International Journal of
Financial Studies

MDPI

Article

Bitcoin as an Inflation Hedge and Safe Haven in the Euro Area: Evidence from Inflation and Financial Stress

Abdelkader Aguir * and Ramzi Farhani

MOFID LAB, Faculty of Economics and Management of Sousse, University of Sousse, Sousse 4000, Tunisia
* Correspondence: a.aguir@groupe-espi.fr

Abstract

This paper separately evaluates whether euro-denominated Bitcoin returns are associated with a recursive euro-area HICP forecast error and whether Bitcoin offers protection during European equity-market stress. Monthly inflation innovations are treated as ex post time-series forecast errors rather than announcement surprises. Controlled H1 coefficients are not statistically significant at t, t + 1 or t + 2. Full-sample daily conditional equity betas are positive, but moving-block bootstrap intervals include zero and exclusion of 2020 removes statistical significance at the 10%, 5% and 1% lower-tail thresholds. A one-day shift in equity-return dating also reverses the sign of the 1% estimate. CISS-based tests provide broader, but equity-inclusive, stress classifications. The findings do not establish a reliable short-run inflation hedge or a stable, timing-invariant safe haven for euro-area investors. They underscore the importance of sample composition, return synchronization and transparent replication. The analysis is reconstructed from archived data with documented source-vintage limitations.

Keywords: bitcoin; inflation hedge; safe haven; financial stress; euro area; portfolio diversification

JEL Classification: E31; E44; G11; G15

1. Introduction

Bitcoin has evolved from a decentralized payment experiment into a globally traded financial asset held by retail investors, institutional investors, investment funds, and corporations. Its predetermined supply, decentralized architecture, and independence from central-bank issuance have encouraged comparisons with gold and supported the narrative that Bitcoin may protect investors against monetary depreciation and macro-financial instability. These characteristics have become particularly relevant following the prolonged period of unconventional monetary policy, the liquidity expansion associated with the COVID-19 crisis, the subsequent inflation surge, and the rapid tightening of monetary policy implemented by major central banks.

The investment properties attributed to Bitcoin nevertheless remain controversial. Its limited supply may support its long-term store-of-value narrative, but scarcity alone does not establish that an asset provides effective protection against inflation or financial turmoil. Bitcoin is also characterized by exceptionally high volatility, large drawdowns, speculative trading, regulatory sensitivity, and an increasing connection with conventional financial markets. As institutional participation has expanded, Bitcoin may have become more responsive to changes in market liquidity, interest rates, and investors' risk appetite.

Check for updates

Academic Editor: Babacar Sène

Received: 10 September 2026
Revised: 9 November 2026
Accepted: 11 November 2026
Published: 15 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Consequently, the same process that has increased its financial legitimacy may also have weakened its ability to behave independently of traditional risky assets.

This distinction is essential because the concepts of store of value, inflation hedge, diversifier, and safe haven are not interchangeable. A store of value is expected to preserve purchasing power over a relatively long horizon. An inflation hedge should generate returns that are positively associated with actual or expected inflation. A diversifier is weakly correlated with another asset or portfolio under normal market conditions. By contrast, a safe-haven asset should remain uncorrelated or become negatively correlated with risky assets specifically during episodes of severe market stress (Baur & Lucey, 2010; Baur & McDermott, 2010). Therefore, observing an increase in the Bitcoin price during an inflationary or expansionary monetary period is insufficient to demonstrate safe-haven behavior.

The empirical literature provides conflicting evidence on these properties. Early studies emphasized Bitcoin's similarities with both gold and speculative currencies and suggested that it could contribute to portfolio diversification (Dyhrberg, 2016). Bouri et al. (2017) found that Bitcoin could hedge certain forms of global uncertainty, although its effectiveness varied across market conditions and distributional quantiles. More adverse conclusions emerged during the COVID-19 market collapse. Conlon and McGee (2020) showed that Bitcoin declined alongside equity markets during the crisis, casting doubt on its ability to protect portfolios against extreme financial losses. Similarly, Choi and Shin (2022) concluded that Bitcoin may respond positively to inflation shocks but fails to provide reliable protection against financial uncertainty. These results demonstrate that an inflation hedge is not necessarily a safe haven.

Blau et al. (2021) provide an early direct time-series examination of the relationship between Bitcoin and inflation. They document a positive association between Bitcoin returns and changes in expected inflation, while emphasizing that this evidence remains descriptive and does not establish a stable causal hedging relationship. Their analysis therefore supports the potential inflation-hedge narrative but leaves open the question of whether this property persists across currencies, samples, and market conditions.

Extending this debate to countries characterized by relatively high cryptocurrency adoption, Phochanachan et al. (2022) compare the inflation-hedging properties of Bitcoin with those of gold, equities, and oil under stable and turbulent market conditions. Their results indicate that hedge effectiveness is both asset-specific and regime-dependent. This evidence cautions against classifying Bitcoin as an unconditional inflation hedge and supports an empirical design that explicitly considers instability across macro-financial environments.

More recent evidence further challenges the idea of unconditional protection. Smales (2024) finds that cryptocurrencies hedge inflation only under a restricted set of circumstances and do not constitute a reliable alternative to gold when protection is most needed. Rodriguez and Colombo (2025) similarly show that Bitcoin's response to inflationary shocks is context-dependent and may weaken as the asset becomes increasingly integrated into mainstream financial markets. Taken together, these findings suggest that Bitcoin's hedging properties are nonlinear, time-varying, and dependent on the nature of the shock under consideration.

The distinction between inflation protection and crisis protection is especially important in the euro area. Bitcoin may respond to inflation conditions without protecting portfolios during equity-market downturns, while an asset with low average stock-market correlation may become positively exposed to equities precisely in stressed states. A joint framework is therefore needed to evaluate the two defensive properties using measures and frequencies appropriate to each hypothesis.

Despite the rapid expansion of this literature, three limitations remain. First, much of the evidence focuses on the United States and dollar-denominated Bitcoin prices, whereas the return relevant to a euro-area investor also reflects euro exchange-rate movements. Second, inflation hedging and financial safe-haven behavior are frequently conflated even though they concern different risks and horizons. Third, average correlations may conceal the dependence that emerges during severe equity losses and episodes of systemic stress.

The paper addresses these limitations through two deliberately separate tests. The first asks whether euro-denominated Bitcoin returns display a positive and statistically reliable association with new inflation information as captured by a recursive HICP forecast error. The second asks whether Bitcoin's exposure to European equities becomes non-positive during severe market losses or episodes classified by a broader CISS measure that itself includes equity-market information. Neither exercise is presented as a structural causal estimate; the contribution lies in disciplined measurement, frequency matching, and explicit decision rules for two claims that are often conflated.

To answer these questions, the study adopts a frequency-matched empirical framework covering January 2015 to December 2024. Monthly BTC/EUR returns are used for inflation analysis, while daily BTC/EUR and EURO STOXX 50 returns are used for conditional safe-haven tests. Gold is included as a financial-market control. The ECB Composite Indicator of Systemic Stress (CISS) incorporates equity-market information (Holló et al., 2012). To address potential mechanical overlap with equity-market stress, an additional robustness exercise classifies stress using the sum of three ECB-published contributions associated with money, bond, and foreign-exchange markets. This is an alternative non-equity-segment contribution proxy, not an officially reconstructed equity-excluded CISS; the published contributions may retain dependence on the full CISS correlation structure. The distinction between these two stress measures is explicit throughout.

The analysis has direct implications for investors, portfolio managers, and regulators. The relevant question is not whether Bitcoin occasionally appreciates during inflationary or uncertain periods, but whether it delivers systematic protection against purchasing-power erosion and during the episodes in which conventional portfolios experience their largest losses. The results support a cautious distinction between possible diversification benefits in normal markets and reliable defensive performance under stress.

The remainder of the paper is organized as follows. Section 2 reviews the theoretical and empirical literature and develops the research hypotheses. Section 3 presents the data and empirical methodology. Section 4 reports the baseline empirical results and diagnostic tests. Section 5 presents the robustness analysis and discusses the findings. Section 6 develops the portfolio and policy implications. Section 7 concludes.

2. Literature Review and Hypothesis Development

2.1. Bitcoin and Inflation: Theory and Empirical Evidence

Bitcoin's potential ability to hedge inflation is commonly linked to its predetermined and inelastic supply. Unlike fiat currencies, whose supply can be adjusted by central banks in response to macroeconomic conditions, the total number of bitcoins that can ultimately be created is limited by protocol to 21 million units (Nakamoto, 2008). This scarcity has encouraged the comparison of Bitcoin with gold and supported the argument that it may preserve purchasing power when monetary expansion reduces confidence in conventional currencies.

The theoretical basis for this argument can be understood through the interaction of scarcity, monetary credibility, and portfolio allocation. An unexpected increase in inflation reduces the real value of nominal monetary assets. Investors seeking to protect their purchasing power may consequently reallocate wealth toward assets whose supply cannot

be expanded discretionarily. If Bitcoin is perceived as such an asset, positive inflation shocks or upward revisions in inflation expectations should increase its demand and generate positive returns. This mechanism may be particularly relevant when inflation exceeds the central bank's target or when investors question the ability of monetary authorities to restore price stability.

However, Bitcoin's fixed supply is not sufficient to guarantee inflation-hedging effectiveness. Its market price also depends on speculative demand, financial-market liquidity, technological developments, regulatory announcements, and changes in global risk appetite. In addition, the limited supply of Bitcoin concerns its nominal quantity and does not establish a stable relationship between its price and the general price level. A credible inflation hedge must exhibit a sufficiently systematic positive response to inflationary shocks rather than merely appreciating over selected periods.

Empirical evidence remains inconclusive. Choi and Shin (2022) find that Bitcoin appreciates following positive inflation shocks, which supports its potential role as an inflation hedge. Nevertheless, the same study shows that Bitcoin does not provide reliable protection against financial uncertainty. This result highlights the need to distinguish protection against inflation from protection against market turmoil.

Other studies report substantially more conditional results. Smales (2024) finds that cryptocurrency returns are positively associated with inflation expectations only under restricted circumstances, particularly when inflation remains relatively low. Cryptocurrencies respond negatively to some inflation announcements and do not consistently outperform gold as inflation hedges. Rodriguez and Colombo (2025) similarly conclude that Bitcoin's inflation-hedging capacity is context-specific and may diminish as the asset becomes more closely integrated with conventional financial markets.

Recent evidence based on derivatives markets reinforces this state-dependent interpretation. Liu and Valcarcel (2024) report a positive relationship between cryptocurrency-futures prices and inflation expectations over their full sample, but the sign turns negative between the collapses of Luna and FTX. This reversal is important because it shows that an apparent inflation hedge can disappear when crypto-specific stress changes margin constraints, liquidity, and investor composition. The result also cautions against transferring evidence from US-dollar futures directly to spot Bitcoin returns experienced by euro-area investors.

Taken together, the literature does not support a purely scarcity-based inference. Although Blau et al. (2021) and Choi and Shin (2022) provide evidence consistent with some inflation-hedging capacity, Phochanachan et al. (2022), Smales (2024), Liu and Valcarcel (2024), and Rodriguez and Colombo (2025) show that this property depends on the inflation measure, market conditions, institutional setting, and sample period. A fixed protocol supply may support a narrative of protection, but it does not guarantee a stable empirical response to inflation. These conflicting findings motivate the use of a euro-area-specific inflation measure and require the conclusions to remain conditional on the selected proxy and observation frequency.

Differences across studies may originate from the choice of inflation measure, sample period, data frequency, currency denomination, and information set. Realized inflation is the observed year-on-year HICP rate. A change in inflation expectations measures a revision in a survey or market-implied forecast. An announcement surprise is the released value minus the market consensus recorded immediately before publication. By contrast, a model-based HICP innovation is a forecast error relative to an explicitly specified time-series information set. These objects cannot be treated as interchangeable. Because a consistent real-time consensus series is not available for the full euro-area sample, the baseline analysis uses a

recursive autoregressive innovation and interprets it narrowly as inflation unexplained by past HICP dynamics, not as new information received by financial markets.

