

Performance of the CoinShares Gold and Cryptoassets Index Under Different Market Regimes

Aikaterini Koutsouri
Department of Computing
Imperial College London
London, SW7 2AZ, UK
k.koutsouri@imperial.ac.uk

Michael Petch
CoinShares
Octagon Point, 5 Cheapside
London, EC2V 6AA, UK
mpetch@coinshares.com

William J. Knottenbelt
Department of Computing
Imperial College London
London, SW7 2AZ, UK
wjk@doc.ic.ac.uk

Abstract—Regime-switching models are frequently used to explain the tendency of financial markets to change their behavior, often abruptly. Such changes usually translate to structural breaks in the average means and volatilities of financial indicators, and partition their time-series into distinct segments, each with unique statistical properties. In this paper, we address the problem of identifying the presence of such regimes in the constituents of diversified, cryptoasset-containing portfolios, ultimately to define high-risk market conditions and assess portfolio resilience. For each portfolio component, we first consider a Gaussian Hidden Markov Model (HMM) in order to extract intermediate trend-related states, conditional on the weekly returns distributions. We further apply a Markov-switching GARCH model to the demeaned daily returns to describe changes in the conditional variance dynamics and isolate volatility-related states. We combine the former approaches to generate a number of price paths for each constituent, simulate the portfolio allocation strategy and obtain a risk profile for each combination of the trend and volatility regimes. We apply the proposed method to the CoinShares Gold and Cryptoassets Index, a diversified, monthly-rebalanced index which includes two main risk-weighted components; a cryptoassets basket and physical gold. Results demonstrate an overall stable risk-reward profile when compared against the individual components and suggest a superior performance in terms of Omega ratio for investors that target wealth preservation and moderate annual returns. We detect underperformance regions in bear-low volatility market regimes, where diversification is hindered.

Index Terms—Cryptoassets, Gold, Index, Hidden Markov Model, Regime-switching GARCH

I. INTRODUCTION

The study of market regimes is important in many fields, including governmental policy, financial markets and regulation. Use-cases span from Investment Banks who attempt to determine when market regime changes occur (e.g. Morgan Stanley Regime Switching Index, MSCEEMRI) to Central Banks aiming to estimate the occurrence probability of high-stress scenarios [1] in order to signal and mitigate financial risk. Methods including estimation of conditional volatilities are key for risk-monitoring processes and, while the original works of Engle [2] and Bollerslev [3] have been widely adopted and expanded by risk managers, studies highlight the

presence of structural breaks in the dynamics of financial time-series.

The first application of regime switching approaches is found in the works of Hamilton [4] which examine how economic activity fluctuates between states of expansion and recession. Since then, Markov-switching approaches have extended to different specifications, with Hamilton and Susmel [5] presenting a conditional heteroskedasticity (ARCH) setting with a Markov-switching specification in the state alteration. Other extensions are found in the works of Bauwens et al. [6] and Haas [7] which investigate stock market indices and categorise time periods according to volatility changes. Other applications include attempts to forecast stock prices [8], [9], portfolio allocation methodologies [10], [11] and univariate Value-at-Risk (VaR) estimations [12].

In the cryptoasset space, previous studies have investigated single-regime generalised conditional heteroskedasticity (GARCH) model estimations [13], [14]. Nevertheless, the weaknesses of single-regime models have also been highlighted in existing literature, with Molnár and Thies [15] detecting structural changes in Bitcoin pricing data. Ardia et al. [16] address the switch in the Bitcoin returns process through a Markov-switching GARCH model (MSGARCH) whose parameters adapt to variations in the unconditional variance. They further show that the 2-state MSGARCH approach in volatility modeling improves one-day ahead VaR predictions. Koki et al. [17] further study cryptoasset prices via a Non-Homogeneous Pólya Gamma Hidden Markov (NHPG) model. Their findings identify two states – high and low volatility, with frequent transitions between the two – and the proposed model is characterised by good in-sample performance but poor posterior out-of-sample predictions.

The inclusion of cryptoassets in investment portfolios has also been a point of interest in recent studies. Henriques and Sadorsky [18] refer to gold as a safe haven asset with important diversification capabilities and recognise that its elimination from investment portfolios can potentially negatively impact their risk-return profile. Motivated by the occasional referral to Bitcoin as digital gold, they examine whether Bitcoin can replace physical gold in traditional investment portfolios and how such scenario would impact the risk-adjusted returns. Their approach uses multivariate GARCH models to min-

Imperial College of London gratefully acknowledges the support given by CoinShares. Their funding provides Imperial with the support needed to conduct the research found within this paper.

imise variance, given a target return and is applied to a US benchmark portfolio that includes gold and a portfolio that substitutes gold for bitcoin. They conclude that the bitcoin-containing portfolio ranks higher in terms of risk-adjusted returns. Nevertheless, their model assumes daily rebalancings with no restrictions on short sales and is limited in historical data up to 2018, therefore disregarding a large portion of the crypto-market’s recent dynamics.

