Bet on a bubble asset ? An optimal portfolio allocation strategy

Gilles de Truchis*

Department of Economics (LEO), University of Orléans

and

Elena-Ivona Dumitrescu

Department of Economics (CRED), University of Paris-Panthéon-Assas

and

Sébastien Fries

and

Arthur Thomas

Department of Economics (LEDa), Paris Dauphine University - PSL

February 5, 2024

Abstract

We discuss portfolio allocation when one asset exhibits phases of locally explosive behavior. We model the conditional distribution of such an asset through mixed causal-non-causal models which mimic well the speculative bubble behaviour. Relying on a Taylor-series-expansion of a CRRA utility function approach, the optimal portfolio(s) is(are) located on the mean-variance-skewness-kurtosis efficient surface. We analytically derive these four conditional moments and show in a Monte-Carlo simulation exercise that incorporating them into a two-assets portfolio optimization problem leads to substantial improvement in the asset allocation strategy. All performance evaluation metrics support the higher out-of-sample performance of our investment strategies over standard benchmarks such as the mean-variance and equally-weighted portfolio. An empirical application on three bubble hedging portfolios that rely on CO2, Brent, and WTI price series respectively as the speculative assets confirms these findings.

Keywords: non-causal process, α-stable, asset allocation, utility function, out-of-sample performance

*We are grateful to Jean-Michel Zakoian, Christian Francq, Olivier Scaillet, Christophe Pérignon, Serge Darolles, Gaëlle le Fol and David Ardia. We also thank the seminar participants at CREST-ENSAE, Dauphine University - PSL, University of Paris Nanterre, as well as the participants at the 4th 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.

1 Introduction

The question of portfolio allocation with bubble assets is a highly relevant empirical question nowadays in the context of the emergence of private pension systems, which increases competition within fund management industry by pushing more and more individuals to choosing among funds with different characteristics, in terms of realized portfolio performance, for example.

A close look at the dynamics of various asset prices, that are sometimes called speculative assets, reveals the presence of phases of locally explosive behaviors, i.e. increasing patterns followed by a burst. Called rational asset pricing bubbles when due to rational deviations from the fundamental value (see Blanchard and Watson, 1982; Tirole, 1985), these phenomena have been detected more and more accurately in the financial markets across the world together with the more traditional properties of heavy-tailed marginal distributions and volatility clustering.

A rich theoretical literature has been focusing on two aspects of this phenomenon: the investment problem and the financial economic implications, (see e.g. Davis and Lleo, 2013). To explain how a bubble originates in the market, researchers generally relied on standard martingale theory of bubbles (Biagini et al., 2014; Jarrow et al., 2010; Protter, 2012), or added further assumptions such as portfolio constraints or defaultable claims, (see Biagini and Nedelcu, 2015; Hugonnier, 2012; Jarrow et al., 2012). The impact of bubbles on economic growth (Martin and Ventura, 2012; Carvalho et al., 2012) or on unemployment (Hashimoto and Im, 2016, 2019; Miao et al., 2016) has also been scrutinized recently. But this phenomenon has not been thoroughly gauged so far through the lens of portfolio allocation, although optimal portfolio selection has been a major topic in finance since the works of Markowitz (1952).

The scarcity of this literature may be explained by the distributional specificities of bubble asset prices and the risk they incur although, from a financial perspective, investors are certainly interested in constructing portfolios hedging bubble burst risk. Indeed, traditional portfolio theory is consistent with expected utility and its von Neumann-Morgenstern axioms of choice when either asset returns are normally distributed (i.e., higher moments are irrelevant), or investors have a quadratic utility function (see e.g. Samuelson, 1967). But these assumptions were shown not to be empirically justified (see Mandelbrot, 1963; Ang et al., 2006; Massacci, 2017; Ingersoll, 1975; Scott and Horvath, 1980, among others).

This lead researchers and practitioners to intensively work on new portfolio allocation strategies, which, among others, pay attention to higher order moments, namely asymmetry and fat-tailness (see Briec et al., 2013; Kolm et al., 2014, for literature reviews). A wide variety of approaches have been proposed in this literature: Taylor expansion of the expected utility (Jondeau and Rockinger, 2006, 2012; Guidolin and Timmermann, 2008; Martellini and Ziemann, 2010), Gram-Charlier expansion of downside risk measures (Favre and Galeano, 2002; León and Moreno, 2017; Zoia et al., 2018; Lassance and Vrins, 2021), the shortage function of (Briec et al., 2007, 2013, to name but a few).

A portfolio strategy based on an accurate characterization of the mechanism generating financial bubbles seems necessary to avoid misleading outcomes. However, neither ARMA nor (G)ARCH / stochastic volatility models, traditionally used to characterize the predictive distribution of returns, are able to mimic such bubble behaviours. To our knowledge, only the paper by Ghahtarani (2021) discusses portfolio allocation in bubble conditions. The author introduces a new portfolio risk measure and shows that it can perform better than classical risk measures in bubble situations. To this aim, he uses a fuzzy neural network model to compute scenario paths of end-horizon market value. But the approach is not specifically designed for portfolio allocation, it operates in multiple steps and requires accounting for uncertainty surrounding the fundamental and market value predictions.

We contribute to this literature by documenting the attractiveness of portfolio strategies that account explicitly for the distributional characteristics of bubble assets. More precisely, we exploit very recent theoretical results on non-causal models to appropriately characterize the conditional distribution of asset prices exhibiting bubble behaviour. Indeed, non-causal autoregressive processes with stable distributed errors appear to be fit to model speculative financial bubbles as they mimic well locally explosive patterns (see e.g. Gourieroux and Zakoian, 2017).

Our approach is anchored in the classical theoretical rational-expectations bubble framework proposed by Blanchard and Watson (1982). A bubble occurs when prices temporarily deviate from the fundamental value. But if Blanchard and Watson's model features successive bubble/burst cycles, the non-causal model may generate more realistic price dynamics where bubble events intersperse calmer periods. Besides, the gradual collapse in the dynamics of mixed causal-noncausal model (hereafter MAR) reconciles the rational expectations bubbles with regular variation tail indexes above 1, a well-documented statistical property of financial data, (see Lux and Sornette, 2002).[1] Most importantly, they exhibit surprising features such as a predictive distribution with lighter tails than the marginal distribution, which allows one to obtain predictions of higher-moments that are expected to be of crucial importance for the (non-)investment decision. Indeed, this framework relaxes the finite variance constraint while insuring the stationarity of the process, (see Gourieroux et al., 2020, on the existence of multiple stationary nonlinear equilibria in bubble models).

By relying on the results of Fries (2021), we derive the first four conditional moments of an α-stable MAR(1,1) process and show that incorporating them into a two-assets portfolio optimization problem can lead to substantial improvement in the asset allocation strategy. For this, we consider the standard Taylor-series-expansion of a CRRA utility function approach à la Jondeau and Rockinger (2006, 2012), (see also Martellini and Ziemann, 2010). The optimal portfolio(s) is(are) located on the mean-variance-skewness-kurtosis efficient surface in the sense that no other portfolio can dominate it on all four moments. But since there is evidence that standard utility functions are locally quadratic and higher-order moments may not significantly impact portfolio selection (see e.g. Markowitz, 2014), we also consider, as a robustness check, a polynomial-goal-

[1] We defer the reader to Fries (2021) for further discussion on the link between non-causal models and rational bubbles à la Blanchard and Watson (1982).

programming approach to find a portfolio on the higher-moment efficient surface without the need to specify a utility function. An advantage of our bubble portfolio optimisation approach (hereafter BP) over the benchmarks is that the optimal strategies take the form of couples – investment share and investment horizon –. The endogenous character of the latter leads to fewer rebalancings over the global investment horizon and simplifies portfolio management operations relative to the standard daily rebalancing approach. The economic value of our strategy is compared with standard benchmarks such as the mean-variance and equally-weighted portfolios.

In contrast to Ghahtarani (2021), if his machine learning framework were to be used for portfolio allocation, our approach is not scenario-dependent. Besides, the non-causal framework also presents the advantage of ease of interpretability, in the sense that the solution(s) of the portfolio allocation problem can be traced back to the conditioning value of the bubble asset dynamics and the higher-order moments of the conditional return distribution.