The euro area provides a relevant setting for this analysis. The period following 2015 includes unconventional monetary-policy measures, negative interest rates, the COVID-19 liquidity expansion, the sharp inflation increase of 2021–2022, and the subsequent tightening cycle. These distinct episodes make it possible to examine whether Bitcoin's response is stable or limited to particular inflationary environments. Moreover, using Bitcoin returns denominated in euros captures the actual performance experienced by a euro-area investor and avoids treating dollar-denominated returns as directly representative of European purchasing power. The monetary-policy environment provides relevant institutional context for interpreting inflation-related asset returns, although it does not itself establish Bitcoin's hedging properties. Aguir (2013) examines the relationship between de facto central-bank independence and inflation using cross-country evidence, while Aguir et al. (2017) study inflation targeting and macroeconomic stability in emerging and developing economies following the global financial crisis. These studies highlight the roles of monetary credibility and policy frameworks in inflation dynamics. Their institutional focus differs from the present analysis, which tests the association between BTC/EUR returns and euro-area HICP forecast errors.

The post-COVID inflationary environment also raises questions about the effectiveness of monetary-policy frameworks. Aguir (2025) compares inflation-targeting and alternative monetary regimes across 22 emerging economies and reports differences in macroeconomic performance during the pandemic period. Although this evidence concerns emerging economies rather than the euro area, it provides institutional context for the inflationary environment examined here; it does not establish whether Bitcoin hedges inflation risk.

A related euro-area contribution by Aguir and Fragny (2026) studies the transmission of cryptocurrency adoption, volatility, and regulatory developments to real-estate investment and transactions using a Bayesian structural VAR and quarterly observations. Its evidence of financial transmission beyond cryptocurrency markets reinforces the relevance of investor confidence, liquidity, and regulatory conditions in interpreting Bitcoin's risk exposures. Because the outcome variables and identification strategy differ, this study is used as contextual evidence rather than as a direct test of Bitcoin's inflation-hedging or safe-haven properties.

The first hypothesis is therefore formulated as follows:

H1. *Bitcoin provides short-run inflation protection for euro-area investors: euro-denominated Bitcoin returns respond positively to a model-based euro-area HICP innovation.*

H1 is supported only if the model-based HICP innovation has a positive, statistically reliable coefficient under the reported timing specifications and fixed-AR(1) sensitivity check. Unreported alternative inflation definitions and subperiod interactions are not interpreted as estimated robustness evidence. The recursive forecast error is not a market-consensus announcement surprise.

2.2. Bitcoin as a Safe Haven During Equity-Market Downturns

The concepts of hedge, diversifier, and safe haven describe different relationships between assets. Following Baur and Lucey (2010) and Baur and McDermott (2010), an asset can be considered a hedge when it is, on average, uncorrelated or negatively correlated with another asset or portfolio. A diversifier may remain positively correlated with that portfolio, provided that the correlation is relatively low. A safe haven represents a more demanding property: the asset must remain uncorrelated or become negatively correlated with the reference market specifically during periods of extreme losses or financial stress.

This definition has an important empirical consequence. Bitcoin cannot be classified as a safe haven merely because its average correlation with equity markets is low. The relevant question is whether this relationship remains non-positive in the lower tail of the equity-return distribution, when portfolio protection is most valuable. A conventional linear model estimated over the full sample may obscure this conditional behavior.

Several characteristics could theoretically support Bitcoin's safe-haven role. It operates outside the conventional banking system, is traded continuously, can be transferred internationally, and is not directly issued by a government or central bank. These characteristics may attract investors when confidence in traditional financial institutions deteriorates. Bitcoin could therefore benefit from a flight-to-alternative-assets mechanism similar, at least conceptually, to the flight toward gold.

Nevertheless, other characteristics point in the opposite direction. Bitcoin exhibits extreme volatility, lacks an intrinsic cash flow, and remains highly sensitive to market sentiment, leverage, and liquidity conditions. During episodes of market panic, investors may sell volatile and liquid assets to cover losses, satisfy margin requirements, or reduce portfolio risk. Under these conditions, Bitcoin may behave as a speculative risk asset rather than as a defensive asset.

The empirical literature reflects this ambiguity. Dyhrberg (2016) identifies some similarities between Bitcoin, gold, and the US dollar and suggests that Bitcoin may have useful portfolio-management properties. Bouri et al. (2017) find that its hedging and safe-haven performance depends on the market considered, the data frequency, and the severity of market movements. The evidence is therefore conditional rather than universal.

The COVID-19 crisis provided a particularly demanding test. Conlon and McGee (2020) show that Bitcoin declined alongside the US equity market during the initial pandemic-related collapse, increasing rather than reducing portfolio downside risk. Choi and Shin (2022) also find that Bitcoin falls in response to financial-uncertainty shocks. More recent international evidence generally suggests that Bitcoin is, at best, a weak and episodic safe haven, whereas gold exhibits more stable defensive properties.

Complementing the evidence on pandemic-period market stress, Dardouri et al. (2023) examine the relationship between COVID-19 transmission and the prices of Bitcoin, Ethereum, and Dogecoin using an autoregressive distributed lag framework over January 2020–March 2021. They report long- and short-run links between pandemic indicators and cryptocurrency prices. Their findings support treating the pandemic as a distinct source of crypto-market vulnerability, although their price-level and pandemic-transmission specifications do not directly establish Bitcoin's conditional safe-haven beta relative to euro-area equities.

Post-pandemic studies do not resolve the debate; they make its conditional nature more visible. Tarchella et al. (2024), using G7 equity markets and nonlinear GARCH models, find that cryptocurrencies can provide safe-haven benefits during COVID-19 and that Bitcoin performs particularly well for several European and UK equity markets under stress. Jin and Tian (2024) likewise document an enhanced safe-haven role during the Silicon Valley Bank collapse. These findings imply that Bitcoin may protect portfolios in specific, relatively bounded episodes without qualifying as a universal safe haven.

A second strand reaches a more skeptical conclusion. Yae and Tian (2024) characterize Bitcoin as a highly volatile safe-haven asset whose apparent diversification benefit is sensitive to time-varying dependence and the horizon over which protection is evaluated. Using multifractal cross-correlation methods across several asset classes, Acikgoz (2025) finds that neither Bitcoin nor gold consistently qualifies as a hedge or safe haven, apart from a limited risk-reduction role for the gold-bond pair. Evidence from the Russia-Ukraine

and Israel-Palestine conflicts also indicates that safe-haven performance depends on the pre-crisis correlation structure and the event window (Pastén-Henríquez et al., 2025).

This recent evidence identifies three sources of disagreement: the crisis considered, the equity-market benchmark, and the econometric definition of tail protection. The relevant test is whether Bitcoin's conditional equity exposure becomes non-positive across lower-tail thresholds and alternative high-stress states. The aggregate CISS includes equities; a further test based only on published contributions associated with non-equity market segments is reported as a sensitivity exercise, without claiming that those contributions are statistically independent of equities.

The empirical literature has developed three main approaches to evaluating hedge and safe-haven properties. The first relies on dynamic conditional correlation models, particularly the DCC-GARCH framework introduced by Engle (2002). This approach captures changes in conditional covariance and correlation over time and is therefore well suited to identifying periods in which the dependence between Bitcoin and conventional assets strengthens or weakens. Several studies employ DCC-type models to assess whether Bitcoin behaves as a hedge, diversifier, or safe haven under changing volatility conditions. However, a low or negative conditional correlation does not necessarily establish that Bitcoin provides protection specifically on days characterized by extreme equity losses.

A second approach uses quantile regression (Koenker & Bassett, 1978) to examine whether the relationship between Bitcoin and conventional financial assets varies across the conditional-return distribution. Quantile methods are particularly useful because average regressions may conceal asymmetric dependence in bearish and bullish market states. Bouri et al. (2017) show that Bitcoin's hedging properties depend on the quantile considered, while Wu et al. (2019) combine GARCH modeling and quantile regression to compare Bitcoin and gold as hedges against economic-policy uncertainty. More recently, Bahloul et al. (2023) jointly use DCC-GARCH and quantile models to assess the hedge and safe-haven properties of Bitcoin and gold across equity markets.

The third approach originates from the conditional safe-haven framework of Baur and McDermott (2010). It interacts benchmark-market returns with indicators identifying extreme adverse market observations. Unlike DCC-GARCH, which models the evolution of conditional correlations, and conventional quantile regression, which evaluates relationships at different points of a conditional-return distribution, the tail-interaction specification directly tests whether the candidate asset's exposure to the benchmark becomes non-positive when the benchmark itself experiences its largest losses. Because the present study defines a safe haven from the perspective of a euro-area equity investor, this benchmark-tail criterion constitutes the baseline empirical approach.

For euro-area investors, the relevant benchmark should be a broad European equity index, such as the EURO STOXX 50, rather than the S&P 500 alone. This distinction matters because European financial stress, ECB policy decisions, exchange-rate dynamics, and the composition of European equity markets differ from those observed in the United States. Bitcoin's safe-haven capacity should consequently be evaluated against European equity losses and euro-area financial-stress indicators.

The second hypothesis is formulated as follows:

H2. *Bitcoin acts as a safe haven for euro-area equity investors: the conditional relationship between euro-denominated Bitcoin returns and European equity returns becomes non-positive during extreme equity-market downturns and periods of elevated financial stress.*

A significantly negative relationship supports a strong safe-haven classification. A relationship that is statistically indistinguishable from zero is compatible with, but does not establish, a weak safe-haven classification. A significantly positive relationship during

severe equity losses rejects both strong and weak safe-haven classifications and supports the interpretation of Bitcoin as a risky or speculative asset.

Gold returns enter the Bitcoin regressions as a control. This study does not estimate matched gold-dependent-variable inflation-hedge or safe-haven regressions; therefore, it makes no empirical claim that Bitcoin provides more or less protection than gold under the present design.

2.3. Synthesis and Research Gap

The literature does not support treating Bitcoin's defensive properties as a single, unconditional characteristic. An asset may respond positively to inflation without protecting portfolios during market crashes, while low average correlation may disappear precisely when systemic stress is elevated. Inflation hedging and safe-haven behavior must therefore be evaluated separately with definitions, frequencies, and decision rules suited to each claim.

This study develops an integrated euro-area framework based on two hypotheses. H1 evaluates the association between Bitcoin returns and a recursive HICP innovation. H2 tests protection during extreme European equity-market losses, aggregate CISS stress, and an alternative classification based on money-, bond-, and foreign-exchange-market contributions. The aggregate CISS incorporates equity-market information, and the alternative contribution-based proxy is not an officially recomputed CISS excluding equities.

The updated literature reveals a clear empirical gap. Many studies examine either inflation hedging or crisis protection, often from a US-dollar perspective, and frequently infer safe-haven status from average correlations. This paper evaluates both defensive claims for a euro-area investor while keeping their definitions, frequencies, and decision rules separate.

3. Data and Empirical Methodology

3.1. Empirical Design and Sample Period

The empirical analysis covers the period from January 2015 to December 2024. This period encompasses markedly different monetary and financial environments in the euro area, including the Eurosystem's asset-purchase programs, the negative-interest-rate period, the COVID-19 market collapse, the inflation surge of 2021–2022, and the subsequent monetary-policy tightening cycle.

The analysis is based on two complementary datasets. First, a monthly dataset is used to examine whether Bitcoin protects euro-area investors against the relationship between Bitcoin returns and a recursive model-based HICP inflation forecast error (rather than observed market inflation expectations). Second, a daily dataset is constructed to assess Bitcoin's safe-haven properties during severe European equity-market losses and periods of systemic financial stress.

This multi-frequency structure is necessary because inflation and financial stress operate at different frequencies. Converting annual or quarterly macroeconomic variables into artificial monthly observations would create spurious precision, while aggregating daily financial-market movements would conceal the extreme market conditions relevant to safe-haven behavior.