The diversification properties of cryptoassets are also examined by Antipova [19]. Empirical results show that global portfolios display better performance when they utilise the crypto-market as a diversification mechanism. Additionally, it is shown that better results are achieved when portfolios are exposed to multiple cryptoassets rather than solely Bitcoin. Optimization approaches on crypto-containing portfolios have also been studied, with Brauneis et al. [20] examining a traditional mean-variance framework. Castro et al. [21] suggest maximising the Omega ratio when optimising cryptoasset-based portfolios and consider four investment portfolios, two of which contain stock market indices in addition to cryptoassets. They conclude that, while crypto-exposure improves returns, it also increases risk. A detailed, comprehensive overview and analysis of further research work around cryptocurrency trading is presented by Fang et al. [22].

Notwithstanding past studies on the crypto-market dynamics and the diversification properties of cryptoassets, it is not clear how traditional allocation strategies (a) behave in relation to different market conditions and (b) meet individual investor’s expectations. The primary goal of this paper is the application of regime-switching models to unveil high-risk market states for the diversification strategy employed by the CoinShares Gold and Cryptoassets Index (CGCI) and assess how appropriate it is for investors with different annual return targets. For this purpose, we isolate the index’s two main market determinants, namely the crypto-basket and gold, and study the dynamics of their price evolution both in terms of volatility and intermediate trend. In the proposed setting, we isolate four trend states and three volatility states for the crypto-basket and three trend states and three volatility states for gold. We produce 1 000 paths for each CGCI component and report the Omega ratio, both in isolation and combined, following the index rebalancing scheme. We also detect regions where the weighting scheme does not improve risk-adjusted returns due to limited diversification opportunities.

The remainder of this paper is organised as follows. Section II presents the essential background on the principles behind the design of the index. In Section III, we elaborate on the models that are used for the extraction of the market regimes. Intermediate trend – referring to general movement in the prices that can last up to a few weeks – is examined on the demeaned weekly logarithmic returns of each constituent through a Gaussian Hidden Markov Model. Volatility regimes are identified through a Markov-switching GARCH approach on the demeaned daily logarithmic returns, where parameters change according to switches in time-varying conditional variance. Section IV demonstrates the results of 1 000 simulations

for price paths, including all 144 combinations of regimes. Section V concludes.

II. BACKGROUND

The CoinShares Gold and Cryptoassets Index (CGCI) [23] is a monthly rebalanced index that aims to provide diversified exposure to the cryptoasset space in a way that strengthens resilience to unforeseen drawdowns during stressful crypto-market conditions. The index employs a weighting scheme based on the historical volatilities of its two components, an equally-weighted basket of five cryptoassets, and physical gold. The crypto-basket constituents are selected on a monthly basis and include the top five eligible cryptoassets based on the 6-month rolling mean of the free-float market capitalisation.

The design of the CGCI is based on the fact that the cryptoasset space is characterised by high levels of volatility and high intraclass correlation, therefore providing limited diversification opportunities for pure cryptoasset-containing portfolios. Taking into further consideration the lack of significant rolling correlation with physical gold, the CGCI utilises the concept of volatility harvesting and uses a weighted-risk contribution scheme as a rebalancing mechanism between a basket of five cryptoassets and gold.

The weighting among the crypto-basket and gold in the CGCI is computed through:

$$x_c = \frac{\rho - \frac{c}{g}}{c + g}; \quad (1)$$

$$x_g = 1 - x_c;$$

where $\rho = 4$, x_c and x_g are the weights for the crypto-basket and gold and c and g the historical volatilities of the crypto-basket and gold logarithmic daily returns.

The index is calculated following a two-stage allocation scheme that involves:

- 1) Computation of the 6-month rolling volatility of (a) the equally-weighted crypto-basket, and (b) gold;
- 2) Asset allocation among the crypto-basket and gold expressed as the bivariate weighted risk contribution problem presented in Eq. 1. The risk contribution ratio level, set as $\rho = 4$, indicates that 80% of the total risk emanates from the crypto-basket.

Since the persistent high levels of cryptoasset volatility and the lack of correlation among the crypto-market and gold are the main design hypotheses of the CGCI, it is important to consider the effect of changes in both factors. In [24], several scenarios are considered in a stress testing framework, with different specifications in conditional variance and correlation. The CGCI constituents are modelled individually using ARMA and GJR–GARCH processes, while their joint evolution is described through a t-Copula.

Several specifications, including historical and hypothetical plausibility-constrained simulated scenarios yield results that support the hypothesised superiority of diversified strategies, such as the one employed by the CGCI. Nevertheless, the considered framework is limited in its ability to consider the regime-conditional performance of the examined strategy.