A set of Monte-Carlo simulations emphasizes the reliability of our BP approach. As shown in the Web-Appendix, the portfolio strategies based on estimates of the MAR(1,1) parameters are clustered around the theoretical optimal ones, i.e. based on the true parameters, and their dispersion reduces quickly as the sample size increases. This indicates that estimation uncertainty does not affect much the portfolio allocation problem. Note, however, that the starting value of the speculative asset $X_t = x$ matters a lot in the selection of the optimal investment share and horizon, which is not the case of the no-bubble asset. This is expected to lead to investment strategies that outperform standard benchmarks such as the mean-variance and equally-weighted portfolios. We study this intuition by simulation and find that both terminal wealth and opportunity cost performance evaluation measures support the superiority of our portfolio allocation strategy in out-of-sample. The difference is particularly significant when the conditioning values $X_t = x$ are in the tails of the marginal distribution of the process, i.e. when the first asset is undergoing a bubble episode, which is of outmost importance for the investor.

An empirical application on three bubble hedging portfolios that use the CO2, the Brent and the WTI prices respectively, as speculative assets confirms these simulation-based results. As a preliminary step to select candidate assets, we test for the presence of bubbles in asset price dynamics by relying on the recent generalized-sup ADF test of Phillips et al. (2015) that is appropriate for rational bubble frameworks, among others.[2] The pseudo-out-of-sample performance of our allocation approach is then compared to that of the two benchmark models for each of the three hedging porfolios considered. A battery of robustness checks is performed and all findings support the superiority of our BP approach to the benchmarks whatever the investor's allocation program and preferences specification. A more realistic setting that accounts for transaction costs is also discussed as well as the case where portfolio optimization relies on the first two moments of the price distribution. The latter setup highlights the superiority of using conditional moments in the portfolio allocation optimisation process rather than unconditional ones.

The paper is structured as follows. In Section 2 we introduce the proposed allocation problem. In Section 3 we conduct a simulation-based out-of-sample horce-race with standard benchmarks to evaluate the relative economic value of our approach, while Section 4 details the empirical application. Finally, Section 5 concludes and the Appendix gathers all proofs.

2 Bubble-riding allocation problem

Hedging bubble asset risk is nowadays a particularly important issue for an investor handling speculative assets. In this section, we provide a unified framework to solve the allocation problem in presence of an asset exhibiting a bubble behaviour. First, we formally introduce the portfolio allocation problem and then provide the necessary quantities to compute the conditional moments of portfolio return distribution when the speculative asset price is modeled as a mixed causal-noncausal process. Finally, we briefly review the methods that will be used to evaluate the economic value of the optimal portfolio strategies.

[2]Early tests for rational bubbles relied on Shiller (1981)'s variance bounds test, West (1987, 1988)'s two step procedure or cointegration tests, but these approaches are subjected to multiple issues.

2.1 Optimal portfolio allocation

We investigate the asset allocation problem in the context with a speculative asset price X_t, for which the dynamics of higher order conditional moments is of particular importance, and a bubble-free one, S_t. Two approaches that account for higher-order moments in the choice of the optimal portfolio have gained investors' attention to date and are considered in our analysis. The first is based on a Taylor expansion of the expected utility function, while the latter consists in the Polynomial Goal Programming (PGP) model. We privilege the CRRA utility function because it is probably economically the most relevant preference family, as it realistically assumes that risk aversion is relatively constant over wealth levels, (see also Jondeau and Rockinger, 2006, 2012, and references therein). A complementary analysis, based on the PGP approach, is available in the Web-Appendix (see De Athayde and Flôres Jr, 2004, for arguments in favour of this approach).

We consider an investor endowed with wealth W_t at present date t, who allocates her portfolio constituted of these two assets to maximize the expected utility $U(W)$ over her end-of-period wealth W_{t+H}. The initial wealth is innocuous to the optimization problem and arbitrarily set to one. The investor has an investment horizon H: at date t, she will decide of the share ω (resp. $1-\omega$) to invest in the speculative asset (resp. bubble-free asset), and of the intermediate horizon $h \leq H$ at which she commits to liquidate its holding of speculative asset and to invest the proceedings in the bubble-free asset until $t+H$. Short selling is allowed, hence portfolio weights can take both positive and negative values. This leads to an optimization problem of the terminal wealth W_{t+H} or, equivalently, of the overall return $R_{t+H} = (W_{t+H} - W_t)/W_t$ in both the allocation ω and the intermediate horizon h, which is new in the financial literature.

We assume that the speculative asset's price X_t follows a mixed causal-noncausal stable AR process, i.e. MAR(1,1), with a non-zero location parameter. This choice is motivated by the recent econometric literature that proved non-causal models to be a convenient way to model locally explosive phenomena such as speculative bubbles, while featuring heavy-tailed marginals and conditional heteroscedastic effects generally encountered in financial data (see e.g. Cavaliere et al., 2020; Fries and Zakoian, 2019; Gourieroux and Jasiak, 2018; Gourieroux and Zakoian, 2017). The bubble-free asset is assumed to follow a Geometric Brownian Motion (GBM) dynamics with drift υ and volatility ς. The price processes (X_t) and (S_t) will be assumed independent, which provides a nice framework for hedging purposes.[3]

For a given strategy (ω, h), the terminal wealth can be expressed as

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

or alternatively, in terms of returns, $W_{t+H} = 1 + R_{t+H}$, where the terminal portfolio return R_{t+H} writes

$$R_{t+H} = \Big(1 + r^S_{t+h,t+H}\Big)\Big(\omega r^X_{t,t+h} + (1-\omega) r^S_{t,t+h} + 1\Big) - 1,$$

with $r^S_{t+h,t+H} := S_{t+H}/S_{t+h} - 1$, $r^S_{t,t+h} := S_{t+h}/S_t - 1$, and $r^X_{t,t+h} := X_{t+h}/X_t - 1$, the asset's returns in-between the key investment events.

In this framework, we follow Jondeau and Rockinger (2006, 2012) to approximate the allocation problem.[4] The CRRA utility maximization program of the fourth order Taylor approximation around the expected terminal wealth is

$$\max_{(\omega,h)} \mathbb{E}[U(W_{t+H}|X_t, S_t)] \approx \sum_{k=0}^{4} \frac{U^{(k)}(\overline{W}_{t+H})}{k!}\mathbb{E}\Big[(W_{t+H} - \overline{W}_{t+H})^k | X_t, S_t\Big], \tag{2}$$

with $U(c) = c^{1-\gamma}/(1-\gamma)$ for a risk aversion parameter $\gamma > 0$ and $\overline{W}_{t+H} = \mathbb{E}\Big[W_{t+H}|X_t, S_t\Big]$.

[3] 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.

[4] Lhabitant et al. (1998) has shown that the infinite Taylor series expansion converges to the expected utility in the CRRA case for wealth levels between 0 and $2\overline{W}$ that appear to be large enough for stocks and bonds regardless of the degree of non-normality, in particular when short-selling is prohibited.

The investor's preference (or aversion) toward the k^{th} moment is directly given by the k^{th} derivative of the utility function. The effects of the third and fourth moments on the approximated expected utility are positive and negative, respectively, and correspond to financial theory (see Scott and Horvath, 1980). The expected utility also depends on the central conditional moments of the distribution of terminal wealth, which can be expressed in terms of conditional moments of the portfolio return distribution as $\mathbb{E}\Big[(W_{t+H} - \overline{W}_{t+H})^k | X_t, S_t\Big] = \mathbb{E}\Big[(R_{t+H} - \overline{R}_{t+H})^k | X_t, S_t\Big]$, since $\overline{W}_{t+H} = 1 + \overline{R}_{t+H}$ with $\overline{R}_{t+H} := \mathbb{E}[R_{t+H}|X_t, S_t]$. It is just a matter of algebra using the independence between (X_t) and (S_t) to express the objective functions in terms of the conditional moments of the speculative asset price, $\mathbb{E}\Big[X_{t+h}^p | X_t\Big]$, $p = 1, 2, 3, 4$, that are detailed in the next subsection, and the parameters (see 6.1 for further computational details).