All Bitcoin prices are denominated in euros to reflect the returns experienced by a euro-area investor. Daily Bitcoin observations are merged with European equity, gold, exchange-rate, and financial-stress series using common trading dates. Because the Bitcoin daily close and the European equity-market close are not perfectly synchronized in clock time, calendar-date matching is treated as an approximate daily alignment rather than as exact interval synchronization. Daily Bitcoin observations are merged with European

equity, gold, exchange-rate, and financial-stress series using common trading dates. Monthly observations are aligned chronologically and used only in the inflation models.

3.2. Variable Definitions

The principal variables, frequencies, sources, and transformations are presented in Table 1.

Table 1. Core variables, frequencies, sources, and transformations.

Variable	Symbol	Definition	Frequency	Transformation	Source
Bitcoin price in euros	P^{BTc}	Bitcoin market price expressed in EUR	Daily	Daily or end-of-month price	Yahoo Finance BTC-EUR (historical/chart series)
Bitcoin return	r^{BTc}	Continuously compounded return (%)	Daily and monthly	$100[\ln(Pt) - \ln(Pt-1)]$	Authors' calculation
Bitcoin trading volume	VOL^{BTc}	BTC-EUR trading-volume field; denomination/units require vendor verification	Daily	Raw volume; not used in controlled monthly H1 regressions	Yahoo Finance BTC-EUR
Euro-area HICP	HICP	All-items Harmonised Index of Consumer Prices	Monthly	Original index; annual log change	Eurostat prc_hicp_midx; CP00
Realized inflation	π	Annual euro-area HICP inflation (%)	Monthly	$100[\ln(HICP_t) - \ln(HICP_{t-12})]$	Authors' calculation
Model-based HICP innovation	INFShock	Annual HICP forecast error relative to lagged inflation dynamics	Monthly	Recursive one-step-ahead AR forecast error; not a market-consensus surprise	Authors' calculation
EURO STOXX 50	$P^{E}Q$	Euro-area blue-chip price index in EUR	Daily	Closing price on common trading dates	Yahoo Finance ^STOXX50E (EURO STOXX 50 price index; SX5E family)
European equity return	$r^{E}Q$	Continuously compounded equity return (%)	Daily	$100[\ln(P_t) - \ln(P_{t-1})]$	Authors' calculation
Gold futures price (USD)	P^{GOLD}	Yahoo Finance GC=F, USD per troy ounce; futures, not spot	Daily	Quoted USD futures price (no EUR conversion in workbook)	Yahoo Finance GC=F
Gold futures return (USD)	r^{GOLD}	Log return on GC=F USD futures price	Daily/monthly	$100[\ln(GC=F_t) - \ln(GC=F_(t-1))]$	Authors' calculation from Yahoo Finance GC=F
Systemic financial stress	CISS	Composite Indicator of Systemic Stress	Daily	Level; standardized level; tail dummy	ECB CISS.D.U2.Z0Z.4F.EC.SS_CI.IDX
ECB deposit facility rate	DFR	ECB deposit facility rate (%)	Daily/monthly	Level and change in percentage points	ECB Data Portal
USD/EUR exchange rate	FX	US dollars per euro	Daily	100 log difference	ECB EXR.D.USD.EUR.SP00.A

Bitcoin prices and trading-volume observations are obtained from the Yahoo Finance BTC-EUR historical series. Bitcoin prices are quoted directly in euros; no BTC/USD conversion is applied. Daily observations are indexed by the vendor's reported calendar dates. The precise historical intraday closing convention is not identified in the available series metadata.

The European equity benchmark is the EURO STOXX 50 price index, accessed through Yahoo Finance under ticker ^STOXX50E. The series represents index-price changes rather than total returns. Accordingly, the empirical results concern the price-index benchmark and should not be interpreted as estimates for a dividend-reinvested equity portfolio.

Index identifier and trading calendar. The EURO STOXX 50 (SX5E) observations are matched with Bitcoin and other financial variables on common available trading dates. European market holidays are excluded rather than filled with artificial zero returns. Because vendor calendar dates do not establish identical intraday closing intervals, the lead/lag tests in Table A1 assess sensitivity to observation timing.

Gold enters the regressions as a financial-market control, not as a separate safe-haven benchmark. The gold series is Yahoo Finance GC=F, quoted in US dollars per troy ounce and representing gold futures rather than LBMA spot gold. The daily and monthly gold-return variables are logarithmic returns on this USD-denominated futures series. A EUR-denominated futures-price proxy is defined by:

$$P^{\mathrm{GOLD,\,EUR}}{}_{t} = P^{\mathrm{GOLD,\,USD}}{}_{t} / FX^{\mathrm{USD}}{}_{/\mathrm{EUR}}{}_{t}$$

Here FXUSD/EUR denotes US dollars per euro. Thus, P_gold, futures, EUR=P_gold, futures, USD/FXUSD/EUR and, on matching dates, r_gold, futures, EUR=r_gold, futures, USD − r_FX. The baseline gold control is the USD-denominated futures return; Table A6 reports the sensitivity of H1 to the EUR-converted futures-return proxy. This conversion does not constitute an LBMA spot-gold series.

Gold-return measurement and currency consistency. The gold control in the supplied H1 and H2 working data is r_GC=F, USD = 100[ln(GC=F_t) − ln(GC=F_(t − 1))]. A euro-area futures-price proxy can be computed on matching dates as r_GC=F, EUR=r_GC=F, USD − 100[ln(USD/EUR_t) − ln(USD/EUR_(t − 1))]. The daily H2 regressions include the USD/EUR return as a separate control, so a consistent linear reparameterization from USD to EUR gold-futures returns (retaining the FX regressor) leaves the fitted values and conditional equity betas unchanged when observations and controls are identical; it changes the interpretation of the gold and FX coefficients. The monthly H1 specifications do not contain the FX return as an additional regressor, and this invariance does not apply to them. Neither construction establishes the behavior of EUR spot gold.

3.3. Monthly Dataset for the Inflation Hedge Hypothesis

The inflation analysis covers monthly observations from January 2015 to December 2024. Bitcoin prices are quoted in euros, while the gold-futures control is quoted in US dollars. Monthly returns are computed as successive log price differences, expressed in percent. The baseline H1 model includes USD gold-futures returns, and the currency-converted alternative is reported in Table A6.

$$r^{\mathrm{BTc}}{}_{\mathrm{m}} = 100[ln(P^{\mathrm{BTc}}{}_{\mathrm{m}}) - ln(P^{\mathrm{BTc}}{}_{\mathrm{m-1}})]$$

Euro-area inflation is measured from the all-items HICP:

$$\pi_{\mathrm{m}} = 100[ln(HICP_{\mathrm{m}}) - ln(HICP_{\mathrm{m-12}})]$$

Realized inflation may already be anticipated by financial markets. The baseline inflation variable is therefore constructed as a recursively generated one-step-ahead forecast error rather than as the first difference of annual inflation. Let π_t denote the year-on-year euro-area HICP inflation rate, calculated as the twelve-month log difference of the all-items HICP index.

The model-based HICP innovation is defined as the difference between observed annual inflation and its one-step-ahead prediction based exclusively on lagged inflation information.

To avoid unstable forecasts from very short recursive samples, no innovation is generated until 36 usable monthly year-on-year inflation observations are available. Thereafter, the autoregressive model is re-estimated recursively at each date. The lag order is re-selected at every recursive step using the Bayesian Information Criterion (BIC), allowing $p = 1, \ldots, 12$. For comparability of the BIC across lag orders, candidate models at a given recursive date are evaluated over a common estimation span.

Annual euro-area inflation is defined as 100 times the twelve-month logarithmic change in the HICP index. The innovation is obtained from recursive one-step-ahead AR(p) forecasts with an initial window of 36 annual-inflation observations. At each forecast origin, BIC selects p from 1 to 12 on a common AR(12)-compatible estimation window. For the 72 forecast origins, the selected orders are AR(1) in 49 cases, AR(2) in 14, AR(3) in 3, and AR(8) in 6. The resulting innovation has a mean of −0.0311 and a sample standard deviation of 0.5476 percentage points. This log-index measure differs from the simple year-on-year percentage-change field.

A positive INFShock_t indicates that inflation exceeds the value predicted from its own lagged dynamics. This variable is an econometric time-series innovation. It is not a market-consensus announcement surprise and should not be interpreted as an event-study measure of new information arriving at the HICP release. The HICP observations in the replication database correspond to the revised Eurostat series rather than reconstructed real-time vintages; consequently, the innovation may incorporate later statistical revisions.

The timing of the Bitcoin return is treated explicitly. Because the final HICP for reference month t is normally published during month t + 1, the contemporaneous regression linking INFShock_t to r^BTC_t is interpreted only as an ex post hedging-covariance specification. We therefore report a lagged-information specification using r^BTC_{t + 1}. As an additional conservative timing check, we also report r^BTC_{t + 2}, for which the final HICP for month t is available before the return month begins. None of these regressions is described as an HICP announcement-surprise regression.

Information set and interpretation. The month-t HICP innovation is calculated from realized inflation and lagged inflation observations. As the complete HICP observation is published after the reference month, the contemporaneous model measures an ex post monthly association, not a reaction to a market announcement. Specifications for returns at t + 1 and t + 2 examine alternative information timing without identifying announcement-day effects.

$$\pi_t = 100[\ln(HICP_t) - \ln(HICP_{t-12})]$$

$$\pi_t = \alpha_t + \sum_{j=1}^{p_t} \rho_{j,t}\pi_{t-j} + u_t$$

$$INFShock_t = \pi_t - \hat{\pi}_{t|t-1}$$

The monthly variables, their transformations, and their roles in the inflation-hedge specification are summarized in Table 2.

Table 2. Descriptive Statistics of Monthly Variables.

Role	Variable	Frequency	Transformation
Dependent variable	Bitcoin return in EUR	Monthly	End-of-month log return
Main explanatory variable	Model-based HICP innovation	Monthly	Recursive AR forecast error after a 36-observation initialization; BIC-selected lag order (1–12) re-estimated recursively; not an announcement surprise
Alternative inflation measure	HICP annual inflation	Monthly	Twelve-month log difference

Table 2. *Cont.*

Role	Variable	Frequency	Transformation
Market control	EURO STOXX 50 return	Monthly	End-of-month log return
Financial-stress control	CISS	Monthly	Monthly average and first difference
Monetary control	ECB deposit facility rate	Monthly	End-of-month level and change
Financial-market control	Gold-futures return (USD; Yahoo GC=F)	Monthly	End-of-month log return

GDP per capita, government expenditure, and M2 are excluded from the baseline model. GDP per capita and government expenditure are not genuinely available monthly, while interpolation would manufacture observations without adding economic information. M2 may be introduced only in an additional specification because it is highly persistent and requires separate stationarity and potential cointegration tests.

3.4. Daily Dataset for the Safe Haven Hypothesis

The safe-haven analysis uses daily observations on Bitcoin, the EURO STOXX 50, gold, and the CISS. Bitcoin weekend observations are removed, and the remaining series are merged by common trading date. If a conventional market is closed, the corresponding date is excluded rather than filling the missing price mechanically.

The daily Bitcoin series is based on a consistent UTC closing convention, whereas the EURO STOXX 50 is observed at the European market close. The available daily database does not contain intraday Bitcoin observations sampled exactly at the European close. Consequently, matching the two series by calendar date does not imply exact clock-time synchronization. The same-day specification is therefore interpreted as a daily comovement measure. To assess whether the safe-haven classification is driven by this timing mismatch, we additionally estimate one-trading-day lead/lag specifications and cumulative-return models over two to five trading days.

Daily returns are calculated as:

$$r^{i}_{t} = 100[ln(P^{i}_{t}) - ln(P^{i}_{t-1})]$$

where i denotes Bitcoin, European equities, or gold. Extreme European equity-market conditions are defined using the empirical distribution of EURO STOXX 50 returns:

$$D^{TAIL}_{t} = 1 \; if \; r^{E}Q_{t} \leq q\tau, \; \tau \in \{0.10, \; 0.05, \; 0.01\}$$

The 10%, 5%, and 1% thresholds identify moderate, severe, and extreme equity-market downturns. A complementary financial-stress dummy equals one when the CISS is above its 90th percentile. The 95th percentile is used as an alternative threshold.