III. METHODOLOGY

A. Regime-Switching Intermediate-Trend

We let V_t denote the daily value of a financial instrument at day t and we express its weekly log-return with $R_t = \ln(V_{t:n}/V_{t:1})$, where $V_{t:1}$ is the value on the first day and $V_{t:n}$ is the value on the last day of the week. In an attempt to model its intermediate-trend, we assume that R_t can be described through a pair of stochastic processes $f(S_t; R_t); t \in \mathbb{N}$ that follow a Hidden Markov Model (HMM) specification [25].

In this case, $fS_t; t \in \mathbb{N}$ represents a Markov chain that is not directly observable and $fR_t; t \in \mathbb{N}$ is a sequence of independent random variables conditional on S_t . At every time point t , the next state S_{t+1} is dependent only upon the current state S_t and the conditional distribution of R_t only depends on S_t . For the purposes of this study, the output variable R_t is assumed to follow a Gaussian model $(\mu_k; \sigma_k)$, conditional on state $S_t = k \in \{1, \dots, K\}$ so that:

$$R_t \sim \begin{cases} N(\mu_1; \sigma_1); & S_t = 1 \\ \vdots \\ N(\mu_K; \sigma_K); & S_t = K \end{cases} \quad (2)$$

The hidden state variable S_t is assumed to be defined on the discrete space $\{1, \dots, K\}$. The $K \times K$ transition matrix is time-invariant and denoted with P :

$$P = \begin{bmatrix} p_{1,1} & \dots & p_{1,K} \\ \vdots & \ddots & \vdots \\ p_{K,1} & \dots & p_{K,K} \end{bmatrix}$$

where $p_{i,j}$ is the probability of transition from state $S_{t-1} = i$ to state $S_t = j$, $0 < p_{i,j} < 1 \forall i, j \in \{1, \dots, K\}$ and $\sum_{j=1}^K p_{i,j} = 1 \forall i \in \{1, \dots, K\}$.

A main assumption of the HMM is that any observation of the response is statistically independent of the previous outputs. For the purposes of this study, autocorrelation effects in the returns time series are eliminated through filtering with a first-order autoregressive process, AR(1).

The joint distribution of a sequence of a series of T observations $fS_{1:T}; R_{1:T}g$ is written as:

$$P(S_{1:T}; R_{1:T}) = P(S_1) \prod_{t=2}^T P(S_t | S_{t-1}) P(R_t | S_t) \quad (3)$$

The definition of the probabilistic network described by Eq. 3 requires specifying the probability distribution over the initial state $P(S_1)$, the $K \times K$ transition matrix that describes the evolution of the state variable, $P(S_t | S_{t-1})$, and the output model $P(R_t | S_t)$.

The prior, transition and response parameters $(\mu_k; \sigma_k)$ are estimated through the expectation-

maximisation (EM) algorithm by maximising their expected joint log-likelihood:

$$\log(P(S_{1:T}; R_{1:T})) = \log(P(S_1)) + \sum_{t=2}^T \log(P(S_t | S_{t-1})) + \sum_{t=1}^T \log(P(R_t | S_t)) \quad (4)$$

B. Regime-Switching Volatility

We denote the daily value of a given financial instrument at time t by V_t and the daily logarithmic returns by r_t , which satisfy the moment conditions $E[r_t] = 0$ and $E[r_t r_{t-l}] = 0$ for $l \neq 0$ and $t > 0$. In an attempt to capture the time-varying volatility behavior, we express r_t in terms of a process that follows a regime-switching specification in its conditional variance h_t . The general mixture model, which allows for categorisation of the conditional variance dynamics in low, moderate and high-volatility periods, can be expressed as:

$$r_t | (S_t = k; I_{t-1}) \sim D(0; h_{k;t}; \nu_k) \quad (5)$$

where I_{t-1} represents any information about r_t that is observed up to time $t-1$ and $D_k(0; h_{k;t}; \nu_k)$ is a continuous distribution, corresponding to state S_t , with zero mean and time-dependent conditional variance $h_{k;t}$. Additionally, ν_k denotes the shape parameters of the distribution of the independent and identically distributed standardised innovations $\epsilon_{k;t} \stackrel{i.i.d.}{\sim} D_k(0; 1; \nu_k)$ of the conditional variance process. Assuming that there are K different specifications of $h_{k;t}$, the latent variable S_t is defined on the discrete space $\{1, \dots, K\}$. The definition of the volatility regime-switching model defined by Eq. 5, requires specifying (i) the conditional variance dynamics, unique in each state $S_t = k$ and (ii) the state dynamics, driving the evolution of the variable S_t .

In regards to the conditional variances $h_{k;t}$ of r_t , we adopt the approach of Haas et al. [7] and assume that they follow K separate GARCH type processes which evolve in parallel. Therefore, given time t , $h_{k;t}$ follows a GARCH specification, conditional on regime $k \in \{1, \dots, K\}$ that prevails at time t :

$$h_{k;t} = h(r_{t-1}; h_{k;t-1}; \nu_k) \quad (6)$$

where h is a function that denotes the GARCH expression of the conditional variance and ensures positivity and covariance stationarity and ν_k are the model-specific parameters.