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

In this subsection we discuss the existence and derivation of the first four conditional moments of the speculative asset price. As the econometric literature has identified MAR processes to be appropriate for financial bubble modelling, we rely on them, (see e.g. Hecq and Voisin, 2021, and references therein).[5]

Let (X_t) be the α-stable solution of the MAR(1,1) process $X_t = \varphi^{\circ} X_{t+1} + \varphi^{\bullet} X_{t-1} + \varepsilon_t$, with i.i.d. α-stable errors, $\varepsilon_t \overset{i.i.d.}{\sim} \mathcal{S}(\alpha, \beta, \sigma, \mu)$ and $\alpha \neq 1$ (for simplicity), $\beta \in [-1, 1]$, and $\sigma > 0$. The process is well defined and strictly stationary for $|\varphi^{\circ}| < 1$, $|\varphi^{\bullet}| < 1$, and $\varphi^{\circ} \neq \varphi^{\bullet}$. It then has a $MA(\infty)$ representation $X_t = \sum_{k\in\mathbb{Z}} a_k \varepsilon_{t+k}$, whose coefficients satisfy $\sum\limits_{k\in\mathbb{Z}} |a_k|^s < +\infty$ for some $s \in (0, \alpha) \cap [0, 1]$. Without loss of generality, in the following we assume that the shift μ is null, but in practice we handle the possibility of $\mu \neq 0$ by relying on a simple transformation of the conditional moments obtained with zero location parameter to those associated with a non-null shift (see Section 2 in Fries, 2021).

[5]More generally, any MARMA model could be used, but this would engender an approximation cost related to the truncation of the two-sided infinite sum of $MA(\infty)$ coefficients (see Fries, 2021). We prefer a more parsimonious approach for which we can obtain the formulas for the coefficients in closed form.

Now let $\mathbf{X}_t = (X_t, X_{t+h})$ denote the bivariate stable vector obtained from X_t for horizon $h \geq 1$. Proposition 3.1 $i)$ in Fries (2021) then applies and states the condition of existence of higher-order conditional power moments, although the marginal variance of the process X_t is infinite. In particular, the conditional moments up to integer order p, $\mathbb{E}[|X_{t+h}|^p|X_t]$, may exist as long as $\nu \geq 0$ exists such that $\sum_{k\in\mathbb{Z}}(a_k^2 + a_{k-h}^2)^{\frac{\alpha+\nu}{2}}|a_k|^{-\nu} < +\infty$ and $0 \leq p < \min(\alpha+\nu, 2\alpha+1)$, where $(a_k,\ a_{k-h})$ are the coefficients of the infinite moving average representation of the process. The more *anticipative, i.e. noncausal* the process, the larger $\nu \geq 0$, which insures the existence of all conditional moments up to order $2\alpha+1$ at all prediction horizons when (a_k) decays geometrically or hyperbolically for example.

Proposition 1 *For $\alpha \neq 1$, the moments $\mathbb{E}[X_{t+h}^p|X_t]$, $p \leq 4$, when they exist, are given by Theorems 2.1 and 2.2 in Fries (2021) as a function of four quantities, σ_1^α, β_1, κ_p, and λ_p and a family of functions $\mathcal{H}$. We demonstrate that in the case of a MAR(1,1) process these constants are equal to*

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

where $y^{<\alpha>} = \text{sign}(y)|y|^\alpha$ for any $y \in \mathbb{R}$. σ_1 and β_1 denote the scale and asymmetry parameters of the marginal 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*

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

Proof 1 *See 6.2*

Remark 1 *The conditional moments can be easily computed for $p \leq 4$ and $h \geq 1$ once the functions $\mathcal{H}(n, \theta; x)$ are evaluated for $n = 2, 3, 4$ by following the approach discussed by Fries (2021) and originally proposed by Samorodnitsky et al. (1996) for the conditional expectation.*

Remark 2 *The asymptotic expressions for the conditional moments with respect to the conditioning variable, i.e. when X_t becomes large, given in Proposition 2.1 of Fries (2021) and stated below, remain valid in the MAR(1,1) case when σ_1^α, β_1, κ_p, and λ_p are replaced by the expressions given in Proposition 1 above. To be more precise, if the conditional moment of order p of a bivariate α-stable vector exists and $|\beta_1| \neq 1$, then,*

$$x^{-p}\mathbb{E}[X_{t+h}^p | X_t = x] \underset{x\to\infty}{\to} \frac{\kappa_p + \lambda_p}{1+\beta_1}, \quad x^{-p}\mathbb{E}[X_{t+h}^p | X_t = x] \underset{x\to-\infty}{\to} \frac{\kappa_p - \lambda_p}{1-\beta_1}. \tag{3}$$

3 Economic Value

In this section, we illustrate the usefulness of the BP approach to provide high-performing portfolio allocation strategies. As shown in Section 3 of the Web-Appendix, while the starting value of S_t does not matter, that of X_t deeply modifies the investment landscape. The optimal investment strategies vary according to the conditioning values of the marginal distribution of the process. For this reason, they are expected to outperform standard mean-variance and equally-weighted portfolios, that cannot take into account the current state of the nature at the moment of investing.

We study this intuition in a Monte Carlo setup which has similarities with the empirical application. More precisely, we generate 1000 trajectories of $N = 350$ observations (250 in-sample and 100 out-of-sample) from the MAR(1,1) process $(1 - 0.9F)(1 - 0.1B)X_t = \varepsilon_t$ where $\varepsilon_t \overset{i.i.d.}{\sim} S(1.7, 0.3, 1.5, 15)$[6], which is inspired by the Brent dataset. For each trajectory, we use the first two thirds of the data, $\{1, 2, \ldots, T\}$, labeled as in-sample, to estimate the conditional moments of returns and identify the optimal investment strategies in the form of couples (ω, h) for conditioning values covering the whole marginal distribution of X_t implied by the DGP. The remaining one third of the data, $\{T + 1, \ldots, T + k, \ldots N\}$, labeled as out-of-sample, is used as conditioning values for a new investment. Said otherwise, we assume the investor wishes to invest his wealth in the two assets at a certain date, say $T + k$, within the out-of-sample period. To select the optimal share of the bubble asset in the portfolio, ω, and the duration of this risky investment, h, out of the $H = 26$ periods of the overall investment, she searches for the closest quantile of the theoretical distribution of X_t just below the actual conditioning price at the selected date. The couple(s) (ω, h) estimated in-sample for this quantile by using the *CRRA* utility function with $\gamma = 10$ will then be used to construct the portfolio strateg(y/ies). For each strategy, the portfolio is rebalanced once, at period $T + k + h$. Consisting only in an investment in the no-bubble asset, it is then held constant up until date $T + k + H$.

For comparison reasons, we compute also the mean-variance and the equally-weighted portfolios over the same periods. In the case of the *MV* benchmark portfolio, we use the in-sample data to estimate the optimal investment share in the bubble asset. Then, we use it to construct a buy and hold strategy over H periods for each out-of-sample starting date $T + k$. Finally, the computation of the *EW* portfolio for the same investment horizon is immediate.

To compare the economic value of these strategies we rely on the terminal wealth and on the opportunity cost (also known as performance fee, hereafter OC), (see e.g. Jondeau and Rockinger,

[6] This choice of parameters make the process satisfy the rational bubble condition, i.e. $\left[(\varphi^{\circ})^{\alpha-1} + \varphi^{\bullet}(1 - (\varphi^{\circ})^{\alpha})\right]^{-1} < 1$, (see Remark 4.1 in Fries, 2021)

2006, 2012; González-Pedraz et al., 2015). More precisely, we check whether our approach performs better that the two traditional benchmarks, the equally-weighted portfolio (EW) and the standard mean-variance (MV) portfolio, by testing the equality of the medians (across simulations) of the terminal wealth of the strategies over the out-of-sample. At the same time, the OC corresponds to the amount that needs to be added to the return of a competing benchmark strategy so that the investor becomes indifferent to the portfolio decision based on our framework. As our approach may lead to multiple optimal strategies for a given conditioning value, we report three statistics: the median one, labeled BP_{med}, the one leading to the highest terminal wealth, labeled BP_{sup}, and the one leading to the lowest terminal wealth, labeled BP_{inf}.

