.

H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

J. H. McCulloch. Simple consistent estimators of stable distribution parameters. *Communications in Statistics-Simulation and Computation*, 15(4):1109–1136, 1986.

R. C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, 1969.

R. C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971.

P. C. B. Phillips, Y. Wu, and J. Yu. Explosive behavior in the 1990s nasdaq: When did exuberance escalate asset values? *International Economic Review*, 52(1):201–226, 2011.

P. C. B. Phillips, S. Shi, and J. Yu. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the s&p 500. *International Economic Review*, 56(4):1043–1078, 2015.

R. C. Scott and P. A. Horvath. On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, 35(4):915–919, 1980.

Appendix A. Proofs

Appendix A.1. Conditional moments of terminal wealth

This subsection derives the conditional moments of terminal wealth used in the portfolio optimization problem described in Section 2.1. Note that in the empirical implementation in Section 5.1 returns are computed from the observed price and not the detrended component (as in Section 2.1), which leads to the adjusted expressions derived below.

Following Section 5.1, the observed price of the speculative asset is decomposed as

$$Y_t = T(t) + X_t,$$

where $T(t)$ is a deterministic trend and (X_t) is a stationary MAR(1,1) process. The safe asset (S_t) follows a geometric Brownian motion with drift v and volatility ς.

The portfolio optimization hence requires computing

$$\mathbb{E}\big[U(W_{t+H}) \mid Y_t, S_t\big],$$

which depends on the conditional raw moments

$$M_p := \mathbb{E}\big[W_{t+H}^p \mid Y_t, S_t\big], \quad p = 1, 2, 3, 4.$$

Given the independence of (Y_t) and (S_t), conditional expectations factorize with respect to X_t and S_t.

Decomposition of terminal wealth. Recall from Section 2.1 that terminal wealth can be written as

$$W_{t+H} = \frac{S_{t+H}}{S_{t+h}} \left(\omega \frac{Y_{t+h}}{Y_t} + (1-\omega) \frac{S_{t+h}}{S_t} \right).$$

Let

$$A := \frac{S_{t+H}}{S_{t+h}}, \quad C := \frac{Y_{t+h}}{Y_t}, \quad B := \frac{S_{t+h}}{S_t},$$

where C denotes the gross return of the speculative asset and A, B those of the safe asset over non-overlapping intervals. Then

$$W_{t+H} = A\big(\omega C + (1-\omega)B\big).$$

Theorem 1. *Suppose that the processes (Y_t) and (S_t) are independent, and that (S_t) follows a geometric Brownian motion with independent increments. Then, the p-th conditional raw moment of terminal wealth is*

$$M_p = \sum_{j=0}^{p} \binom{p}{j} \omega^j (1-\omega)^{p-j}\, \mathbb{E}\big[A^p \mid S_t\big] \mathbb{E}\big[B^{p-j} \mid S_t\big] \mathbb{E}\big[C^j \mid Y_t\big]. \tag{A.1}$$

Proof. Expanding W_{t+H}^p gives

$$W_{t+H}^p = A^p \sum_{j=0}^{p} \binom{p}{j} \omega^j (1-\omega)^{p-j} C^j B^{p-j}.$$

Taking conditional expectations given (Y_t, S_t) yields

$$M_p = \sum_{j=0}^{p} \binom{p}{j} \omega^j (1-\omega)^{p-j} \mathbb{E}\left[A^p C^j B^{p-j} \,\middle|\, Y_t, S_t\right].$$

Since $A = S_{t+H}/S_{t+h}$ and $B = S_{t+h}/S_t$ depend on non-overlapping increments of the geometric Brownian motion over $(t+h, t+H]$ and $(t, t+h]$, respectively, they are conditionally independent given S_t. Moreover, the independence of (Y_t) and (S_t) implies that C is independent of (A, B) conditional on (Y_t, S_t). Hence

$$\mathbb{E}\left[A^p C^j B^{p-j} \,\middle|\, Y_t, S_t\right] = \mathbb{E}[A^p \mid S_t]\, \mathbb{E}[B^{p-j} \mid S_t]\, \mathbb{E}[C^j \mid Y_t].$$

Substituting into the previous expression gives (A.1).