Assuming the ARCH model of Engle [2], we obtain:

$$h_{k;t} = a_{0,k} + a_{1,k} r_{t-1}^2 \quad (7)$$

where $k \in \{1, \dots, K\}$, $\nu_k = [a_{0,k}; a_{1,k}]$, $a_{0,k} > 0$ and $a_{1,k} < 0$. Additionally, we require $a_{1,k} < 0$ to ensure covariance stationarity in regime k . We further assume that the state process $S_t \in \{1, \dots, K\}$ evolves according to a first-order ergodic homogeneous Markov chain, with $K \times K$ probability matrix P .

Similarly, if we assume that the conditional variance follows the GARCH model of Bollerslev [3], we have:

$$h_{k;t} = a_{0;k} + a_{1;k}r_{t-1}^2 + \alpha_k h_{k;t-1}; \quad (8)$$

where $k \in \{1, \dots, K\}$, $\alpha_k = [a_{0;k}; a_{1;k}; \alpha_k]$, $a_{0;k} > 0$, $a_{1;k} > 0$ and $\alpha_k < 1$. Covariance stationarity is ensured through $a_{1;k} + \alpha_k < 1$.

Estimation of the specified model can be achieved through Maximum Likelihood approaches as described by Ardia et al. [26]. Given the Markov switching GARCH model parameters $\theta = (\alpha_1; \alpha_2; \dots; \alpha_K; \alpha_K; P)$, where α_i are the GARCH parameters and the parameters of the conditional distribution of the standardised innovations of state i and P the transition matrix, the likelihood function is given by:

$$L(\theta | j | T) = \prod_{t=1}^T f(r_t | j; |_{t-1}); \quad (9)$$

where $f(r_t | j; |_{t-1})$ denotes the density of r_t given its past observations $|_{t-1}$. Maximisation of the logarithm of $L(\theta | j | T)$ gives the ML estimator $\hat{\theta}$.

C. Market Regimes Simulation

The two presented models can be used to describe and simulate the evolution of V_t according to the regime switching behaviour of the market, both in terms of volatility and intermediate trend. In the case of regime-switching volatility, a Markov-switching ARCH(1) model is fitted on the daily logarithmic returns, while intermediate trend regime changes are observed in the weekly logarithmic returns.

Assuming the specifications of the previous sections and a discrete space of K_T distinct trend states $S; i \in \{1, \dots, K_T\}$ and K_V volatility states $s; i \in \{1, \dots, K_V\}$, the fitting and simulation process can be summarised in the following steps:

Model Fitting

- 1) Given V_t , observable on day t , obtain the daily logarithmic returns r_t and demean using an AR(1) filter to eliminate autocorrelation effects.
- 2) Fit the regime-switching ARCH(1) model to the demeaned r_t and obtain the $K_V \times K_V$ transition matrix P_V and the ARCH parameter vector $\alpha_k = [a_{0;k}; a_{1;k}]$ for each state s , $k \in \{1, \dots, K_V\}$.
- 3) Translate V_t to weekly logarithmic returns R_t and demean using an AR(1) filter to eliminate autocorrelation effects.
- 4) Fit the Gaussian HMM to demeaned returns and obtain the $K_T \times K_T$ transition matrix P_T and the distribution parameter vector $\beta_k = [\mu_k; \sigma_k]$ for each state S_k , $k \in \{1, \dots, K_T\}$.

Simulating Process

- 1) Given a simulation horizon of T days and parameters θ , simulate process r_t for $t = 1; \dots; T$, according to the fitted regime-switching ARCH(1) model, and translate to daily returns r_t according to the previously fitted AR(1).
- 2) Simulate process R_t for $t = 1; \dots; T$, according to the Gaussian HMM with distribution parameter vector $\beta_k = [\mu_k; \sigma_k]$, and translate to weekly returns R_t through the previously fitted AR(1).

3) For day t , offset r_t by $(1 + R_t)^{1=7}$

4) Convert the simulated daily logarithmic returns where $r_t \in \{1; \dots; T\}$, to daily values $V_t = V_0 e^{r_t}$, where $r_t = \sum_{i=1}^t r_i$ are the daily cumulative logarithmic returns and V_0 the initial price level on day $t = 0$.

The produced time series evolve according to $K_T \times K_V$ combinations of the volatility states $s; i \in \{1, \dots, K_V\}$ and trend states $S; i \in \{1, \dots, K_T\}$, with the two Markov processes progressing independently (Fig. 1).