Table 1: Relative performance of portfolio strategies

				Strategy			
		BP_{med}	BP_{inf}	BP_{sup}	EW	EW	
Wealth	μ	$1.019^{*/*}$	1.009	$1.069^{*/*}$	1.003	1.002	
	σ	0.038	0.061	0.058	0.052	0.006	
			EW vs			MV vs	
		BP_{med}	BP_{inf}	BP_{sup}	BP_{med}	BP_{inf}	BP_{sup}
OC	μ	0.023	-0.006	0.036	0.001	0.005	0.006
	σ	0.057	0.072	0.056	0.036	0.046	0.047

Notes: Our MAR(1,1)-based strategies are compared with the equally weighted (*EW*) and mean-variance (*MV*) ones in terms of terminal wealth. The opportunity cost (*OC*) relatively to the two benchmark portfolios is also provided. The results take the form of out-of-sample average and standard deviation over the 1000 simulations. Asteriks $(*/*)$ indicate the rejection of the null hypothesis of Wilcoxon's test of equality of medians at the 95% level relatively to each of the two benchmark strategies, *EW* and *MV*, respectively.

Table 1 reports the average, μ, and standard deviation, σ, of the median terminal wealth of each of the five portfolio strategies over the 1000 simulated out-of-sample trajectories. Asteriks $(*/*)$ associated with the estimated μ of each of our strategies indicate that the average of the median terminal wealth is statistically different from that of the (EW/MV) portfolios according to Wilcoxon's test. The average of the median terminal wealth for the three MAR(1,1)-based portfolios is similar and always well above that of the benchmark portfolios. Wilcoxon's test always rejects the null of equal averages, suggesting that our approach performs best in terms of terminal wealth. The positive averages of the opportunity cost also support these findings. A smaller amount needs to be added to the MV strategy than to the EW one to provide the same expected utility as our BP strategies. As our approach is specifically designed for investors that

Table 2: Relative performance of portfolio strategies in positive bubble period

		Strategy					
		BP_{med}	BP_{inf}	BP_{sup}	EW	MV	
Wealth	μ	$1.045^{*/*}$	$1.042^{*/*}$	$1.048^{*/*}$	1.031	1.019	
	σ	0.047	0.046	0.048	0.030	0.020	
		vs EW			vs MV		
		BP_{med}	BP_{inf}	BP_{sup}	BP_{med}	BP_{inf}	BP_{sup}
OC	μ	0.028	0.028	0.028	0.015	0.015	0.015
	σ	0.052	0.052	0.052	0.013	0.013	0.013

Notes: see note to Table 1. The results are based only on the cases where the investment is performed while the first asset exhibits a bubble period, i.e. $X_{T+k} = x$ is beyond the 95% quantile of the theoretical distribution of the price process.

wish to take advantage of bubble periods, in Table 2 we focus on this setup. We assume that one invests only when the unconditional price process seems to exhibit a locally explosive behaviour, i.e. the conditioning values $X_{T+k} = x$ are above the 95% quantile of the theoretical distribution of the process. The results are similar to the previous case, but with a much more concentrated terminal wealth across quantiles, which is particularly favorable to risk averse investors, i.e. BP_{inf} is much higher than when the whole distribution is considered.

All in all, these results indicate that our method may prove useful for the investor that includes a bubble asset in her portfolio. They hold when investigating the case of negative bubbles, i.e. looking only at conditioning values beyond the 95% quantile of the theoretical distribution of the process. Finally, they are qualitatively similar when fixing the risk-aversion parameter γ to 5 or when we rely on the PGP framework (see Section 3.1 in the Web-Appendix).

As a robustness check, in a subsequent Subsection of the Web-Appendix we also investigate the performance of the BP portfolio strategies when relying only on the first two conditional moments. The main difference here with respect to the benchmarks is that our portfolio optimisation approach relies on the conditional moments whereas the EW and MV ones use the unconditional ones. Overall, this version of our portfolio approach performs better than the benchmarks, but it is nevertheless dominated by the BP strategies that rely on all four conditional moments, which is consistent with the DGP used. It hence emphasizes the advantage of using conditional moments in the portfolio allocation optimisation process rather than unconditional ones.

4 Empirical application

4.1 Data and in-sample analysis

This section discusses the performance of the proposed portfolio allocation strategy on real data. We consider three candidates for the bubble asset: the price of CO2 as well as Brent and WTI oil prices. The oil prices are long established in the non-causal econometric literature as having a mixed-causal dynamics, while evidence of explosive behaviour in the CO2 price has been provided by Friedrich et al. (2019). At the same time, we use *EUREX-Euro bund* settlement price series as the bubble-free asset on the European market and the *30-year US Treasury-bond* settlement price for the US market, i.e. in the case of the WTI series. Weekly data, ranging from 2017-01 to 2023-03, have been obtained from Refinitiv Eikon for all series and splitted into an in-sample part (2017-01 to 2021-05) and an out-of-sample one (2021-06 to 2023-03).

The results of the generalized-sup ADF test of Phillips et al. (2015) in Table 3 clearly indicate the presence of a locally explosive behaviour in the oil prices, while the CO2 is on the boundary of the significance level as the series exhibits a more pronounced bubble behaviour in the out-of-sample part.

Table 3: GS-ADF Test

		Finite Sample Critical Value		
	Test Stat.	90%	95%	99%
EU ETS CO2 price	1.481	1.489	1.711	2.351
Brent	2.071			
WTI	2.505			

Notes: Generalized-sup ADF test for the presence of multiple bubbles developed by Phillips et al. (2015). The critical values are based on 1000 simulations.

Existing econometric literature has also emphasized a certain level of deterministic nonstationarity in these series that needs to be tackled. In particular, a trending time varying fundamental part must be extracted before estimating Mixed AR models on the in-sample data.[7] Two detrending methods have gained interest in this literature. A polynomial function has been used by Hencic and Gouriéroux (2015) and by Hecq and Voisin (2019), while Hecq and Voisin (2021) use the Hodrick - Prescott filter. As the first approach is more direct to implement in the out-of-sample, we follow Hecq and Voisin (2019) and use a polynomial trend of order four to capture trending patterns in the Brent series. Similarly, a polynomial trend of order four seems to fit the CO2 series, while a polynomial trend of order three better captures the trend of the WTI series. The bubblish behaviour of the detrended series is clearly distinguishable in Figure 1, the gray region corresponding to the out-of-sample and the dense vertical grid to two-week intervals.

We then rely on the procedure of Lanne and Saikkonen (2011) based on the AIC information criterion to perform model selection on causal-non-causal models and identify the MAR(1,1) as the best specification for the oil series. Given the less bublish behaviour of CO2 and that a purely non-causal AR(1) model seems to fit it better, we decide to consider it as a case of a slightly misspecified data generating process in the analysis and estimate a MAR(1,1) specification as for the other two series. A 6-months investment horizon is defined in all cases by fixing $H = 26$.

Table 4 reports the estimation results for the MAR(1,1) parameters that drive the dynamics of the bubble assets. The coefficients of the polynomial trend function are significant, supporting the use of the detrending strategy. Most importantly, the non-causal component dominates the causal one, revealing, for example, the forward-looking steady increase in the oil-price data followed by quite abrupt bubble bursts (see panels B and C). A similar pattern, although with less asymmetry around the shock, is observed in the case of the CO2 series (see panel A). At the same time, Table 5 displays the mean and standard deviation of the bubble-free assets, that are set to follow a GBM

[7] As discussed in Hencic and Gouriéroux (2015), it is important to detrend the series by a deterministic function of time and to avoid standard filtering and smoothing procedures that may induce spurious noncausal effects in the data.