Bitcoin qualifies as a strong safe haven if its relationship with European equities is significantly negative during extreme downturns. A relationship that is statistically indistinguishable from zero is compatible with, but does not prove, a weak safe-haven classification. A significantly positive relationship rejects both classifications. Gold is included as a control, but its own safe-haven status is not estimated in this study.

Table 3 summarizes the daily variables and stress thresholds used to evaluate the safe-haven hypothesis.

Table 3. Daily Variables for the Safe Haven Analysis.

Role	Variable	Frequency	Transformation or Threshold
Dependent variable	Bitcoin return in EUR	Daily	Log return
Principal market variable	EURO STOXX 50 return	Daily	Log return

Table 3. *Cont.*

Role	Variable	Frequency	Transformation or Threshold
Financial-market control	Gold-futures return (USD; Yahoo GC=F)	Daily	Log return
Systemic-stress indicator	CISS	Daily	Level and standardized level
Moderate downturn	EURO STOXX 50 lower tail	Daily	Return below 10th percentile
Severe downturn	EURO STOXX 50 lower tail	Daily	Return below 5th percentile
Extreme downturn	EURO STOXX 50 lower tail	Daily	Return below 1st percentile
High systemic stress	CISS upper tail	Daily	Above 90th or 95th percentile
Currency control	USD/EUR exchange rate	Daily	Log change
Dynamic control	Lagged Bitcoin return	Daily	Previous matched trading-day BTC/EUR return

3.5. Subperiod Definitions

The three macroeconomic subperiods in Table 4 provide institutional and descriptive context. No H1 interaction estimates for these subperiods are reported; therefore, no empirical claim of coefficient stability across these regimes is made.

Table 4. Baseline classification of macroeconomic subperiods.

Subperiod	Period	Economic Interpretation
Pre-pandemic low-rate period	January 2015–February 2020	Negative policy rates and asset purchases before the pandemic
Pandemic and PEPP period	March 2020–June 2022	Pandemic shock, PEPP purchases, and exceptionally abundant liquidity
Inflation-tightening period	July 2022–December 2024	Rapid policy-rate increases followed by the beginning of normalization

3.6. Data Cleaning and Alignment

The underlying financial and macroeconomic observations were obtained from the data sources identified above; AI-generated outputs were not treated as independent data or empirical evidence. The authors retain responsibility for independently verifying the computations, source metadata, and reported estimates.

The series are transformed at their observed frequencies. Daily financial observations are merged on common trading dates; non-trading dates and missing observations are excluded from complete-case regressions. The following conventions determine the estimation samples:

Bitcoin and equity prices are matched by reported calendar date rather than identical intraday closing time. Table A1 evaluates the effect of shifting the equity-return series by one matched trading observation, and Table A2 considers overlapping cumulative returns over two to five matched observations.

Weekend Bitcoin observations are excluded from the daily cross-asset dataset.
Conventional market series are merged using common trading dates.
Missing prices are not linearly interpolated.
Monthly financial variables use the last common trading observation of the month, except the CISS, for which the monthly mean is also used.
Returns are calculated only after prices have been denominated in euros.
Extreme-market thresholds are estimated from the empirical distribution of the estimation sample.
Lagged variables are constructed after chronological sorting and merging.

No annual or quarterly macroeconomic series is converted artificially into monthly data.

Tables A1 and A2 examine, respectively, one-observation lead/lag conventions and overlapping cumulative-return horizons. These tests characterize timing and horizon sensitivity; they do not identify exchange-level intraday synchronization.

The descriptive appendix reports sample size, mean, standard deviation, quantiles and range, supplemented by skewness, excess kurtosis and Jarque–Bera statistics in Table A11. Unconditional correlation matrices are reported separately. The recursively estimated HICP innovation is presented graphically in the diagnostic results section. Additional time-series plots and a complete monthly residual-diagnostic battery are not reported and are not claimed as evidence for the hypotheses.

3.7. Econometric Specifications

3.7.1. Inflation-Hedge Model

Hypothesis H1 is evaluated at the monthly frequency. The dependent variable is the BTC/EUR logarithmic return. The main explanatory variable is the recursive model-based HICP innovation defined in Section 3.3. Three timing conventions are distinguished.

Model A is the contemporaneous specification. Because the complete HICP observation for month t is not available throughout month t, this regression is interpreted only as an ex post hedging-covariance test and not as an investable response to newly released inflation information.

Model B is the principal lagged-information specification requested to align the information set with the return timing. It relates INFShock_t to the subsequent Bitcoin return in month t + 1.

Model C is a conservative timing robustness test using the Bitcoin return in month t + 2. Since the final HICP for reference month t is normally released during month t + 1, this specification ensures that the complete month-t HICP information precedes the entire return month.

For each timing convention, a parsimonious specification is reported first, followed by the controlled specification. The controlled model includes the lagged Bitcoin return, the contemporaneous EURO STOXX 50 return, the gold return, the change in CISS, and the change in the ECB deposit facility rate measured at the inflation reference month t for all three response horizons (t, t + 1, and t + 2). Newey–West HAC standard errors use four monthly lags. The coefficient β is interpreted as a conditional association, not as a structural causal effect. H1 requires a positive β; failure to obtain a positive statistically supported coefficient is recorded as lack of support for H1 rather than evidence of a negative causal inflation effect.

Controlled H1 specification. Each regression includes the recursive HICP innovation, lagged BTC/EUR return, monthly EURO STOXX 50 return, USD-denominated GC=F gold-futures return, the change in monthly mean CISS, and the change in the ECB deposit facility rate. All controls retain reference-month-t dating when the dependent Bitcoin return is advanced to t + 1 or t + 2. Trading volume is excluded. Newey–West HAC standard errors use four monthly lags. The contemporaneous, t + 1 and t + 2 regressions contain 72, 71 and 70 observations, respectively. The fixed-AR(1) alternative uses the same controls.

$$r_t^{BTC} = \alpha + \beta INFShock_t + \gamma' X_t + \varepsilon_t$$

$$r_{t+1}^{BTC} = \alpha + \beta INFShock_t + \gamma' X_t + \varepsilon_{t+1}$$

$$r_{t+2}^{BTC} = \alpha + \beta INFShock_t + \gamma' X_t + \varepsilon_{t+2}$$

$$X_t = \left[r_{t-1}^{BTC},\ r_t^{SX5E},\ r_t^{Gold},\ \Delta CISS_t,\ \Delta DFR_t \right]'$$

3.7.2. Safe-Haven Model

Hypothesis H2 is tested with the daily trading-day panel. Following the conditional safe-haven logic, Bitcoin must exhibit a non-positive sensitivity to European equities during extreme stock-market losses. For each empirical threshold q (10%, 5%, and 1%), the estimated model is:

$$r_t^{BTC} = \alpha + \beta r_t^{SX5E} + \theta D_t^q + \varphi(r_t^{SX5E} \times D_t^q) + \delta' X_t + \varepsilon_t$$

The choice of a tail-interaction model is determined by the estimand implied by H2. The hypothesis does not merely ask whether Bitcoin's average correlation with equities changes over time or whether equity exposure differs across quantiles of Bitcoin returns. It asks whether Bitcoin protects an investor on days when the EURO STOXX 50 records exceptionally large losses. Accordingly, the stress state is defined from the empirical distribution of the benchmark equity return:

$$D_(q,t) = \mathbb{1}[r_(S,t) \leq Q_q(r_S)], q \in \{0.10,\ 0.05,\ 0.01\}.$$

Here, Qq(rS) denotes the empirical q-th quantile of EURO STOXX 50 returns. The thresholds are pre-specified in the empirical design, but they are sample quantiles and are not described as exogenous. The existing interaction equation therefore estimates the change in Bitcoin's equity exposure during benchmark-defined adverse market states.

Sparse-tail inference. For each equity-return threshold, the analysis reports the empirical cutoff, number of tail observations and their distribution across calendar years. The 1% tail contains only 25 observations; consequently, conventional HAC confidence intervals are supplemented by moving-block bootstrap inference and crisis-year deletion exercises (Tables A5 and A6).

The equity beta outside the identified tail state is β, whereas the conditional equity beta during extreme market losses is:

$$\beta_q\hat{}(tail) = \beta + \varphi_q.$$

The safe-haven assessment is based on the conditional tail beta and its 95% confidence interval. Two distinct inferential questions are reported separately. First, evidence for a strong safe haven is evaluated with the one-sided test H0: βq(tail) ≥ 0 against H1: βq(tail) < 0. Only a sufficiently negative estimate can reject this null in favor of a strong safe-haven effect. Second, the tables report the conventional two-sided test H0: βq(tail) = 0 against H1: βq(tail) ≠ 0, together with a 95% confidence interval. A significantly positive conditional beta is therefore evidence against both strong and weak safe-haven classifications, whereas a confidence interval containing zero is compatible with, but does not establish, a weak safe-haven classification.

$$H_0:\ \beta_q\hat{}(tail) \geq 0\ against\ H_1:\ \beta_q\hat{}(tail) < 0.$$

Accordingly, a significantly negative conditional beta under the one-sided test supports a strong safe-haven classification. A conditional beta whose 95% confidence interval contains zero is statistically compatible with a weak safe haven, but this should not be interpreted as affirmative evidence of protection. A significantly positive conditional beta under the two-sided zero test rejects both strong and weak safe-haven classifications. The relevant *p*-values are explicitly labeled by their null hypotheses in the empirical results below.

CISS is a composite of several market segments including equities. Its upper-tail classification is an alternative broad-stress definition, not a stress measure independent of equity-market information. The CISS interaction regressions include the equity return, a CISS state dummy, their interaction, gold futures returns (USD), USD/EUR returns, the ECB

deposit-rate change, and lagged Bitcoin returns; ΔCISS is omitted from the implemented CISS specification.

DCC-GARCH and quantile regression remain relevant competing approaches, but they answer related rather than identical questions. DCC-GARCH captures smooth time variation in second-moment dependence but does not automatically isolate observations corresponding to the most severe equity losses. Quantile regression captures distributional heterogeneity, but its interpretation depends on the outcome whose conditional quantiles are modeled and does not necessarily coincide with conditioning on the lower tail of the benchmark equity market. The tail-interaction model is therefore retained because it maps directly into H2. Sensitivity to a single threshold is reduced by reporting the 10%, 5%, and 1% equity-return tails and CISS stress thresholds defined by a different rule from the EURO STOXX 50 lower-tail threshold, but not independent of equity-market information.

Model-specific control vectors are as follows. (i) Equity-tail baseline: USD gold-futures return, USD/EUR return, first difference of CISS, first difference of the ECB deposit facility rate and lagged daily BTC/EUR return, in addition to the EURO STOXX 50 return, the lower-tail indicator and their interaction. (ii) CISS-state model: the same controls except ΔCISS is omitted, because the state dummy is defined from the CISS level; the model includes the equity return, CISS high-stress indicator and equity-by-CISS interaction. (iii) Shifted-equity timing models: the equity return and corresponding tail indicator are shifted by one matched trading observation in either direction, with all non-equity controls at their original dates. (iv) Cumulative-return models: Bitcoin, equity, USD gold-futures and FX returns are accumulated over each 2–5-observation matched horizon; ΔCISS and ΔDFR span the corresponding window, and lagged BTC is measured before the window. The equity-tail models include ΔCISS, whereas the CISS-state models do not. Newey–West HAC(5) inference is used for daily baseline and timing models; the cumulative specifications use horizon-adjusted HAC lags. These timing exercises do not synchronize intraday market closes.