IV. ANALYSIS AND RESULTS

A. Datasets

The aim of this study is to test the resilience of the CoinShares Gold and Cryptoassets Index (CGCI) [23] during different market regimes. For this purpose, we examine the return profile in relation to the volatility and trend states of its main market determinants. The CGCI is a low-volatility index that aims to maintain a prudent risk profile through diversification and regular rebalancing. The two uncorrelated risk factors driving the value of CGCI are the crypto-basket and the gold component, with the crypto-basket being an equally-weighted basket of 5 cryptoassets. The weighting among the crypto-basket and gold in the CGCI is computed through Eq. 1.

The estimation of the regime-switching models is performed on a sample of 1231 price observations, ranging from July 1st 2015 to May 31st 2020. Both the crypto-basket and gold time series correspond to the prices used for the calculation of the CGCI, with the crypto-basket price being calculated using historical tick-by-tick trade data provided by Kaiko and the gold price corresponding to the LBMA Gold Price PM data provided by ICE Benchmark Administration (IBA). For a detailed view on the full pricing methodology of both the crypto-constituents and the CGCI, readers can consult the official methodology document.

B. Trend Regime Estimation

We denote by $P_{c;t}$ the price of the crypto-basket component on day t and by $P_{g;t}$ the daily gold prices. Both time-series are expressed in USD. For the detection of the trend regimes we compute the weekly logarithmic returns $R_{c;t}$ and $R_{g;t}$ respectively. As mentioned in previous sections, we aim to model the weekly return series as an HMM with Gaussian distributions for the observations, while the regimes change according to a discrete Markov Process. A Durbin–Watson test reveals the presence of serial autocorrelation in the weekly logarithmic returns of both the crypto-basket ($d = 1:338$; $p < 0:001$) and gold ($d = 1:574$; $p < 0:001$) time-series. Therefore, in order to respect the output independence assumption, $R_{c;t}$ and $R_{g;t}$ are demeaned prior to the HMM fitting, using an AR(1) filter. The same test on the demeaned time series rejects autocorrelation (Crypto-basket: $d = 2:078$; $p = 0:7347$, gold: $d = 1:971$; $p = 0:4080$).

¹Available online by the index owner, CoinShares (Holdings) Limited (www.coinshares.com), and the Benchmark Administrator and Calculation Agent, Compass Financial Technologies (www.compassft.com).

Fig. 1. Market regime simulation procedure

In regards to the crypto-basket, taking into consideration that the sample dataset includes the historic 2017 price run, we assume the presence of $K_{E,T} = 4$ states, namely the bear, sideways, bull and outlier regimes. We expect the former states to reflect a downwards, stable, upwards and extreme intermediate trend respectively. The parameters of the four regimes are displayed in Table I. Kolmogorov-Smirnov tests across regimes confirm that observations, conditional to the prevailing latent state, are normally distributed, as expected by the HMM specification (Bear: $D = 0:140$, $p = 0:1260$; sideways: $D = 0:067$, $p = 0:7901$; bull: $D = 0:105$, $p = 0:3385$; outlier: $D = 0:073$, $p = 0:9993$). ANOVA test confirms that the sample's means differ significantly across regimes ($F = 90:423$, $p < 0:001$) (Fig. 2). Latent state 4 corresponds to the outlier regime, and describes the dynamics of the 2017 crypto-market price run, displaying extreme volatility and a mean similar to the one of the bull regime. Overall, the bear, sideways, bull and outlier states each constitute 26:46%, 35:02%, 29:96% and 8:56% of the 256 analysed weeks respectively.

For the evolution of gold's prices we follow the same approach and assume a simple HMM with $K_{E,T} = 3$ states that correspond to bear, sideways and bull market regimes (Fitting diagnostics in Table I). Response normality across the three regimes is confirmed through Kolmogorov-Smirnov tests (Bear: $D = 0:071$; $p = 0:9176$; sideways: $D = 0:073$; $p = 0:4155$; bull: $D = 0:124$; $p = 0:3736$). Welch's ANOVA test further confirms that the sample means differ significantly across all regimes ($F = 54:548$, $p < 0:001$) (Fig. 3) and the bear, sideways and bull latent states each constitute 20:23% and 57:59% of the 256 analysed weeks respectively.

Fig. 2. Welch's ANOVA – Crypto-basket Trend Regimes

Fig. 3. Welch's ANOVA – Gold Trend Regimes

C. Volatility Regime Estimation

For the detection of the volatility states in the historical dataset, we transform the constituent prices $P_{c,t}$ and $P_{g,t}$ to daily logarithmic returns $r_{c,t}$ and $r_{g,t}$. To account for different specifications in volatility dynamics of the daily returns, we express them in terms of a Markov-switching ARCH model, as specified in Section III-B. We assume that the ARCH residuals are Student-t distributed (as it is commonly used in practice and considered adequate for most financial applications [27]). Prior to the fitting process, we iteratively estimate both time-series with an AR(1) filter. A Durbin–Watson test on the demeaned time-series rejects autocorrelation (Crypto-basket: $d = 1.9708$, $p = 0.5080$; gold: $d = 1.9708$, $p = 0.4076$).