Figure 1: Data

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

Table 4: Speculative assets: MAR(1,1) and trend estimation

Panel A : CO2					
Detrending method: polynomial of order 4					
(Intercept)	τ^1	τ^2	τ^3	τ^4	
9.015***	−0.5339***	0.01541***	−0.0001143***	0.0000002647***	
(6.88E-01)	(4.09E-02)	(7.15E-04)	(4.62E-06)	(9.89E-09)	
α-stable MAR(1,1)					
$\varphi^\bullet$	φ°	α	β	σ	μ
0.08***	0.716***	1.625***	-0.103***	0.732***	0.816***
(1.60E-04)	(1.54E-04)	(2.83E-05)	(2.08E-04)	(7.99E-05)	(2.81E-04)
Panel B : Brent					
Detrending method: polynomial of order 4					
(Intercept)	τ^1	τ^2	τ^3	τ^4	
54.41***	−0.413***	0.01916***	−0.0001765***	0.0000004488***	
(2.33E+00)	(1.40E-01)	(2.46E-03)	(1.61E-05)	(3.47E-08)	
α-stable MAR(1,1)					
$\varphi^\bullet$	φ°	α	β	σ	μ
0.049***	0.908***	1.65***	0.26***	1.48***	1.379***
(9.90E-06)	(1.99E-06)	(6.03E-06)	(1.56E-05)	(5.26E-06)	(2.78E-06)
Panel C : WTI					
Detrending method: polynomial of order 3					
(Intercept)	τ^1	τ^2	τ^3		
36.06***	0.9801***	−0.01017***	0.00002772***		
(2.27E+00)	(8.48E-02)	(8.52E-04)	(2.43E-06)		
α-stable MAR(1,1)					
$\varphi^\bullet$	φ°	α	β	σ	μ
0.158***	0.955***	1.927***	0.944**	1.709***	0.251***
(6.63E-04)	(3.49E-03)	(1.67E-02)	(4.27E-01)	(1.75E-02)	(3.16E-02)

Notes: Estimated parameters of the polynomial trend function and of the α-stable MAR(1,1) process associated with the detrended speculative series for the period 01.2017 - 06.2021. Standard deviations are in parentheses. Asterisks *, **, and *** indicate significance at the 90%, 95% and 99% level, respectively.

Table 5: Safer asset returns

EUREX-Euro bund		US 30-year Treasury-bond	
μ	σ	μ	σ
0.0001	0.0077	0.0001	0.0119

Notes: Mean and standard-deviation of the bubble-free asset returns for the period 01.2017 - 06.2021.

as commonly hypothesised in the financial literature.

Based on the in-sample data, we then compute the quantiles and conditional moments of the detrended series. The latter are fed to the CRRA portfolio optimization program with $\gamma \in \{5, 10\}$, which covers two levels of relative risk aversion.

Table 6: Optimal portfolio strategies under CRRA

Panel A: CO2

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 5$)	(0.01,1)	(-0.04,6)	(-0.08,5)	(-0.11,5)	(-0.16,4)	(-0.15,4)	(-0.08,4)	(-0.06,4)	(-0.05,4)
(w^*, h^*)	(0,0)	(0,0)	(0,0)		(-0.13,5)	(-0.12,5)			

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 10$)	(0.01,1)	(-0.02,6)	(-0.04,5)	(-0.06,5)	(-0.07,5)	(-0.08,4)	(-0.04,4)	(-0.03,4)	(-0.02,5)
(w^*, h^*)	(0,0)	(0,0)	(0,0)			(-0.06,5)			

Panel B: Brent

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 5$)	(0.07,1)	(-0.03,23)	(-0.07,20)	(-0.12,17)	(-0.15,15)	(-0.14,15)	(-0.09,14)	(-0.06,15)	(-0.05,15)
(w^*, h^*)	(0.02,17)	(0.03,1)	(0,0)	(0,0)					
	(0,0)	(0,0)							

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 10$)	(0.06,1)	(-0.01,26)	(-0.03,22)	(-0.06,18)	(-0.07,17)	(-0.07,15)	(-0.04,16)	(-0.03,15)	(-0.03,14)
(w^*, h^*)	(0.01,17)	(0.05,1)	(-0.04,19)	(0,0)					(-0.02,18)
	(0,0)	(0,0)	(0.03,1)						
			(0,0)						

Panel C: WTI

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 5$)	(0.02,1)	(-0.02,26)	(-0.06,26)	(-0.13,26)	(-0.22,26)	(-0.27,26)	(-0.22,26)	(-0.09,26)	(-0.09,26)
(w^*, h^*)	(0,0)	(0.02,1)	(0.02,1)						
		(0,0)	(0,0)						

$X_t = x$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	$q_{0.99}$	$q_{0.999}$	$q_{0.9999}$
CRRA ($\gamma = 10$)	(0.02,1)	(-0.01,26)	(-0.03,26)	(-0.06,26)	(-0.11,26)	(-0.14,26)	(-0.11,26)	(-0.05,26)	(-0.04,26)
(w^*, h^*)	(0,0)	(0.03,1)	(0.04,1)						
		(0,0)	(0,0)						

Notes: The Table displays the optimal portfolio strategies (w^*, h^*) for each portfolio conditional on the quantile of the in-sample distribution in which the bubble asset is at the time of the investment. w is reported in percentages of the investment and h in weeks.

Table 6 displays the optimal portfolio strateg(y/ies) in the form of couples (ω, h) identified using the in-sample data for the three portfolios under analysis for a range of quantiles in the upper half of the distribution of the bubble asset. The amount invested, w, may vary quite a lot across quantiles and investors. Apart from some cases that belong to the center of the distribution, the investor has incentives to take short positions. These are generally small in terms of volume of trading (low w) and over relatively short time periods (about one month) independent of the quantile of X_t in the case of the CO2, which corresponds to the particularities of the series. In contrast, the more bubblier oil series are characterized by larger investment horizons, of around 4 months for the Brent and of 6 and a half months for WTI (meaning that very few rebalancings are performed). This finding seems reasonable for investments in oil price, as the largest actors generally take long term hedging positions and the adjustment mechanisms take a long time to be put in place given the strong links with the underlying industry. Besides, for a given quantile, the more risk-averse the investor, the smaller the optimal allocation to the bubble asset.

4.2 Out-of-sample performance analysis

To check the performance of these strategies, we turn to the out-of-sample data. First, we rely on the polynomial coefficient estimates to extend the trend dynamics and remove it from the data.[8] This allows us to match each out-of-sample detrended value that plays the role of a conditioning price, X^d_{T+k}, with the closest floor empirical in-sample quantile of the detrended series and identify the associated portfolio allocation strateg(y/ies). To make the *BP* approach realistic, we allow for portfolio rebalancing. The investment horizon h being endogenous at each rebalancing step, an irregular, path dependent set of conditioning prices is obtained and the procedure described above to identify the optimal strateg(y/ies) is applied at each time. The position in the bubble asset is closed either when the most recent allocation strategy points to a rebalanced investment horizon h which goes beyond the full investment horizon H, or when a $(\omega^*, h^*) = (0,0)$ strategy is identified as the optimal. To simplify the understanding of the rebalancing procedure, a decision path is displayed in Figure 2.

As for a given conditioning value an investor can choose among various optimal strategies and (s)he can rebalance the portfolio several times up to the investment horizon H, the perfor-

[8] The polynomial coefficients are reestimated in a recursive framework for each new out-of-sample period, i.e. based on all the previous observed values of the bubble series.

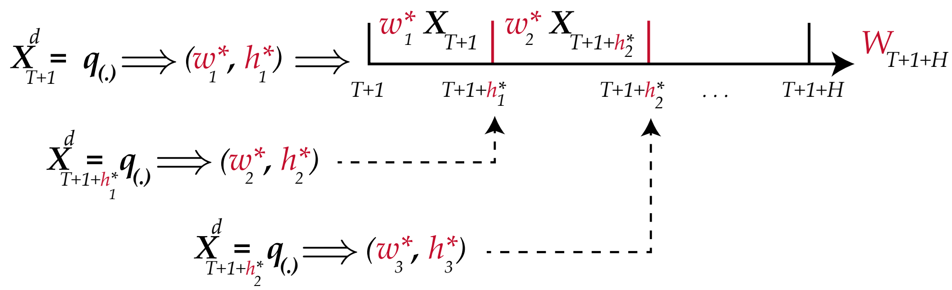

Figure 2: Rebalancing path

mance results for the *BP* approach for a given initial conditioning price are expressed in terms of quantiles of the distribution of the terminal wealth and other performance measures, and are denoted by $BP_{10\%}$, $BP_{50\%}$, and $BP_{90\%}$, respectively. These results are subsequently aggregated over the out-of-sample (based on the full set of initial conditioning prices).