For cumulative-return sensitivity, forward log returns are summed over two to five consecutive complete-case matched trading observations, with horizon-specific equity lower-tail cutoffs. Gold and exchange-rate returns are summed over the same window; CISS and policy-rate changes are measured from the observation preceding the window to its end. The lagged Bitcoin return precedes the window. The overlapping windows and sparse extreme tails limit inference; these tests assess horizon sensitivity rather than providing an intraday-synchronized estimate.

Implemented CISS equation: $rBTC_t = \alpha + \beta\ rEQ_t + \delta\ D(CISS_t > Qc) + \gamma\ [rEQ_t \times D(CISS_t > Qc)] + \theta1\ rGoldUSD_t + \theta2\ rFX_t + \theta3\ \Delta DFR_t + \theta4\ rBTC_{(t-1)} + \varepsilon_t$, with $c = 0.90$ or 0.95. The conditional beta is $\beta + \gamma$; HAC(5) inference uses the 2447-row complete-case daily sample. No ΔCISS term enters this CISS-specific model.

4. Empirical Results and Diagnostics

This section reports the monthly inflation-hedge estimates and the daily conditional safe-haven estimates. The evidence does not support Bitcoin as an inflation hedge or a safe-haven asset for euro-area investors. Diagnostic tests are reported separately from hypothesis decisions.

4.1. Inflation-Hedge Results

Economic significance is evaluated using the estimation-sample-matched standardized association $|\beta| \times SD(INFShock)/SD(BTC\ return)$. The estimated standardized magnitudes for the controlled BIC models are 0.0004 (t), 0.1422 (t + 1), and 0.0510 (t + 2). These de-

scribe conditional sample associations, not purchasing-power protection or implementable trading gains.

Table 5 presents the controlled H1 results. The coefficient on the BIC-selected HICP innovation is 0.014 (HAC SE 4.251; $p = 0.997$; $N = 72$) contemporaneously, −4.876 (SE 3.100; $p = 0.116$; N = 71) for t + 1, and −1.749 (SE 4.498; $p = 0.697$; N = 70) for t + 2. None of these coefficients provides statistically significant evidence of a positive short-run inflation association. The underlying innovation follows the recursive procedure described in Section 3.3.

Table 5. Bitcoin Returns and Model-Based HICP Innovations under Alternative Timing Conventions.

Variable	BIC Innovation t	BIC Innovation t + 1	BIC Innovation t + 2	Fixed AR(1) t + 1
Model-based HICP innovation	0.014 (4.251)	−4.876 (3.100)	−1.749 (4.498)	−3.392 (3.113)
Lagged Bitcoin return	0.133 (0.107)	−0.060 (0.107)	0.026 (0.130)	−0.069 (0.104)
EURO STOXX 50 return	1.026 (0.421)	−0.490 (0.363)	0.076 (0.487)	−0.485 (0.364)
Gold return	0.358 (0.679)	−1.123 (0.658)	−0.092 (0.596)	−1.080 (0.641)
ΔCISS	−23.956 (35.297)	−53.297 (37.794)	−59.644 (53.445)	−48.069 (41.205)
ΔDeposit facility rate	−4.493 (11.879)	−7.673 (8.641)	−4.720 (9.397)	−6.965 (9.160)
Constant	3.150 (2.243)	6.724 (3.077)	4.781 (2.833)	6.393 (3.018)
Observations	72	71	70	71
Adjusted R^2	0.059	0.005	−0.062	−0.007

Notes: Coefficients and HAC(4) standard errors (in parentheses). Columns 1–3 use the BIC-selected recursive innovation; column 4 uses the fixed-AR(1) innovation at t + 1. Controls are lagged BTC/EUR returns, EURO STOXX 50 monthly returns, USD GC=F gold-futures returns, ΔCISS and ΔDFR. No Bitcoin trading-volume regressor is estimated. The HICP innovation is an ex post forecast error, not an announcement surprise.

Sparse-tail inference. The full-sample H2 lower-tail states contain 245 (10%), 123 (5%) and only 25 (1%) observations (N = 2447). Exploratory moving-block bootstrap percentile intervals, based on 400 resamples of 10 consecutive matched observations with fixed full-sample cutoffs and seed 20261021, include zero at all three thresholds (Table A4). These intervals should not be treated as precise tail-confidence limits. The full-sample positive HAC estimates are sensitive to excluding 2020 and to equity-return dating; they do not establish a crisis-invariant safe-haven failure.

4.2. Safe-Haven Results

Conditional equity beta estimates are positive in the contemporaneous full sample (10%: 1.650; 5%: 2.180; 1%: 3.531), but this result is highly sample-sensitive. Excluding calendar year 2020 (N = 2201, same full-sample cutoffs) yields 0.214 ($p = 0.696$), 0.265 ($p = 0.736$), and −0.584 ($p = 0.658$), respectively. Neither insignificance nor a positive point estimate establishes reliable defensive protection. Table 6 and the supplementary bootstrap/event-deletion results qualify the interpretation.

Table 6. Conditional Bitcoin Equity Beta during Market Downturns.

Downturn Threshold	Crisis Beta	HAC SE	Two-Sided *p*-Value H0: βtail = 0	95% CI	Tail N	Cutoff (%)
10% lower tail	1.650	0.794	0.038	[0.094; 3.206]	245	−1.3017
5% lower tail	2.180	0.924	0.018	[0.370; 3.990]	123	−1.8801
1% lower tail	3.531	1.152	0.002	[1.274; 5.789]	25	−3.6878

Notes: The crisis beta equals $\beta + \varphi$ from the interaction model. The reported two-sided *p*-value tests H0: $\beta q(\text{tail}) = 0$ against H1: $\beta q(\text{tail}) \neq 0$. The one-sided strong-safe-haven *p*-value corresponds to H0: $\beta q(\text{tail}) \geq 0$ against H1: $\beta q(\text{tail}) < 0$. The 95% confidence intervals are based on Newey–West HAC standard errors. A significantly negative conditional beta supports a strong safe haven; an interval containing zero is compatible with, but does not establish, a weak safe haven; and a significantly positive beta is evidence against both classifications.

The 1% estimate is particularly uncertain because the full-sample tail contains only 25 observations, falling to 11 after excluding 2020. The completed moving-block bootstrap (Table A4) yields a percentile interval including zero; the completed 2020-deletion exercise (Table A5) eliminates conventional significance. A one-matched-observation shift in the equity-return series reverses the 1% point-estimate sign (Table A1), but this sensitivity is not an intraday-clock-synchronized estimate and cannot establish a strong safe haven. Taken together, the completed checks prevent a stable generalization from the significant full-sample HAC result.

Table 6 reports the conditional equity betas for the three lower-tail thresholds and their associated inference.

Figure 1 displays the conditional equity-beta estimates and their HAC-based confidence intervals for the three equity-market downturn thresholds.

Tail-state frequency and empirical cutoffs (computed from the 2447 complete-case daily observations underlying the model variables): 10%: 245 days, cutoff −1.3017%; 5%: 123 days, cutoff −1.8801%; 1%: 25 days, cutoff −3.6878%. Thresholds use the empirical quantiles of EURO STOXX 50 daily log returns in percent, with the tail defined by return ≤ cutoff. These counts quantify the scarcity of tail observations, particularly at the 1% cutoff; they are descriptive and cannot establish the stability of the estimated conditional betas. Bootstrap and crisis-year deletion results are assessed separately

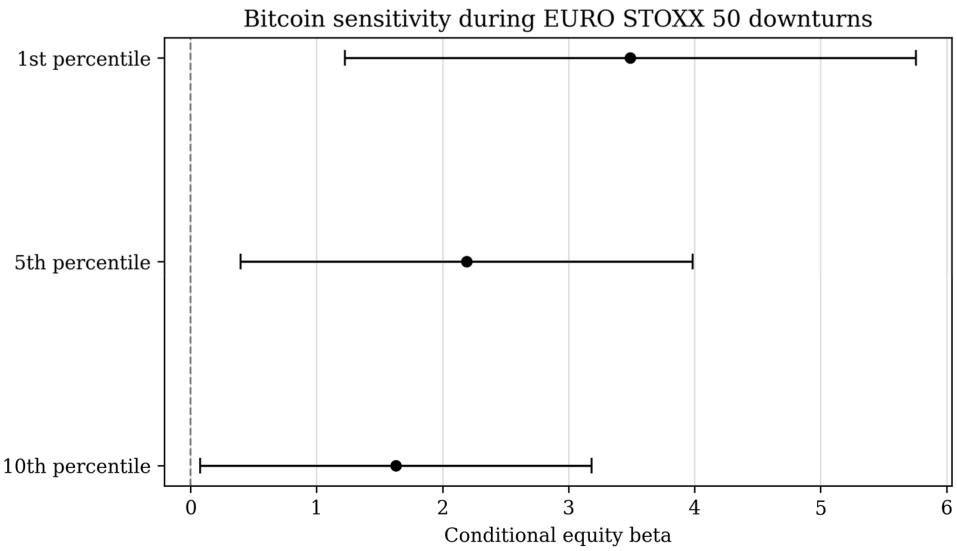


Figure 1. Conditional Bitcoin equity betas. Notes: Points report conditional equity betas; horizontal bars are 95% confidence intervals based on HAC standard errors. The dashed vertical line marks zero.

4.3. Diagnostic Tests

The recursively generated HICP innovation contains 72 observations from January 2019 to December 2024. Its mean is −0.0311 percentage points, and its sample standard deviation is 0.5476 percentage points, calculated from the 72 recursively generated observations used in H1. The first-order autocorrelation is 0.223. Ljung–Box tests yield $p = 0.498$ at six lags and $p = 0.083$ at twelve lags. These diagnostics indicate that the recursive AR procedure removes a substantial part of the persistence in year-on-year HICP inflation, although the innovation should not be interpreted as a structural shock or a market-consensus surprise.

Figure 2 displays the 72 recursively estimated monthly HICP forecast errors from January 2019 to December 2024. The series has a mean of −0.0311 and a sample standard deviation of 0.5476 percentage points. The figure illustrates the model-based inflation measure used in H1; it does not display or compare conditional equity betas. Cross-model coefficient comparisons appear instead in Figure 3 and should be interpreted with their distinct frequencies, scales and estimands in mind.

Durbin–Watson statistics close to two offer only limited descriptive evidence concerning first-order residual autocorrelation; they do not establish the absence of serial dependence, particularly with heavy-tailed daily returns. Inference relies primarily on HAC standard errors and the reported sensitivity checks. The monthly models do not reject residual normality or homoskedasticity. The daily models, however, strongly reject both assumptions, which is unsurprising for cryptocurrency returns. HAC inference is therefore essential and is complemented by alternative stress thresholds and sample restrictions in Section 5.

Figure 3 compares the reported inflation-hedge and safe-haven robustness coefficients, which refer to different estimands and scales.

Table 7 presents residual diagnostics for the daily safe-haven regressions; these diagnostics do not replace HAC-based inference.