We assume that the crypto-basket time-series contains $K_{c,V} = 3$ states, namely the low, moderate and high volatility states. The ARCH parameters of the three regimes and the diagnostics of the expected log-likelihood maximisation are displayed in Table I. In this case our aim has been to specify a regime switching set of rules for the daily returns only based on volatility dynamics. Indeed, Levene's test rejects variance homogeneity across the low, moderate and high volatility regimes ($F = 209.52$; $p < 0.001$). Welch's ANOVA test further fails to reject mean equality across the specified regimes ($F = 1.199$; $p = 0.3037$). The identified low, moderate and high volatility latent states each constitute 24.88%, 64.59%, 13.53% of the 256 analysed weeks respectively.

Similarly for the gold prices, we assume a Markov-switching ARCH model with $K_{g,T} = 3$ volatility specifications and Student-t distributed innovations (Table I). Variance homogeneity is rejected ($F = 95.228$; $p < 0.001$) and mean equality across regimes cannot be rejected ($F = 0.428$; $p = 0.653$). The low, moderate and high volatility latent states each constitute 59.45%, 37.55%, 3.00% of the 256 analysed weeks.

D. Simulation Results

Given the fitted model parameters and a 7-year simulation horizon, we produce $N = 1000$ price paths for the crypto-basket and gold components, with the initial prices set to the last recorded prices of the historical dataset on May 31st 2020. We use the weighted risk contribution allocation scheme (Eq. 1) and produce $N = 1000$ corresponding paths for the CGCI, each one containing numerous of the possible combinations of $K_{c,T} \times K_{c,V} \times K_{g,T} \times K_{g,V} = 108$ possible combinations of its constituents' states. We are ultimately interested in examining both the overall performance of the three time series as well as the index risk-adjusted return profile in relation to its constituents' ongoing trend and volatility regimes.

First, we assess the performance of the CGCI, crypto-basket and gold in terms of their annualised Sharpe ratio, taking into consideration the entire dataset, regardless of the prevailing trend and volatility regimes. Overall, the diversified index strategy yields a superior risk profile, with an average annualised Sharpe ratio equal to 0.9641 (SD 0.3540) and a positive value for all simulation paths. The crypto-basket displays a less competitive performance and greater variability, yielding an average annualised Sharpe ratio of 0.3828 (SD = 0.4001). Gold's risk profile is more stable, with $M = 0.2819$ and $SD = 0.1982$. Welch's ANOVA test further confirms the highly significant difference in the Sharpe ratios ($F = 1422.037$; $p < 0.001$, Fig 4).

While the Sharpe ratio is a widely used risk–return assessment metric by investors, its main drawback is the fact that it takes into consideration only the first two moments of the returns distribution. One way to account for all return distribution moments is through the Omega ratio, as introduced by Keating and Shadwick [28]. The Omega ratio is defined as the probability-weighted ratio of gains over losses given a

HMM & MSGARCH Fitting				
	Crypto-basket Trend Regimes	Crypto-basket Volatility Regimes	Gold Trend Regimes	Gold Volatility Regimes
Transition Parameters Trend States: 1: Bear, 2: Sideways, 3: Bull, 4: Outlier (Crypto-basket only) Volatility States: 1: Low, 2: Moderate, 3: High	$p_{1,1} = 0:3887, p_{1,2} = 0:2049$ $p_{1,3} = 0:4064, p_{1,4} = 0:0000$ $p_{2,1} = 0:0000, p_{2,2} = 0:8007,$ $p_{2,3} = 0:1993, p_{2,4} = 0:0000$ $p_{3,1} = 0:4927, p_{3,2} = 0:0000$ $p_{3,3} = 0:4523, p_{3,4} = 0:0551,$ $p_{4,1} = 0:0041, p_{4,2} = 0:0000$ $p_{4,3} = 0:1468, p_{4,4} = 0:8491,$	$p_{1,1} = 0:9575, p_{1,2} = 0:0425$ $p_{1,3} = 0:0000, p_{2,1} = 0:0202$ $p_{2,2} = 0:9243, p_{2,3} = 0:0555$ $p_{3,1} = 0:0000, p_{3,2} = 0:2268$ $p_{3,3} = 0:7732$	$p_{1,1} = 0:0506, p_{1,2} = 0:1202$ $p_{1,3} = 0:8294, p_{2,1} = 0:9994$ $p_{2,2} = 0:0006, p_{2,3} = 0:0000$ $p_{3,1} = 0:1253, p_{3,2} = 0:2586$ $p_{3,3} = 0:6161$	$p_{1,1} = 0:9956, p_{1,2} = 0:0043$ $p_{1,3} = 0:0001, p_{2,1} = 0:0000$ $p_{2,2} = 0:9878, p_{2,3} = 0:0122$ $p_{3,1} = 0:0265, p_{3,2} = 0:0431,$ $p_{3,3} = 0:9304$
Parameters Trend States: $N (\kappa; \kappa)$ Volatility States: ARCH: $(a_0; a_1)$	$(\kappa_1 = 0:088, \kappa_1 = 0:060)$ $(\kappa_2 = 0:003, \kappa_2 = 0:027)$ $(\kappa_3 = 0:057, \kappa_3 = 0:069)$ $(\kappa_4 = 0:045, \kappa_4 = 0:202)$	$(a_{0,1} = 0:0005, a_{1,1} = 0:0001)$ $(a_{0,2} = 0:0022, a_{1,2} = 0:0000)$ $(a_{0,3} = 0:0084, a_{1,3} = 0:0027)$	$(\kappa_1 = 0:009, \kappa_1 = 0:008)$ $(\kappa_2 = 0:001, \kappa_2 = 0:025)$ $(\kappa_3 = 0:004, \kappa_3 = 0:011)$	$(a_{0,1} = 0:0000, a_{1,1} = 0:0001)$ $(a_{0,2} = 0:0001, a_{1,2} = 0:0001)$ $(a_{0,3} = 0:0003, a_{1,3} = 0:0001)$
Convergence AIC BIC Log-Likelihood	AIC : 471:5804 BIC : 390:0413 LL : 261:1289	AIC : 423:24121 BIC : 417:09959 LL : 2128:206	AIC : 143:6939 BIC : 138:7252 LL : 732:4695	AIC : 843:69141 BIC : 837:55076 LL : 423:04571