In Table 7 we report some summary statistics including the average over the out-of-sample of the median and the standard deviation of the number of rebalancings and of the number of investment paths. On average, our approach involves a median rebalancing of the portfolio between 2 and 6 weeks over the 6-months horizon *H* depending of the underlying bubble series. In particular, the CO2-EUREX euro bund portfolio is more often rebalanced than the oil portfolios. This result holds quite uniformly over the distribution of the strategies, i.e. the differences between the three *BP* strategies are small. The advantage here is that few rebalancings simplify portfolio management operations relative to the daily rebalancing approach of the benchmark portfolios. In contrast, the median of investment paths differs a lot according to the coefficient of relative risk aversion. For example, a more risk averse investor seems to rebalance more, which may quickly lead to a dense tree of possible out-of-sample investment paths over the 26 weeks horizon when we look at all strategies, i.e. $BP_{50\%}$. At the same time, the benchmarks are exogenously set to rebalance every day, as usually done in the literature, and exhibit a unique optimal strategy, which is driven by the weight of the bubble asset in the portfolio. Empirically we observe a quite steady level of these weights over the investment horizon (see Section 4 in the Web-Appendix

Table 7: Number of rebalancings and strategies

		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
Panel A: CO2						
CRRA ($\gamma = 5$)	$Med_{N^oRebalancings}$	6	6	6	26	26
	$\sigma_{N^oRebalancings}$	0.712	0.622	1.371	0	0
	Med_{N^oPaths}	1	1	1	1	1
	σ_{N^oPaths}	0	1.071	0	0	0
CRRA ($\gamma = 10$)	$Med_{N^oRebalancings}$	4.556	4.182	3.952	26	26
	$\sigma_{N^oRebalancings}$	0.831	0.623	0.518	0	0
	Med_{N^oPaths}	15	149	15	1	1
	σ_{N^oPaths}	47.862	510.238	47.862	0	0
Panel B: Brent						
CRRA ($\gamma = 5$)	$Med_{N^oRebalancings}$	2	2.333	2	26	26
	$\sigma_{N^oRebalancings}$	0.536	0.173	0.558	0	0
	Med_{N^oPaths}	1	3	1	1	1
	σ_{N^oPaths}	0	0.558	0	0	0
CRRA ($\gamma = 10$)	$Med_{N^oRebalancings}$	2.167	2.4	3	26	26
	$\sigma_{N^oRebalancings}$	0.674	0.478	0.875	0	0
	Med_{N^oPaths}	1	5.5	1	1	1
	σ_{N^oPaths}	3.717	37.919	3.717	0	0
Panel C: WTI						
CRRA ($\gamma = 5$)	$Med_{N^oRebalancings}$	2	2	2	26	26
	$\sigma_{N^oRebalancings}$	0.295	0.197	0.635	0	0
	Med_{N^oPaths}	1	1	1	1	1
	σ_{N^oPaths}	0	0.52	0	0	0
CRRA ($\gamma = 10$)	$Med_{N^oRebalancings}$	2	4.182	3.952	26	26
	$\sigma_{N^oRebalancings}$	0.599	0.623	0.518	0	0
	Med_{N^oPaths}	1	149	15	1	1
	σ_{N^oPaths}	0.612	510.238	47.862	0	0

Notes: The table displays the average over the out-of-sample of the median and the standard deviation of the number of rebalancings and portfolio trajectories in the 10% best (resp. worst) performing *BP* strategies, $BP_{90\%}$ (resp. $BP_{10\%}$) as well as over the full sample of MAR strategies with horizon H ($BP_{50\%}$).

for further details). The approach to obtain them is standard in the literature for both benchmarks and shall not be detailed further.

We evaluate the out-of-sample performance of the MAR(1,1)-based strategies through two criteria. First, the out-of-sample terminal wealth, W_{T+H}, associated with each investment path, which is obtained from (1) by applying at each rebalancing date the appropriate optimal strategy (ω, h) to the out-of-sample non-detrended price data. At the same time, we rely on the opportunity cost as a relative measure of performance.

Table 8: Relative performance of portfolio strategies under CRRA: terminal wealth

			Panel A: CO2			
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 5$)	μ	$1.058^{/***}$	$1.069^{/***}$	$1.080^{*/***}$	1.049	0.922
	σ	0.104	0.117	0.127	0.140	0.046
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 10$)	μ	$0.992^{/***}$	$1.140^{***/***}$	$1.361^{***/***}$	1.049	0.922
	σ	0.072	0.104	0.134	0.140	0.046
			Panel B: Brent			
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 5$)	μ	$1.042^{/***}$	$1.049^{/***}$	$1.072^{**/***}$	1.001	0.921
	σ	0.073	0.079	0.112	0.152	0.044
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 10$)	μ	$0.991^{/***}$	$1.085^{***/***}$	$1.234^{***/***}$	1.001	0.921
	σ	0.08	0.100	0.191	0.152	0.044
			Panel C: WTI			
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 5$)	μ	$1.201^{***/***}$	$1.201^{***/***}$	$1.201^{***/***}$	0.990	0.932
	σ	0.177	0.176	0.176	0.167	0.056
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	EW	MV
CRRA ($\gamma = 10$)	μ	$1.145^{***/***}$	$1.227^{***/***}$	$1.312^{***/***}$	0.990	0.932
	σ	0.189	0.177	0.179	0.167	0.056

Notes: The quantiles of our MAR(1,1)-based strategies are compared with the equally weighted (EW) and mean-variance (MV) ones in terms of terminal wealth. The results take the form of out-of-sample average and standard deviation. Asteriks $(*/*)$ indicate the rejection of the null hypothesis of Wilkoxon's test at the 90%, 95% and 99% levels.

The first set of results is reported in Table 8. Each panel (A,B,C) corresponding to a portfolio displays the average and the standard deviation of the terminal wealth over the out-of-sample period corresponding to each of the five portfolio strategies under analysis. Asteriks $(*/*)$ associated with the estimated μ of each of our strategies indicate whether the median performance metric for our approach is statistically different from that of the (EW/MV) portfolio according to Wilcoxon's test.

The results indicate a positive gain in using our approach relatively to the standard *MV* and *EW* portfolios. The terminal wealth indicates a positive return on investment for the *BP* strategies for all three bubble-asset portfolios and both relative risk aversion coefficients. The only exception is that of the worst 10% strategies of the CO2-EUREX Euro Bund and Brent-EUREX Euro Bund portfolios for the most risk averse investors. In contrast, the *EW* strategy is a winning strategy, with a return on investment of about 5%, for the first portfolio and as the degree of non-causality of the speculative asset increases its performance drops (see panels B and C). Note also that there is not much difference in terms of wealth dispersion between our approach and the *EW*, although one notices an increase in the standard deviation with the degree of non-causality. Most strikingly, the mean-variance approach always exhibits the lowest dispersion and a negative return on investment. This result suggests that the *MV* strategy is not fit for bubble assets.

Finally, Table 9 reports the opportunity cost relative performance measure for the three portfolios under analysis in two cases, without and with transaction costs. Recall that a positive value indicates the amount that must be added to the return of the benchmark strategy, such that it leaves the investor indifferent to the decision between it and the corresponding *BP* strategy. The OC is generally positive, hence supporting the superiority of our strategy relative to the benchmarks. The only case where the *EW* portfolio is preferred is when the investor bets on the 10% worst *BP* strategies for CO2 and Brent bubble assets.

All in all, the MAR(1,1) process captures well the bubble behaviour of the series, which is shown to generally materialize in better portfolio performance relative to the traditional benchmarks.

4.3 Robustness Analysis

We now investigate the sensitivity of our results to various changes in the portfolio analysis by: i) relying on a different optimisation algorithm, the PGP, ii) accounting for transaction costs, and

Table 9: Relative performance of portfolio strategies under CRRA: opportunity cost

Panel A: CO2

		EW vs			*MV* vs		
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$
CRRA ($\gamma = 5$)	μ	0.009	0.020	0.031	0.137	0.147	0.158
	σ	0.211	0.220	0.228	0.127	0.138	0.149
CRRA ($\gamma = 10$)	μ	-0.056	0.092	0.312	0.071	0.219	0.440
	σ	0.151	0.155	0.179	0.084	0.114	0.141

Panel B: Brent

		EW vs			*MV* vs		
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$
CRRA ($\gamma = 5$)	μ	0.041	0.048	0.071	0.121	0.128	0.151
	σ	0.217	0.221	0.242	0.099	0.104	0.134
CRRA ($\gamma = 10$)	μ	-0.01	0.084	0.233	0.070	0.164	0.313
	σ	0.188	0.201	0.295	0.105	0.126	0.215

Panel C: WTI

		EW vs			*MV* vs		
		$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$	$BP_{10\%}$	$BP_{50\%}$	$BP_{90\%}$
CRRA ($\gamma = 5$)	μ	0.210	0.210	0.211	0.269	0.269	0.269
	σ	0.315	0.314	0.314	0.220	0.220	0.220
CRRA ($\gamma = 10$)	μ	0.154	0.236	0.322	0.213	0.295	0.380
	σ	0.287	0.249	0.247	0.206	0.187	0.193

Notes: The quantiles of our MAR(1,1)-based strategies are compared with the equally weighted (*EW*) and mean-variance (*MV*) ones in terms of opportunity cost (OC).

iii) ignoring the higher-order moments in the portfolio optimisation.