□

This decomposition highlights that higher-order moments of terminal wealth arise from the interaction between nonlinear return dynamics of the speculative asset and the log-normal structure of the safe asset.

Moments of the speculative asset. Since $T(t)$ is deterministic, conditioning on $Y_t = y$ is equivalent to conditioning on $X_t = x$ with $y = T(t) + x$. Under independence, conditioning on S_t does not affect C and is omitted.

Lemma 1. *For* $p \geq 1$,

$$\mathbb{E}\left[C^p \mid X_t = x\right] = \frac{1}{(T(t)+x)^p} \sum_{j=0}^{p} \binom{p}{j} T(t+h)^{p-j} \mathbb{E}\left[X_{t+h}^j \mid X_t = x\right], \tag{A.2}$$

where the conditional moments $\mathbb{E}[X_{t+h}^j \mid X_t = x]$ *are provided in closed form by Proposition 1.*

Proof. Follows from $Y_{t+h} = T(t+h) + X_{t+h}$ and the binomial expansion. □

This expression shows that the deterministic trend affects not only the mean but also higher-order moments of returns, thereby influencing optimal portfolio decisions through both expected returns and higher-order risk.

Corollary 1. *If* $T \equiv 0$, *then*

$$\mathbb{E}[C^p \mid X_t = x] = \frac{\mathbb{E}[X_{t+h}^p \mid X_t = x]}{x^p}.$$

Moments of the safe asset. Using standard properties of the geometric Brownian motion and the definitions $A = S_{t+H}/S_{t+h}$ and $B = S_{t+h}/S_t$, we obtain

$$\mathbb{E}\left[A^p\right] = \exp\Big(p\upsilon(H-h) + \tfrac{1}{2}p(p-1)\varsigma^2(H-h)\Big),$$
$$\mathbb{E}\left[B^p\right] = \exp\Big(p\upsilon h + \tfrac{1}{2}p(p-1)\varsigma^2 h\Big).$$

These expressions, together with (A.1) and (A.2), yield closed-form formulas for the first four conditional moments of terminal wealth used in the portfolio optimization problem in Section 2.1.

Appendix A.2. Derivation of the constants σ_1^α, β_1, κ_p, and λ_p

This appendix derives the constants $(\sigma_1^\alpha, \beta_1, \kappa_p, \lambda_p)$ entering the conditional moment representation in Proposition 1. These quantities are obtained by specializing the general expressions for two-sided α-stable MA(∞) processes derived in Fries [2022] to the MAR(1,1) case.

Following Fries [2022], consider the stationary component X_t represented as

$$X_t = \sum_{k\in\mathbb{Z}} a_k \varepsilon_{t+k},$$

where $\varepsilon_t \sim S(\alpha, \beta, \sigma, \mu)$ is an i.i.d. α-stable sequence with $0 < \alpha < 2$ and $\alpha \neq 1$. Without loss of generality we set $\mu = 0$. The location parameter only induces deterministic shifts in X_t and X_{t+h} and therefore does not affect the dependence structure driving the conditional moments. Its effect on raw conditional moments can be recovered from the $\mu = 0$ case via the binomial transformation $\mathbb{E}[X_{t+h}^p \mid X_t = x] = \sum_{j=0}^{p} \binom{p}{j} \mu_2^{p-j}\, \mathbb{E}[\tilde{X}_{t+h}^j \mid \tilde{X}_t = x - \mu_1]$, where $\tilde{X}_t := X_t - \mu_1$ and μ_1, μ_2 denote the location shifts induced by μ on X_t and X_{t+h} respectively (see Section 2 in Fries [2022]). Then, for $p \leq 4$, the constants entering the conditional moments stated in Proposition 1 are defined as