TABLE I: Regime Fitting Parameters and diagnostics

Fig. 4. Welch's ANOVA – Annualised Sharpe Ratios

Fig. 5. Omega Ratio Graphical Interpretation

specific level of threshold return r_t). We let X represent the observed returns of an asset and $F(r)$ denote the cumulative distribution function of returns. Given a selected target return r_t , the Omega ratio is given by:

$$\Omega(r_t) = \frac{\int_{r_t}^{\infty} (1 - F(r)) dr}{\int_{r_t}^{\infty} F(r) dr} \quad (10)$$

When r_t is set to be equal to zero, we get the gain-loss ratio of Bernardo and Ledoit [29]. The Omega ratio is used to rank investments similarly to the Sharpe ratio. The threshold is first chosen to a desired target level at will and investments can then be ranked accordingly, with higher values preferred to lower. The metric can further be extended to a portfolio optimisation strategy that aims to maximise quantity for a selected value of the returns threshold.

In this study we inspect the average annual return of each asset's cumulative probability component. Given a simulated price path $P_{i,1}, P_{i,2}, \dots, P_{i,1000}$, the annualised return is obtained through:

$$R_i = \left(\frac{P_{i,1000}}{P_{i,1}} \right)^{\frac{365}{1000}} - 1; \quad (11)$$

where n denotes the length of the simulated prices expressed in number of days. This yields 1000 values of annualised returns for the CGCI, crypto-basket and gold respectively. Fig. 5 presents the index cumulative distribution function and a graphical representation of the Omega ratio. Given a benchmark threshold = 0:01 (translating to a target return of 1% annually), the ratio is defined as the ratio of the blue over the red shaded area. In Fig. 6 the slope of the crypto-basket CDF appears visibly less steep compared to the gold and the index, with significantly heavier tails.

Fig. 6. Empirical Cumulative Distribution Functions

Fig. 7. Omega Ratios - Log Scale

We further investigate the ranking among the three components given different levels of expected probability. For large negative values of α , all three Omega ratios tends to infinity, while for large positive values of α , they tend to zero. Fig. 7 presents on a log scale the Omega ratio of the CGCI, the crypto-basket and gold as a function of the intersection point of the CGCI and gold Omega lines reveals that the index strategy is more appropriate for investors aiming for moderate positive returns up to 10.6% annually. The area for $\alpha < 0$ reveals that gold provides protection against negative market developments more effectively than the other two components. When it comes to the crypto-basket, it overperforms the index when the target is set higher than a 10.6% annual return, making it a more appropriate for investors with high risk tolerance. Examining the three components wealth preservation capabilities, a threshold of value $\alpha = 0$ yields an Omega ratio of $(0)_{CGCI} = 1:4018$ for the index, and $(0)_{CB} = 0:3056$ $(0)_{Gold} = 0:7747$ for the two components respectively.