First, we rely on the PGP as an alternative portfolio allocation approach that explicitly accounts for higher-order moments in asset price distribution. The technical details on PGP as well as the empirical results are available in the Web-Appendix. We find that the PGP investor who weights more the fourth moment appears to choose more extreme allocations in the bubble asset for both long and short strategies. Still, the results indicate a positive gain in using our approach relatively to the standard *MV* and *EW* portfolios.

At the same time, transaction costs play an important role in the investment decisions as they impact portfolio performance proportionally to the number and / or value of rebalancings. For this reason, we gauge their impact on the performance of our portfolio strategies relative to the benchmarks in a simple setup where transaction costs are fixed at 0.05% per unit of investment. The results, available in Section 4.2 of the Web-Appendix are qualitatively similar to those without transaction costs. Note however that the *BP* strategies register a slightly larger drop in terminal wealth than the benchmarks: 0.010 vs 0.003 on average. Indeed, as the transaction costs are proportional to the shares of both assets that are exchanged, although our approach involves sparse rebalancing, the amounts involved in each trade are larger than those rebalanced daily by the benchmark portfolios (see Figure 1 in the Web-Appendix for the dynamics of EW and MV portfolio weights).

Finally, we gauge the optimal portfolio allocation according to the BP strategies when the higher-order moments are ignored. In the spirit of a stress-test, this analysis puts our MAR(1,1)-based approach on a more equal footing with the benchmarks, which make use only of information in the first two moments. The main difference is nevertheless, that the BP approach uses conditional and not unconditional moments. The main insight is that this version of the bubble portfolio approach (BP2) generally performs much better than the benchmarks. It relies on different optimal strategy couples (ω, h), which imply shorter positions on longer horizons and less rebalancing than BP. Its performance seems to be slightly better than that of BP for the CO2 data, but the results are more mitigated for the other two bubble series.[9] We therefore suggest one to rely on the more general four-moments based approach and use BP2 for sensitivity analysis. For space-constraints all tables and detailed analyses are deferred to the Web-Appendix.

4.4 Tree path analysis

To better illustrate an investor's decision-making process, we consider the case of the WTI - US Treasury-bond portfolio and choose the out-of-sample starting point so as to be in a bubble period. More precisely, the investment time is set to 04 March 2022, when the WTI detrended

[9] Note also that in the BP2 case there is only one PGP strategy, as the fourth moment, which differentiated then before, is now irrelevant.

price is in the 80% quantile of its in-sample distribution.

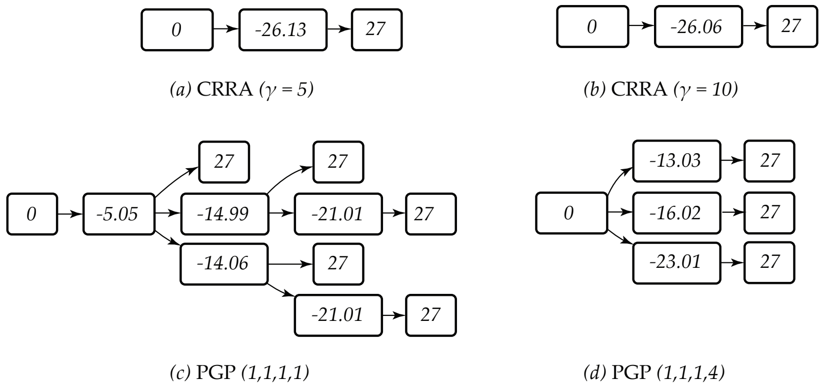

Figure 3: Decision-making trees

Reading key: $\text{sign}(w) * h.w$, with w the investment share in the bubble asset and h the horizon of this investment.

Figure 3 displays the possible investment paths for the two optimisation algorithms and the different investor preferences used in the main analysis. At each rebalancing time, the optimal strateg(y/ies) are expressed as "$\text{sign}(w) * h.w$". 0 denotes the initial investment time and 27 the closure of all positions once the investment horizon H is reached. Panel (a) reveals that in the CRRA (γ=5) case there is only one path, involving a single portfolio rebalancing which occurs the day before closing the positions. A similar path would follow a more risk averse investor, who, in turn, would invest a shorter amount in the bubble asset (w=0.06 in panel (b) instead of 0.13 in the previous case). In the PGP case the tree is more dense, with five possible paths when the four conditional moments of the return distribution are equally important and three such paths when the investor pays more attention to the kurtosis. When the tree involves several paths, they are all locally optimal for a given level of risk aversion and hence the investor relies on exogenous constraints to select among them. For example, an investor with stronger constraints or whose client announces an anticipated liquidation of the positions held, would choose lower investment horizons.

5 Conclusion

In this paper we propose an asset allocation strategy particularly designed for the case of an investors wishing to include a bubble asset in his/her portfolio. For this, we account explicitly for the distributional characteristics of bubble assets through a MAR(1,1) model, which seems to be appropriate to capture locally explosive behaviours. The higher-order conditional moments of the return distribution are then plugged in the Taylor-series-expansion of the CRRA utility function and the PGP algorithm, respectively. The economic value of the *BP* strategy is compared in out-of-sample with standard benchmarks such as the mean-variance and equally-weighted portfolios based on well-known performance measures such as the opportunity cost. Both simulation-based analyses and empirical results using CO2, Brent and WTI data for the bubble asset support the superiority of our approach. The stronger the anticipative character of the bubble behaviour of the asset, i.e. the larger the non-causal parameter, the wider the hedging possibilities and the more substantial the possible gains involved.

6 Appendix: Proofs

6.1 Conditional moments of returns

For the bubble asset we have the non-central moments of returns

$$\mathbb{E}(r_{t,t+h}^{X^j}|X_t) = \frac{1}{X_t}\mathbb{E}(X_{t,t+h}^{p}|X_t), \quad \text{for} \quad p \geq 1,$$

where the conditional moments of the price series are defined through Proposition 1.

For the second asset, using the properties of the GBM, one obtains

$$\begin{aligned}
\mathbb{E}(r^{S}_{t,t+h}|X_t) &= e^{\upsilon h} - 1 \\
\mathbb{E}(r^{S^2}_{t,t+h}|X_t) &= e^{2\upsilon h+\varsigma^2 h} - 2e^{\upsilon h} + 1 \\
\mathbb{E}(r^{S^3}_{t,t+h}|X_t) &= e^{3\upsilon h+3\varsigma^2 h} - 3e^{2\upsilon h+\varsigma^2 h} + 3e^{\upsilon h} - 1 \\
\mathbb{E}(r^{S^4}_{t,t+h}|X_t) &= e^{4\upsilon h+4\varsigma^2 h} - 4e^{3\upsilon h+3\varsigma^2 h} + 6e^{2\upsilon h+\varsigma^2 h} - 4e^{\upsilon h} + 1,
\end{aligned}$$

and similar expressions, using $H - h$ instead of h, can be derived for $\mathbb{E}(r^{S}_{t+h,t+H}|X_t)$. Making use of the mapping relations between central and non-central moments and the independence between the two assets, one can subsequently express $\mathbb{E}\Big[(R_{t+H} - \overline{R}_{t+H})^k|X_t, S_t\Big]$ as a function of these moments.