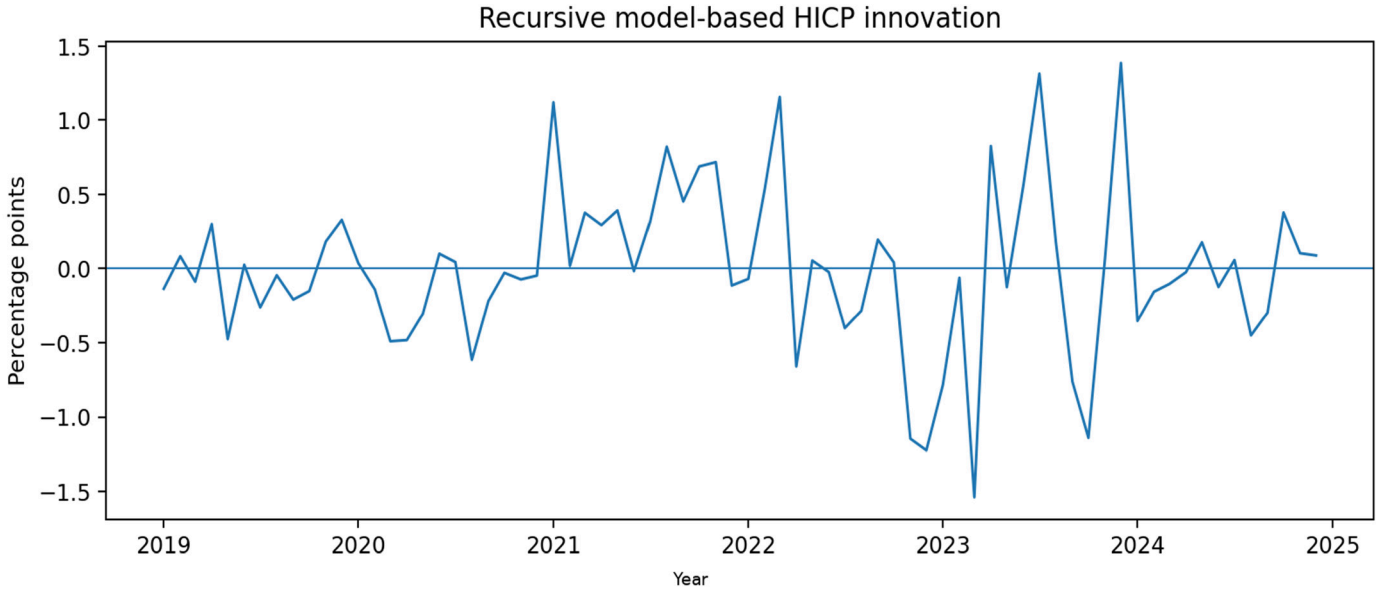


Figure 2. Recursive model-based HICP innovation, January 2019–December 2024.

Table 7. Residual diagnostics for the daily safe-haven regressions.

Model	N	Adjusted R^2	DW	JB *p*-Value	BP *p*-Value
H2 linear	2447	0.029	1.980	<0.001	<0.001
H2 5th percentile	2447	0.045	1.971	<0.001	<0.001
H2 1st percentile	2447	0.051	1.966	<0.001	<0.001

Table 7. *Cont.*

Model	N	Adjusted R^2	DW	JB *p*-Value	BP *p*-Value
H2 10th percentile	2447	0.039	1.971	<0.001	<0.001

Notes: JB denotes the Jarque-Bera normality test; BP denotes the Breusch-Pagan heteroskedasticity test. HAC standard errors are used throughout. Low R^2 values in return regressions are not themselves evidence of misspecification, but they confirm that the models explain only a limited share of Bitcoin return variation.

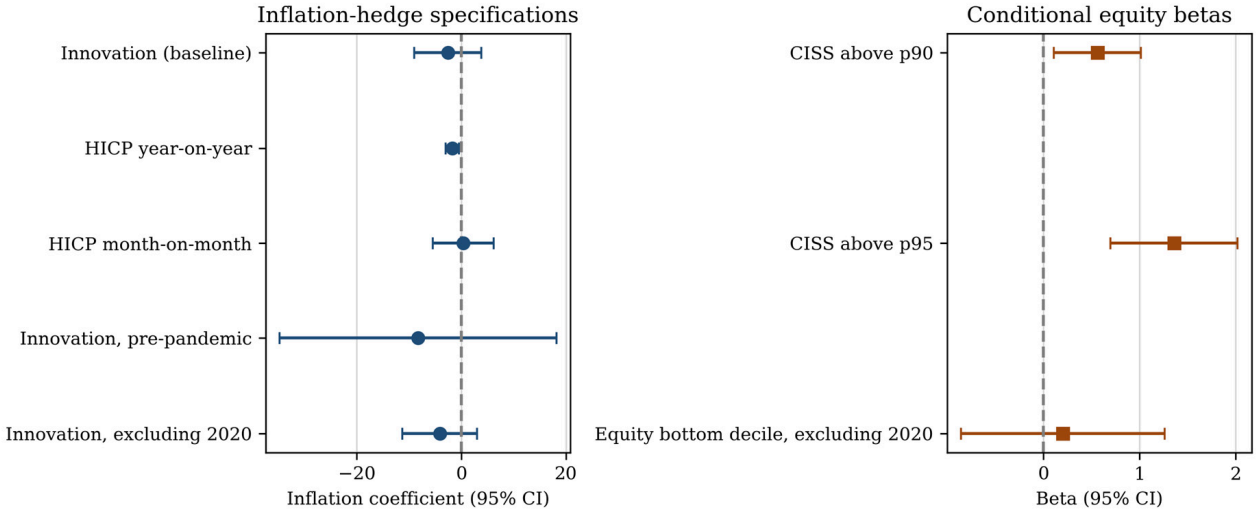


Figure 3. Robustness Coefficients for the Inflation-Hedge and Safe-Haven Models. Notes: Points are coefficient estimates; horizontal bars are 95% confidence intervals based on HAC standard errors. Coefficients in the left panel are not directly comparable in magnitude because the inflation measures use different scales.

4.4. Baseline Assessment of the Hypotheses

This section presents the baseline assessment of the two hypotheses examined in this study (Table 8). First, we evaluate whether Bitcoin provides short-run protection against inflation in the euro area (H1). Second, we examine whether Bitcoin exhibits safe-haven properties during periods of financial stress (H2). The assessment is based on the estimated coefficients, their statistical significance, and the corresponding empirical evidence.

Table 8. Summary of Baseline Hypothesis Tests.

Hypothesis	Empirical Proposition	Decision	Main Evidence
H1	Bitcoin returns respond positively to model-based HICP innovations.	Not supported	The model-based innovation coefficient and its subperiod interactions are insignificant.
H2	Bitcoin is a safe haven during severe EURO STOXX 50 losses.	Not robustly supported; strong safe-haven status not established	Positive HAC full-sample contemporaneous betas, but moving-block percentile intervals include zero; exclusion of 2020 removes significance and timing shifts reverse the 1% estimate. No stable classification across checks.

The controlled H1 estimates do not establish a statistically reliable positive association between Bitcoin returns and the recursive HICP innovation. For H2, the full-sample contemporaneous HAC estimates are positive during large European equity losses, but the moving-block bootstrap percentile intervals include zero; deleting 2020 removes significance, and the extreme-tail estimate changes sign under a one-observation equity-date shift. Thus, the data do not establish robust, timing-invariant safe-haven protection, nor do they justify a general claim that Bitcoin invariably becomes more equity-exposed in crises.

5. Robustness and Discussion

This section discusses the robustness exercises actually reported. H1 does not yield statistically reliable evidence of short-run inflation hedging under the BIC-selected and fixed-AR(1) recursive innovations. Alternative realized-inflation regressions and macroeconomic subperiod interactions are not tabulated and are not invoked as empirical support. H2 results are sensitive to bootstrap uncertainty, the exclusion of 2020, return dating and the definition of financial stress.

5.1. Robustness Design

For H1, the reported robustness exercise replaces the BIC-selected recursive innovation with an innovation from a fixed AR(1) forecasting model. No results for alternative realized-inflation measures or subperiod interaction terms are reported; such specifications are excluded from the evidential claims.

In addition, the recursive innovation is reconstructed with a fixed AR(1) specification. This directly tests whether the H1 conclusion depends on repeated BIC lag-order selection. The fixed-AR(1) innovation is strongly correlated with the BIC-based innovation (correlation = 0.852).

For H2, the aggregate CISS 90th and 95th percentile states are equity-inclusive. A separate sensitivity exercise uses the sum of ECB-published money-market, bond-market and foreign-exchange-market CISS contributions to classify high stress at the same percentiles. This alternative excludes the explicitly labeled equity-market contribution but is not a mathematically reconstructed equity-free CISS. Because contributions are calculated within the published CISS framework, they can still depend on cross-market covariances. The lower-decile equity model is also re-estimated after excluding 2020.

5.2. Robustness of the Inflation-Hedge Result

Under a fixed AR(1) inflation-forecasting specification with otherwise identical controls, the t + 1 coefficient is −3.392 (HAC SE 3.113; $p = 0.276$; N = 71) and the t + 2 coefficient is −2.261 (HAC SE 4.952; $p = 0.648$; N = 70). Neither alternative supports a statistically reliable positive inflation association (Table 9).

Table 9. Robustness of the Inflation-Hedge Estimates.

Specification	Return Timing	N	β: HICP Innovation	HAC SE	*p*-Value	95% CI
Fixed AR(1), full controls	t + 1	71	−3.392	3.113	0.276	[−9.495; 2.710]
Fixed AR(1), full controls	t + 2	70	−2.261	4.952	0.648	[−11.968; 7.445]

5.3. Robustness of the Safe-Haven Result

In the aggregate-CISS interaction models, ΔCISS is omitted from the control vector. The conditional equity beta is 0.549 (HAC SE 0.238; $p = 0.021$; 244 high-stress observations; CISS cutoff 0.291101) above the 90th percentile and 1.298 (SE 0.476; $p = 0.006$; 121 observations; cutoff 0.354851) above the 95th percentile. CISS incorporates equity-market information (Holló et al., 2012). As an additional empirical sensitivity exercise, we use the sum of the ECB's money-, bond- and FX-market CISS contributions, which does not explicitly include the equity-market contribution. The alternative indicator is a contribution-based proxy, not a recomputed equity-excluded CISS. Its HAC estimates are reported in the new Table 10 and can be reproduced from the supplied ECB exports and daily analysis script.

Table 10. Robustness using ECB non-equity-segment contribution proxy (2015–2024).

Proxy Stress	Cutoff	Stress N	Total N	Conditional β	HAC SE	Two-Sided *p*	95% CI
P90	0.31184	245	2447	1.127	0.411	0.006	[0.321, 1.933]
P95	0.36840	122	2447	1.314	0.516	0.011	[0.303, 2.325]

Notes: The stress proxy is the daily sum of ECB-published CISS contributions from the money, bond and FX market segments (SS_MMN.CON, SS_BMN.CON, SS_FXN.CON), observed on the matched sample dates. It is not the ECB's official CISS recalculated without equities. The published contributions can reflect covariance with other segments, including equities. High-stress states are values strictly above the 90th or 95th sample percentile. The dependent variable is BTC/EUR daily log return (%); regressors are EURO STOXX 50 return, stress dummy, their interaction, USD gold-futures return, USD/EUR return, change in ECB deposit rate, and lagged BTC/EUR return. The change in the equity-inclusive aggregate CISS is deliberately omitted from this alternative model. HAC/Newey–West standard errors use five lags. Conditional β is the sum of the equity main effect and its interaction coefficient; two-sided *p-tests* for $\beta = 0$. These are descriptive conditional associations, not causal effects.

Table 10 reports the additional sensitivity test based on the ECB money-, bond-, and foreign-exchange-market contribution proxy, rather than an officially reconstructed equity-excluded CISS.

The left panel reports monthly HICP-innovation coefficients, whereas the right panel reports daily conditional equity betas. The two panels use different scales and estimands, and their coefficient magnitudes are not directly comparable.

Robustness is assessed using fixed full-sample tail cutoffs, exclusion of 2020, and moving-block resampling. The extreme 1% state is too sparse to sustain stable claims across crisis episodes; its sign and precision change materially after excluding 2020.

With 2020 excluded (N = 2201), conditional betas at the full-sample 10%, 5% and 1% equity-tail thresholds are 0.214 (two-sided $p = 0.696$; 206 tail observations), 0.265 ($p = 0.736$; 97 observations), and −0.584 ($p = 0.658$; 11 observations). The verified full-sample and 2020-deletion estimates are reported in Table A5.

Return-timing sensitivity is economically material. When the equity return is lagged by one matched observation, the 1% conditional beta is −1.001 (HAC SE 0.415; two-sided $p = 0.016$; N = 2446), compared with the positive same-date estimate in Table 6. Table A1 reports the three tail thresholds across alternative timing conventions. The shifts are defined on the matched trading-date panel, not by intraday timestamps.