$$\begin{aligned}
\sigma_1^\alpha =& \sigma^\alpha \frac{1 - |\varphi^\circ \varphi^\bullet|^\alpha}{(1 - \varphi^\circ \varphi^\bullet)^\alpha (1 - |\varphi^\circ|^\alpha)(1 - |\varphi^\bullet|^\alpha)},\\
\beta_1 =& \beta \frac{1 - \varphi^{\circ<\alpha>} \varphi^{\bullet<\alpha>}}{1 - |\varphi^\circ|^\alpha |\varphi^\bullet|^\alpha} \frac{(1 - |\varphi^\circ|^\alpha)(1 - |\varphi^\bullet|^\alpha)}{(1 - \varphi^{\circ<\alpha>})(1 - \varphi^{\bullet<\alpha>})},\\
\kappa_p =& \frac{\varphi^{\bullet hp}(1 - |\varphi^\circ|^\alpha) + (\varphi^{\circ -hp} |\varphi^\circ|^{h\alpha})(1 - |\varphi^\bullet|^\alpha)}{1 - |\varphi^\circ \varphi^\bullet|^\alpha}\\
&+ \frac{(\varphi^{\bullet hp} |\varphi^\circ|^\alpha (\varphi^\bullet \varphi^\circ)^{-p} - \varphi^{\circ -hp} |\varphi^\circ|^{\alpha h})(1 - |\varphi^\circ|^\alpha)(1 - |\varphi^\bullet|^\alpha)}{(1 - |\varphi^\circ|^\alpha (\varphi^\bullet \varphi^\circ)^{-p})(1 - |\varphi^\circ \varphi^\bullet|^\alpha)},\\
\lambda_p =& \beta_1 \times \left(\frac{\varphi^{\bullet hp}(1 - |\varphi^\circ|^{<\alpha>}) + (\varphi^{\circ -hp} \varphi^{\circ<\alpha>h})(1 - \varphi^{\bullet<\alpha>})}{1 - \varphi^{\circ<\alpha>} \varphi^{\bullet<\alpha>}} \right.\\
&\left. + \frac{(1 - \varphi^{\circ<\alpha>})(1 - \varphi^{\bullet<\alpha>})}{1 - \varphi^{\circ<\alpha>} (\varphi^\bullet \varphi^\circ)^{-p}} \frac{(\varphi^{\bullet hp} \varphi^{\circ<\alpha>} (\varphi^\bullet \varphi^\circ)^{-p} - \varphi^{\circ -hp} \varphi^{\circ<\alpha>h})}{1 - \varphi^{\circ<\alpha>} \varphi^{\bullet<\alpha>}} \right),
\end{aligned}$$

where, following Fries [2022], $y^{\langle\alpha\rangle} = \text{sign}(y)|y|^\alpha$ for any $y \in \mathbb{R}$, and $\varphi^\bullet \neq 0$.[11] σ_1 and β_1 denote

[11] The case $\varphi^\bullet = 0$ corresponds to a purely noncausal AR(1) process and must be treated separately.

the scale and asymmetry parameters of the marginal α-stable distribution of X_t, whereas the constants κ_p and λ_p, $p > 2$, generalize standard dependence measures invoked in the literature to powers of X_t and X_{t+h} in the asymmetric case. At the same time, $\mathcal{H}$ contains functions related to the marginal density of the stable random variable X_t and for $n \in \mathbb{N}$, $\boldsymbol{\theta} = (\theta_1, \theta_2) \in \mathbb{R}^2$, $x \in \mathbb{R}$ is defined as

$$\mathcal{H}(n, \theta; x) = \int_0^{+\infty} e^{-\sigma_1^\alpha u^\alpha} u^{n(\alpha-1)} \big(\theta_1 \cos(ux - a\beta_1\sigma_1^\alpha u^\alpha) + \theta_2 \sin(ux - a\beta_1\sigma_1^\alpha u^\alpha)\big) du,$$

where $a := \tan\big(\frac{\pi\alpha}{2}\big)$. The quantities σ_1^α and β_1 respectively characterize the effective scale and skewness of the process, while κ_p and λ_p capture the propagation of shocks across leads and lags and govern higher-order dependence and asymmetry.