6.2 Derivation of the constants $\sigma_1^\alpha, \beta_1, \kappa_p$, and λ_p

Fries 2021 shows that if X_t is a α-stable two-sided MA(∞) process with $0 < \alpha < 2, \alpha \neq 1$, $\beta \in [-1, 1]$, and $\sigma > 0$ as defined in Section 2.2, i.e. well defined, stationary process with α-stable errors, and for $h \geq 1$ then one can obtain the conditional moments of the process X_t for $p \leq 4$ with

$$\sigma_1 = \sigma^\alpha \sum_{k\in\mathbb{Z}} |a_k|^\alpha, \qquad \beta_1 = \beta \frac{\sum_{k\in\mathbb{Z}} a_k^{<\alpha>}}{\sum_{k\in\mathbb{Z}} |a_k|^\alpha}, \qquad \kappa_p = \frac{\sum_{k\in\mathbb{Z}} |a_k|^\alpha \left(\frac{a_{k-h}}{a_k}\right)^p}{\sum_{k\in\mathbb{Z}} |a_k|^\alpha}, \qquad \lambda_p = \frac{\sum_{k\in\mathbb{Z}} a_k^{<\alpha>} \left(\frac{a_{k-h}}{a_k}\right)^p}{\sum_{k\in\mathbb{Z}} |a_k|^\alpha},$$

where $y^{<\alpha>} = \text{sign}(y)|y|^\alpha$ for any $y \in \mathbb{R}$. Using his results together with the fact that the coefficients of the $MA(\infty)$ representation of a MAR(1,1) process, $X_t = \sum_{k=-\infty}^{\infty} a_k \varepsilon_{t-k}$, satisfy

$$\begin{aligned}
a_k &= \frac{\varphi^{\circ k}}{1 - \varphi^\circ \varphi^\bullet} && \text{if } k \geq 0, \\
a_k &= \frac{\varphi^{\bullet -k}}{1 - \varphi^\circ \varphi^\bullet} && \text{otherwise,}
\end{aligned}$$

and calculus based on geometric series, one can easily obtain the results in Proposition 1.

References

Ang, A., Chen, J., Xing, Y., 2006. Downside risk. The Review of Financial Studies 19, 1191–1239.

Biagini, F., Föllmer, H., Nedelcu, S., 2014. Shifting martingale measures and the birth of a bubble as a submartingale. Finance and Stochastics 18, 297–326.

Biagini, F., Nedelcu, S., 2015. The formation of financial bubbles in defaultable markets. SIAM J. Fin. Math. 6, 530–558.

Blanchard, O.J., Watson, M.W., 1982. Bubbles, Rational Expectations and Financial Markets. NBER Working Papers 0945. National Bureau of Economic Research, Inc.

Briec, W., Kerstens, K., Jokung, O., 2007. Mean-variance-skewness portfolio performance gauging: A general shortage function and dual approach. Management Science 53, 135–149.

Briec, W., Kerstens, K., Van de Woestyne, I., 2013. Portfolio selection with skewness: A comparison of methods and a generalized one fund result. European Journal of Operational Research 230, 412–421.

Carvalho, V.M., Martin, A., Ventura, J., 2012. Understanding bubbly episodes. American Economic Review 102, 95–100.

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

Davis, M., Lleo, S., 2013. Fractional Kelly strategies in continuous time: recent developments. chapter 37. pp. 753–787.

De Athayde, G.M., Flôres Jr, R.G., 2004. Finding a maximum skewness portfolio—a general solution to three-moments portfolio choice. Journal of Economic Dynamics and Control 28, 1335–1352.

Favre, L., Galeano, J.A., 2002. Mean-modified value-at-risk optimization with hedge funds. The Journal of Alternative Investments 5, 21–25. doi:10.3905/jai.2002.319052.

Friedrich, M., Fries, S., Pahle, M., Edenhofer, O., 2019. Understanding the explosive trend in EU ETS prices–fundamentals or speculation? arXiv preprint arXiv:1906.10572 .

Fries, S., 2021. Conditional moments of noncausal alpha-stable processes and the prediction of bubble crash odds. Journal of Business & Economic Statistics 0, 1–21.

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

Ghahtarani, A., 2021. A new portfolio selection problem in bubble condition under uncertainty: Application of z-number theory and fuzzy neural network. Expert Systems with Applications 177.

González-Pedraz, C., Moreno, M., Peña, J.I., 2015. Portfolio selection with commodities under conditional copulas and skew preferences. Quantitative Finance 15, 151–170.

Gourieroux, C., Jasiak, J., 2018. Misspecification of noncausal order in autoregressive processes. Journal of Econometrics 205, 226–248.

Gourieroux, C., Jasiak, J., Monfort, A., 2020. Stationary bubble equilibria in rational expectation models. Journal of Econometrics 218, 714–735.

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

Guidolin, M., Timmermann, A., 2008. International asset allocation under regime switching, skew, and kurtosis preferences. Review of Financial Studies 21, 889–935.

Hashimoto, K.i., Im, R., 2016. Bubbles and unemployment in an endogenous growth model. Oxford Economic Papers 68, 1084–1106.

Hashimoto, K.i., Im, R., 2019. Asset bubbles, labour market frictions and R&D-based growth. Canadian Journal of Economics/Revue canadienne d'économique 52, 822–846.

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

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

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

Hugonnier, J., 2012. Rational asset pricing bubbles and portfolio constraints. Journal of Economic Theory 147, 2260–2302.

Ingersoll, J., 1975. Multidimensional security pricing. Journal of Financial and Quantitative Analysis 10, 785–798.

Jarrow, R., Protter, P., Shimbo, K., 2010. Asset price bubbles in incomplete markets. Mathematical Finance 20, 145–185. doi:10.1111/j.1467-9965.2010.00394.x.

Jarrow, R.A., Protter, P., Roch, A.F., 2012. A liquidity-based model for asset price bubbles. Quantitative Finance 12, 1339–1349.

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

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

Kolm, P.N., T'fct'fcnc'fc, R., Fabozzi, F., 2014. 60 years of portfolio optimization: Practical challenges and current trends. European Journal of Operational Research 234, 356–371.

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

Lassance, N., Vrins, F., 2021. Minimum rényi entropy portfolios. Annals of Operations Research 299, 23–46.

León, A., Moreno, M., 2017. One-sided performance measures under Gram-Charlier distributions. Journal of Banking & Finance 74, 38–50.

Lhabitant, F.S., et al., 1998. On the (ab)use of Taylor series approximations for portfolio selection, portfolio performance and risk management. Working Paper. University of Lausanne.

Lux, T., Sornette, D., 2002. On rational bubbles and fat tails. Journal of Money, Credit and Banking 34, 589–610.

Mandelbrot, B., 1963. The variation of certain speculative prices. The Journal of Business 36.

Markowitz, H., 1952. Portfolio selection. The Journal of Finance 7, 77–91.

Markowitz, H., 2014. Mean-variance approximations to expected utility. European Journal of Operational Research 234, 346–355.

Martellini, L., Ziemann, V., 2010. Improved estimates of higher-order comoments and implications for portfolio selection. The Review of Financial Studies 23, 1467–1502.

Martin, A., Ventura, J., 2012. Economic growth with bubbles. American Economic Review 102, 3033–58.

Massacci, D., 2017. Tail risk dynamics in stock returns: links to the macroeconomy and global markets connectedness. Management Science 63, 3072–3089.

Miao, J., Wang, P., Xu, L., 2016. Stock market bubbles and unemployment. Economic Theory 61, 273–307.

Phillips, P.C., Shi, S., Yu, J., 2015. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. International economic review 56, 1043–1078.

Protter, P., 2012. A mathematical theory of financial bubbles. SSRN Electronic Journal 2081. doi:10.2139/ssrn.2115895.

Samorodnitsky, G., Taqqu, M.S., Linde, R., 1996. Stable non-gaussian random processes: stochastic models with infinite variance. Bulletin of the London Mathematical Society 28, 554–555.

Samuelson, P.A., 1967. General proof that diversification pays. Journal of Financial and Quantitative Analysis 2, 1–13.

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

Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? American Economic Review 71, 421–436.

Tirole, J., 1985. Asset bubbles and overlapping generations. Econometrica , 1499–1528.

West, K.D., 1987. A specification test for speculative bubbles. The Quarterly Journal of Economics 102, 553–580.

West, K.D., 1988. Bubbles, fads and stock price volatility tests: a partial evaluation. The Journal of Finance 43, 639–656.

Zoia, M.G., Biffi, P., Nicolussi, F., 2018. Value at Risk and Expected Shortfall based on Gram-Charlier-like expansions. Journal of Banking & Finance 93, 92–104.