Table A2 reports overlapping cumulative-return models for horizons of two to five matched observations. Newey–West HAC inference uses $\max(5, h + 2)$ lags for horizon h. Overlapping windows induce dependence, and the reported standard errors should be interpreted subject to the maintained HAC assumptions.

The complete-case daily sample contains 2447 observations. The 90th percentile of CISS is 0.291101 (244 observations strictly above the cutoff), and the 95th percentile is 0.354851 (121 observations). All thresholds are calculated from the same complete-case estimation sample.

Figure-panel comparability. Monthly HICP-innovation coefficients and daily conditional equity betas measure different empirical relationships and are displayed on distinct scales; their graphical magnitudes are not directly comparable.

Table 11 reports the equity-tail and CISS interaction specifications. The equity-tail baseline controls for ΔCISS, whereas the CISS interaction models exclude ΔCISS to avoid including the change in the stress index alongside its level-based stress classification. Reported conditional betas are sums of the main equity-return coefficient and the interaction coefficient.

Table 11. Conditional Equity Betas under Alternative Stress Definitions.

Stress Definition	Cutoff	Stress n	Total N	Conditional Beta	HAC SE	Two-Sided p	95% CI
CISS > p90	0.291101	244	2447	0.549	0.238	0.021	[0.081; 1.016]
CISS > p95	0.354851	121	2447	1.298	0.476	0.006	[0.365; 2.232]
Equity lower 10%	−1.3017%	245	2447	1.650	0.794	0.038	[0.094; 3.206]

Notes: The two-sided p-value tests H0: $\beta tail = 0$ against H1: $\beta tail \neq 0$. The one-sided p-value for a strong safe haven corresponds to H0: $\beta tail \geq 0$ against H1: $\beta tail < 0$. The 95% confidence intervals use HAC standard errors. A significantly positive beta is evidence against both strong and weak safe-haven classifications. An interval containing zero is compatible with, but does not establish, a weak safe haven. CISS thresholds are empirical percentiles over the daily sample.

5.4. Discussion

5.4.1. Bitcoin as an Inflation Hedge

The absence of a positive association between the model-based HICP innovation and Bitcoin returns is robust to alternative information-timing conventions. In particular, moving the Bitcoin return to t + 1 and t + 2 does not produce a positive statistically supported coefficient. The evidence therefore does not support a short-run euro-area inflation-hedge interpretation, while it does not identify a causal response to HICP releases.

A fixed supply rule is not sufficient to produce a short-horizon inflation hedge. Euro-area investors price Bitcoin in a market influenced by global risk appetite, liquidity, leverage, exchange-rate conditions and technology-specific news. The sample spans major changes in market participation and monetary conditions. The negative coefficient on annual inflation is consistent with the possibility that high inflation coincides with restrictive financial conditions, but the present design cannot separate that mechanism from concurrent shocks.

5.4.2. Bitcoin as a Safe Haven

The full-sample contemporaneous H2 HAC estimates indicate positive conditional Bitcoin–equity exposure in severe downturns, which does not support a strong safe-haven interpretation in that specification. However, the moving-block bootstrap 95% percentile intervals include zero at all three thresholds; excluding 2020 eliminates significance; and shifting equity returns by one matched trading observation reverses the sign at the 1% threshold. These sensitivity results rule out a robust general claim of stronger equity exposure in all crises. Nor do they establish a strong safe haven: the one-observation shift is a timing sensitivity, not a precisely synchronized intraday estimate. The 1% tail is especially sparse (25 full-sample observations, 11 without 2020). The 400-draw bootstrap intervals are exploratory rather than definitive tail inference. Cumulative two- to five-observation results characterize horizon sensitivity but cannot resolve close-time mismatch. The findings therefore fail to establish stable, timing-invariant safe-haven protection for euro-area equity investors.

Scope of the gold comparison. Although gold returns are included as a covariate, a coefficient on gold in a Bitcoin-return equation is not a test of gold's own inflation-hedging or safe-haven performance. We do not estimate matched gold-dependent-variable H1 and H2 specifications and therefore cannot draw a within-design comparison between Bitcoin and gold. The interpretation of the empirical results is restricted to Bitcoin. A direct benchmark comparison would require re-estimating the monthly inflation and daily equity-tail models with gold returns as the dependent variable, using appropriately adapted controls and identical euro-area state definitions. The additional ECB contribution-based stress proxy offers a distinct sensitivity check but does not eliminate all possible equity-market dependence; its results should not be represented as an exact equity-excluded CISS estimate (Table 10).

5.5. Robustness-Adjusted Assessment and Remaining Limitations

The robustness evidence does not support a reliable positive short-run association between Bitcoin returns and model-based inflation innovations. The safe-haven evidence is more sample- and timing-dependent: positive conditional equity exposures in the full sample weaken when 2020 is excluded, while a one-observation equity-return shift reverses the sign of the 1% estimate. Consequently, the analysis does not establish stable safe-haven protection, but it also does not rule out defensive behavior in particular episodes.

Table 12 summarizes the final assessment of hypotheses H1 and H2 following the robustness analyses.

Table 12. Hypothesis Decisions after Robustness Analysis.

Hypothesis	Baseline Decision	Robustness-Adjusted Decision	Reason
H1: Inflation protection	Not supported	Not supported	No stable positive coefficient across the model-based HICP innovation and alternative inflation measures.
H2: Safe haven	Not established	Inconclusive; sample- and timing-dependent	Bootstrap intervals include zero, excluding 2020 weakens exposure, and one-day shifts change the 1% result.

Accordingly, the evidence supports two conclusions: Bitcoin is not a reliable short-run inflation hedge under the measures examined, and it does not provide systematic protection against European equity-market or systemic-stress episodes.

The baseline approach does not replace the full information provided by multivariate volatility and distributional models. A DCC-GARCH specification could reveal gradual changes in the conditional correlation between Bitcoin and European equities, while quantile regression could identify broader asymmetries across the conditional distribution of Bitcoin returns. The tail-interaction estimates target a narrower parameter: Bitcoin's contemporaneous equity exposure on days belonging to the lower tail of the EURO STOXX 50 distribution. The results should therefore be interpreted as a test of benchmark-defined crisis protection rather than as a complete characterization of dynamic or distributional dependence.

The 1% equity-loss state contains only 25 observations, so a small number of crisis episodes can materially affect the conditional beta. The crisis-year deletion results and moving-block bootstrap confidence intervals in Tables A5 and A6 therefore receive particular weight in interpreting extreme-tail evidence.

6. Portfolio and Policy Implications

The empirical findings have implications for portfolio construction and financial-stability monitoring. These implications remain proportionate to the evidence. The results do not establish that Bitcoin has no investment value or that its correlation with conventional assets is permanently fixed. They show that, over 2015–2024, Bitcoin did not provide dependable protection against euro-area inflation innovations or severe European market stress.

6.1. Portfolio Allocation and Risk Management

For euro-area investors, Bitcoin should not be assigned the mechanical role commonly attributed to an inflation hedge. The monthly models produce no stable positive response to model-based HICP innovations, while annual HICP inflation is associated with a significant negative coefficient in one alternative specification. A portfolio mandate that requires

preservation of real purchasing power therefore cannot rely on Bitcoin exposure alone. Any allocation to Bitcoin should be justified by expected return, long-horizon diversification, or exposure to digital-asset adoption rather than by an empirically unsupported promise of short-run inflation protection.

The safe-haven results are more restrictive. Bitcoin's conditional equity beta is positive during severe EURO STOXX 50 losses and during episodes in which the CISS exceeds its upper stress percentiles. This pattern means that diversification benefits observed in normal periods can weaken precisely when portfolio insurance is most valuable. Risk managers should therefore evaluate Bitcoin using stressed correlations, conditional betas, drawdowns, liquidity conditions, and portfolio-level expected shortfall rather than full-sample average correlations. Position limits and stress scenarios should allow for joint declines in Bitcoin and equities, particularly when tightening financial conditions trigger deleveraging across speculative assets.

The exclusion of 2020 makes the lower-decile equity beta statistically insignificant, but it does not establish a safe-haven effect. Failure to reject a zero coefficient is not evidence of negative crisis exposure. Accordingly, the defensible interpretation is that Bitcoin's protective capacity is unstable, not that it is universally procyclical. This distinction matters for asset-allocation decisions because it allows for possible diversification benefits without misclassifying Bitcoin as crisis insurance.

6.2. Financial-Stability Implications

The positive relationship between Bitcoin and European equities during high-stress states weakens the argument that crypto-assets automatically absorb shocks originating in conventional finance. As participation by institutional and leveraged investors expands, common funding constraints and synchronized risk reduction may transmit rather than diversify market stress. The relevant concern for authorities is not the volatility of Bitcoin in isolation, but the interaction between crypto-asset exposures, leverage, collateral practices, liquidity mismatches, and links with regulated financial intermediaries.

The results do not quantify systemic risk and cannot establish regulatory necessity. They indicate that joint crypto-asset and equity declines are a plausible scenario for further investigation, not a policy prescription

7. Conclusions

This study examines whether Bitcoin functions as an inflation hedge and a safe haven for euro-area investors. It combines monthly data for the inflation analysis with daily data for the assessment of safe-haven behavior between January 2015 and December 2024. The empirical framework distinguishes purchasing-power protection from crisis protection and evaluates each property using euro-denominated Bitcoin returns.

First, Bitcoin does not display a stable positive short-run association with the recursive model-based HICP innovation. This conclusion holds when the contemporaneous specification is interpreted only as ex post covariance and when Bitcoin returns are shifted to $t + 1$ and $t + 2$ to respect the timing of HICP information.

Second, the safe-haven assessment is sensitive to the crisis composition and the timing convention. Although full-sample contemporaneous equity-tail estimates are positive, their moving-block bootstrap confidence intervals include zero, and exclusion of 2020 yields statistically insignificant estimates. At the 1% tail, lagging equity returns by one common trading observation reverses the sign. The evidence therefore does not establish a stable safe haven, but also does not support a universal rejection of all safe-haven potential.

The contribution is to distinguish inflation-hedge associations from state-contingent equity exposure while explicitly documenting sensitivity to inflation forecasting, crisis

episodes, gold-currency controls, and closing-time alignment. Conclusions are conditional on the archived sample and implemented empirical specifications.

The results are sample- and timing-dependent. H1 fails to show a statistically reliable positive short-run association with the model-based HICP forecast error; this is not evidence that Bitcoin can never preserve purchasing power at longer horizons. The contemporaneous full-sample H2 estimates are positive, but bootstrap intervals cover zero, deletion of 2020 eliminates significance, and a one-day shift of the equity return produces a negative significant 1% estimate. Accordingly, no universal safe-haven classification or definitive rejection is warranted.

The timing analysis reveals a substantial reversal at the 1% threshold: with the equity return lagged by one matched observation, the conditional beta is −1.001 (HAC SE 0.415; two-sided $p = 0.016$; N = 2446). This result limits the robustness of any classification based solely on contemporaneously matched returns (Table A1).

Finally, the empirical design does not directly assess gold's own defensive properties. Gold enters the Bitcoin equations as a control rather than as the dependent variable in matched inflation-hedge and safe-haven tests. Accordingly, no comparative ranking of Bitcoin and gold is warranted by these estimates.

Third, same-day comovement in the safe-haven models does not identify a structural transmission mechanism, and the classification remains sensitive to the chosen market, tail threshold, and sample period. The evidence therefore concerns Bitcoin and euro-area investors over 2015–2024; it is not a universal statement about other crypto-assets or future market conditions.

Future research could examine real-time HICP announcements, portfolio purchasing-power outcomes and out-of-sample hedging performance. The findings do not establish a dependable inflation hedge or systematic crisis insurance in the examined sample, without implying a universal portfolio prohibition.