E. Regime-Conditional Performance

Each of the simulated paths is then factorised according to the constituents' prevailing states and eventually split into 144 subsets of the original time-series. For each subset we compute the average monthly return for both the CGCI and its two constituents, across all 1000 paths. Given the high number of possible regime combinations, especially in the case of CGCI, we expect their duration to correspond to a few days. The average regime duration for the index is indeed approximately 3.44 days, with a total of $K_{c;T} K_{c;V} K_{g;T} = 108$ possible regime combinations. Accordingly, in the case of the crypto-basket, with $K_{c;T} = 4$ trend states and $K_{c;V} = 3$ volatility states, we have $K_{c;T} K_{c;V} = 12$ regime combinations with an average duration of 4.27 days. Finally for gold, we observe $K_{g;T} K_{g;V} = 9$ different regime combinations with an average duration of 9.15 days.

Given the short duration of regimes, we express the return profile of each regime in terms of monthly return instead of annual. If we denote with \mathbf{r}_j the vector of returns that correspond to regime j and simulation path i , we divide it in N_j partitions, each with a duration of 21 days, and obtain the average monthly return of regime j through:

$$R_j^i = \frac{1}{N_j} \prod_{t=1}^{N_j} (1 + r_{j;t}) - 1 \quad N_j = N; \quad i = 1; \dots; 1000 \quad (12)$$

We consider the risk-adjusted return profile of each component per regime combination. As mentioned before, a main drawback of the Sharpe ratio when assessing the performance, is the fact that it fails to consider the entire distribution of returns. An additional disadvantage of the Sharpe ratio is evident in this case because we need to consider the presence of negative returns during bear market periods. A negative Sharpe ratio is generally problematic to interpret when persistent (during market downturns) because a large amount of volatility, given negative excess returns, wrongfully insinuates that the examined performance is not as poor as expected. To this end, when evaluating the regime-conditional risk-adjusted returns, we use the Omega ratio as well. For each regime combination, we calculate the Omega ratio through 1000 values of α , corresponding to the 1000 simulated time-series. We use a benchmark value of $\alpha = 0:0008$ for the monthly returns, which roughly translates to a 1% annual return. The produced heatmap in Fig. 8 reveals high-risk regime combinations for the index diversification and rebalancing strategy. Moreover, it allows for a comparative analysis between the CGCI, the cryptoasset market and gold. For ease of exposition, and since we are interested in the three main trend regimes of bear, sideways and bull market periods, the outlier trend regime is omitted from the heatmap. Overall, the color gradient reveals that the main driver of the CGCI price is the crypto-basket component.

Fig. 8. Average Simulated Omega Ratios (Monthly returns, 0:0008)

The heatmaps in Fig. 9 display the variability of the CGCI Omega ratio across different regimes, in comparison with its two components. All Omega ratios are estimated using the same target return benchmark value, $r = 0:0008$. In bull market conditions, all three components perform better in low volatility regimes. The highest value for the CGCI Omega ratio, $\omega_{CGCI} = 23:8387$, is observed when both components experience low volatility and have an upward intermediate trend. In this case, the Omega ratios for the two index components are $\omega_{CB} = 21:1292$ and $\omega_{Gold} = 3:1493$ respectively. In regions that correspond to bear conditions for both the crypto-asset and gold markets, the CGCI performs best when both components are in a high-volatility state, with $\omega_{CGCI} = 0:3492$, $\omega_{CB} = 0:0788$ and $\omega_{Gold} = 0:2871$. This is in line with the primary goal of the index to protect against unfavorable market conditions and to control and benefit from risk through diversification and frequent rebalancing. Likewise, the regions that negatively affect the index performance the most lie around the bottom-left corner of the heatmap. Those regions correspond to periods when the components simultaneously experience bear markets with low volatilities, therefore, providing limited opportunities for successful diversification. In this case, the Omega ratios of the CGCI and the crypto-basket lie closer to zero, $\omega_{CGCI} = 0:0079$, $\omega_{CB} = 0:0053$. Gold seems to offer better protection in times of market decline, with $\omega_{Gold} = 0:0192$.

V. CONCLUSION

We have proposed a way to describe the dynamic behavior of the two market determinants of the CGCI. We have considered a Gaussian Hidden Markov Model to extract the intermediate trend regimes, given through the weekly logarithmic returns, and a Markov-switching ARCH approach to describe the variability of the conditional variances, through the demeaned daily logarithmic returns. Their combination attempts to produce a realistic set of simulated price paths for the index and its two risk factors, each one containing numerous combinations of the identified regimes. Taking into consideration the evolution of the index and its two risk factors, we report their overall performance in terms of Sharpe and Omega ratio and assess their suitability for investors, according to their individual return targets and willingness to take on certain levels of risk. The Sharpe ratios across the entire datasets demonstrate the overall superiority of the diversified approach when seeking exposure in the cryptoasset market. Computations of the Omega ratio for different values of the target return threshold reveal that the index is more suitable for wealth-preserving investors and investors who target moderate returns, up to 10.6% on an annual basis. For investors with higher risk tolerance, portfolios with cryptoasset components only are more appropriate, whereas gold is the best choice when seeking protection against periods of persistently declining markets.