Substituting the closed-form coefficients of the two-sided MA(∞) representation of MAR(1,1) process

$$a_k = \frac{(\varphi^\circ)^k}{1 - \varphi^\circ\varphi^\bullet}, \quad k \geq 0, \qquad a_k = \frac{(\varphi^\bullet)^{-k}}{1 - \varphi^\circ\varphi^\bullet}, \quad k < 0,$$

into the general expressions given in Fries [2022], and evaluating the resulting sums using standard geometric series arguments, yields the closed-form expressions reported above. Note that $(1 - \varphi^\circ\varphi^\bullet)^\alpha$ is well defined for non-integer α since $\varphi^\circ\varphi^\bullet \in (-1, 1)$ implies $1 - \varphi^\circ\varphi^\bullet > 0$.

Appendix B. Utility Functions

This subsection describes the utility functions used in the portfolio optimization problem. We begin with the standard Constant Relative Risk Aversion (CRRA) specification, and then introduce the class of LIRA (Left Integrable Risk Aversion) utility functions.

In the CRRA case, preferences are represented by

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,\ \gamma \neq 1,$$

where γ denotes the coefficient of relative risk aversion. While this specification is widely used in portfolio choice, it may generate degenerate optimal demands when the return distribution exhibits infrequent but extreme losses.

To address this limitation, we consider the class of LIRA utility functions introduced by Gourieroux and Monfort [2004]. These preferences provide a flexible framework to incorporate higher-order risk considerations while preserving tractability, and are particularly well suited to environments with asymmetric and heavy-tailed risks, such as those generated by MAR dynamics.

Definition. A utility function U is said to be LIRA if it is increasing, concave, and satisfies

$$\lim_{W \to -\infty} U'(W) < +\infty.$$

This condition ensures that marginal utility remains bounded for large losses, thereby avoiding the degeneracy of optimal portfolio choices under CRRA preferences. Any LIRA utility function admits the representation

$$U(W) = aW + b \int_{-\infty}^{W} S(x)\, dx + c, \qquad a \geq 0,\ b > 0, \tag{B.1}$$

where S is a non-increasing function taking values in $[0, 1]$, interpreted as the survivor function of a probability distribution. This representation decomposes utility into a linear component and a downside-risk component related to a put-option payoff.

LIRA preferences yield finite and non-degenerate optimal demands for assets exhibiting rare and severe losses. The ratio a/b controls the relative importance of the linear (risk-neutral) and nonlinear (risk-averse) components, while the shape of S determines sensitivity to downside risk.

Parametric specifications. We consider four parametric specifications corresponding to different choices of S:

- *Gaussian LIRA (LIRA-G): S* is based on the Gaussian distribution, yielding a smooth utility with decreasing absolute risk aversion.
- *Double-exponential LIRA (LIRA-DE): S* is based on a Laplace distribution, allowing for asymmetric sensitivity to gains and losses.
- *Uniform LIRA (LIRA-U): S* corresponds to a uniform distribution on $[u_l, u_h]$, with $u_l < u_h$, leading to a piecewise quadratic utility locally depending only on mean and variance, and which is affine outside this interval.

- *Power-weighted LIRA (LIRA-P):* S is defined on $[u_l, u_h]$ as

$$S(W) = \begin{cases} 1, & W < u_l, \\ \left(\dfrac{u_h - W}{u_h - u_l} \right)^{\alpha_s}, & u_l \leq W \leq u_h, \\ 0, & W > u_h, \end{cases}$$

  where $\alpha_s > 0$ controls the curvature. This specification generalizes the uniform case (obtained when $\alpha_s = 1$) and allows for flexible modeling of downside sensitivity.

For both CRRA and LIRA preferences, expected utility is approximated using a fourth-order Taylor expansion around $\overline{W} = \mathbb{E}[W_{t+H} \mid X_t, S_t]$. The resulting objective depends on the conditional moments of terminal wealth derived in Appendix A.1.