Author Contributions: Conceptualization, A.A. and R.F.; methodology, A.A. and R.F.; software, A.A.; validation, R.F.; formal analysis, A.A.; investigation, A.A.; resources, A.A. and R.F.; data curation, A.A.; writing—original draft preparation, A.A. and R.F.; writing—review and editing, A.A. and R.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The replication materials include Python (3.13) analysis scripts, derived regression tables, the ECB daily contribution exports used for the non-equity-segment proxy, and instructions for reproducing the estimates. The daily Bitcoin, equity, gold and FX vendor price histories in the working dataset may be subject to third-party redistribution restrictions and should not be placed in a public repository without verifying the applicable licenses. The source datasets and matching rules are described in the data dictionary. The ECB component-contribution robustness test is reproducible using the supplied ECB extracts together with the authors' original daily working dataset.

Acknowledgments: During the preparation and revision of this manuscript, the authors used OpenAI ChatGPT (GPT-6) to assist with academic language editing. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Additional Synchronization Robustness Tests

Table A1 reports the sensitivity of conditional Bitcoin equity betas to a one-observation shift in equity returns.

Table A1. Calendar-observation lead/lag sensitivity of conditional Bitcoin equity betas (one matched trading-observation shift; HAC(5)).

Equity Timing	Tail	N	Conditional Beta	HAC SE	*p*-Value	95% CI
equity t + 1	10%	2446	−0.276	0.166	0.096	[−0.602; 0.049]
equity t	10%	2447	1.650	0.794	0.038	[0.094; 3.206]
equity t − 1	10%	2447	−0.532	0.346	0.124	[−1.210; 0.146]
equity t + 1	5%	2446	−0.194	0.217	0.371	[−0.619; 0.231]
equity t	5%	2447	2.180	0.924	0.018	[0.370; 3.990]
equity t − 1	5%	2447	−0.530	0.409	0.195	[−1.332; 0.272]
equity t + 1	1%	2446	−0.402	0.262	0.126	[−0.916; 0.112]
equity t	1%	2447	3.531	1.152	0.002	[1.274; 5.789]
equity t − 1	1%	2447	−1.000	0.414	0.016	[−1.812; −0.188]

Notes: The dependent variable is the daily BTC/EUR log return. The equity-return regressor is evaluated at t − 1, t and t + 1 on the common trading-date sequence; other daily controls remain aligned with the Bitcoin return date. Tail cutoffs are calculated within each timing sample. The specification includes the equity-return coefficient, a tail-state indicator and their interaction, gold-futures returns, USD/EUR returns, ΔCISS, ΔDFR and lagged Bitcoin returns. The conditional beta is the sum of the main equity coefficient and the interaction coefficient. Two-sided *p*-values and 95% confidence intervals use Newey–West HAC(5) inference. The shift is by observation, not by a verified number of intraday hours.

Table A2 reports the cumulative-return sensitivity across two- to five-observation horizons.

Table A2. Overlapping cumulative-return sensitivity of conditional Bitcoin equity betas (h = 2–5 matched observations; HAC max(5, h + 2)).

Horizon	Tail	N	Conditional Beta	HAC SE	*p*-Value	95% CI
2 days	10%	2445	1.224	0.723	0.091	[−0.194; 2.641]
2 days	5%	2445	1.956	0.838	0.020	[0.313; 3.599]
2 days	1%	2445	2.787	1.175	0.018	[0.485; 5.089]
3 days	10%	2444	1.813	0.504	0.000	[0.826; 2.800]
3 days	5%	2444	2.056	0.597	0.001	[0.886; 3.226]
3 days	1%	2444	2.698	0.877	0.002	[0.979; 4.418]
4 days	10%	2443	1.829	0.443	0.000	[0.959; 2.698]
4 days	5%	2443	2.306	0.450	0.000	[1.424; 3.188]
4 days	1%	2443	2.937	0.549	0.000	[1.861; 4.012]
5 days	10%	2442	2.073	0.290	0.000	[1.505; 2.641]
5 days	5%	2442	2.229	0.288	0.000	[1.666; 2.793]
5 days	1%	2442	2.740	0.408	0.000	[1.940; 3.540]

Notes: BTC/EUR, EURO STOXX 50, gold-futures and USD/EUR returns are summed over h consecutive forward complete-case matched observations. The first window is omitted where the preceding CISS observation is unavailable. CISS and policy-rate changes span the full horizon; the lagged Bitcoin control is measured at the start. Equity thresholds are recalculated for each horizon. HAC maxlags = max(5, h + 2). These overlapping-window regressions are sensitive to window construction and do not identify intraday synchronization.

Table A3 reports the empirical equity-tail cutoffs and associated sample sizes.

Table A3. Equity lower-tail sample sizes and empirical cutoffs.

Equity Tail	Cutoff, % Daily Log Return	Tail Observations	Total Regression N
10%	−1.3017	245	2447
5%	−1.8801	123	2447
1%	−3.6878	25	2447

Table A4 reports the moving-block bootstrap confidence intervals.

Table A4. Moving-block bootstrap confidence intervals for conditional equity betas (400 draws; 10-observation blocks; seed 20261021; fixed full-sample thresholds).

Stress Definition	Quantile	State N	HAC Beta	Bootstrap 95% Percentile CI	Valid Draws
Equity	10%	245	1.650	[−0.126; 2.709]	400
Equity	5%	123	2.180	[−0.311; 3.338]	400
Equity	1%	25	3.531	[−1.330; 7.436]	400

Table A5 reports the crisis-year deletion sensitivity.

Table A5. Crisis-year deletion sensitivity (full-sample equity thresholds held fixed; HAC(5)).

Sample	Equity Tail	N	Tail N	Conditional Beta	Two-Sided p
full	10%	2447	245	1.650	0.038
ex_2020	10%	2201	206	0.214	0.696
full	5%	2447	123	2.180	0.018
ex_2020	5%	2201	97	0.265	0.736
full	1%	2447	25	3.531	0.002
ex_2020	1%	2201	11	−0.584	0.658

Table A6 reports the euro-converted gold-futures sensitivity of the monthly inflation estimates.

Table A6. Monthly H1 sensitivity to EUR-converted GC=F gold-futures proxy (HAC(4)).

BTC Return	Gold Control	N	HICP Beta	HAC SE	p-Value	95% CI
t + 0	USD futures	72	0.014	4.251	0.997	[−8.317, 8.345]
t + 0	EUR futures proxy	72	0.105	4.128	0.980	[−7.986, 8.197]
t + 1	USD futures	71	−4.876	3.100	0.116	[−10.952, 1.201]
t + 1	EUR futures proxy	71	−4.653	3.063	0.129	[−10.656, 1.350]
t + 2	USD futures	70	−1.749	4.498	0.697	[−10.565, 7.067]
t + 2	EUR futures proxy	70	−3.065	4.339	0.480	[−11.569, 5.439]

Appendix B. Data Construction and Replication Methods

The monthly H1 reconstruction starts from the Eurostat HICP index in the archived workbook, forms $100 \times [\ln(\text{HICP_t}) - \ln(\text{HICP_(t} - 12))]$, and recursively estimates AR(p), p = 1, ..., 12, with a 36-observation initial annual-inflation training window. BIC is compared over the same estimation span for each p. All H1 tables use this single innovation series. The stored inflation_yoy_pct field differs from the log-index definition and is not substituted for the reconstructed series.

Daily H2 uses the aligned 2447-observation complete-case sample, log returns in percent, the specified equity-tail/CISS interaction models, HAC(5), and full-sample empirical cutoffs. Bootstrap draws use overlapping blocks of 10 consecutive trading observations sampled with replacement (400 draws, fixed seed 20261021), with thresholds held at their original full-sample values. Tables A5 and A6 provide the equity-tail bootstrap and 2020-deletion outputs.

The financial series are identified by Yahoo Finance tickers BTC-EUR, ^STOXX50E and GC=F; the ECB supplies CISS, deposit-rate and USD/EUR series, and Eurostat supplies HICP. The daily analysis is based on common trading-date matching and does not impose intraday synchronization.

A monthly EUR-converted gold-futures sensitivity check uses the last available monthly USD/EUR quote from the archived daily FX panel. This conversion is a futures-price proxy and does not create a matched LBMA spot series. The H1 innovation coefficients remain statistically insignificant across t, t + 1, and t + 2 (Table A7).

The supplementary data and code provide the aligned monthly and daily datasets, recursive HICP innovation algorithm, conditional-beta regressions, block-bootstrap and crisis-deletion analyses, and machine-readable result tables.

Table A1 shifts equity returns by −1, 0 and +1 matched observations. Table A2 considers rolling cumulative returns over h = 2, 3, 4 and 5 observations, with complete-case estimation and horizon-specific equity-return thresholds. The supplementary scripts and CSV outputs report sample sizes, tail counts, HAC lag choices and confidence intervals, including the cumulative-horizon convention used in Table A3.

Appendix C. Descriptive Statistics, Correlation Matrices, and Supplementary Tables

Tables A7–A10 below present descriptive statistics and Pearson correlation matrices for the daily and monthly datasets. These tables describe the archived variables and do not replace the recursive HICP inflation innovation used in the H1 regressions. Table A11 reports the additional distribution diagnostics.

Table A7. Daily Descriptive Statistics.

Variable	N	Mean	SD	Min	P25	Median	P75	Max
BTC return (%)	2447	0.2225	4.3417	−45.7842	−1.4771	0.1721	1.9771	22.7834
EURO STOXX 50 return (%)	2447	0.0181	1.2272	−13.2405	−0.5264	0.0605	0.6225	8.8343
Gold futures return (USD, %)	2447	0.0313	0.9350	−5.1069	−0.4341	0.0390	0.5430	5.7775
USD/EUR return (%)	2447	−0.0042	0.5047	−3.6820	−0.2754	0.0000	0.2763	3.4946
ECB CISS	2447	0.1331	0.1020	0.0229	0.0584	0.0915	0.1746	0.5134
ECB deposit rate (%)	2447	0.4302	1.5929	−0.5000	−0.5000	−0.4000	−0.2000	4.0000

Notes: Descriptive statistics refer to the archived daily variables and their available observations. Gold returns refer to USD-denominated GC=F gold futures, not EUR spot gold.

Table A8. Monthly Descriptive Statistics.

Variable	N	Mean	SD	Min	P25	Median	P75	Max
BTC return (%)	119	5.1677	19.8970	−45.3821	−7.4020	4.9886	18.5584	49.8274
HICP inflation YoY (%)	108	2.6430	2.7202	−0.3134	1.0179	1.7190	2.9071	10.6399
EURO STOXX 50 return (%)	119	0.3139	4.8520	−17.7889	−2.9012	0.7717	3.4402	16.6045
Gold futures return (USD, %)	119	0.6059	3.8158	−8.2510	−1.9995	0.2442	3.1300	10.0071
ECB CISS (monthly mean)	120	0.1315	0.0975	0.0316	0.0605	0.0913	0.1657	0.4542
ECB deposit rate (%)	120	0.4404	1.6035	−0.5000	−0.5000	−0.4000	−0.1500	4.0000

Notes: Descriptive statistics refer to archived monthly variables. The HICP inflation measure shown here is not the recursively generated H1 forecast innovation. Gold returns refer to USD-denominated GC=F gold futures.

Table A9. Daily Correlation Matrix.

Variable	BTC Return (%)	EURO STOXX 50 Return (%)	Gold Futures Return (USD, %)	USD/EUR Return (%)	ECB CISS	ECB Deposit Rate (%)
BTC return (%)	1.0000	0.1613	0.0543	−0.0062	−0.0288	0.0084
EURO STOXX 50 return (%)	0.1613	1.0000	−0.0514	0.0544	0.0317	0.0129
Gold futures return (USD, %)	0.0543	−0.0514	1.0000	0.1853	0.0066	0.0255
USD/EUR return (%)	−0.0062	0.0544	0.1853	1.0000	0.0124	0.0042
ECB CISS	−0.0288	0.0317	0.0066	0.0124	1.0000	0.1105
ECB deposit rate (%)	0.0084	0.0129	0.0255	0.0042	0.1105	1.0000

Notes: Pearson correlations among the archived daily variables; pairwise availability may differ across entries. Gold is USD-denominated GC=F futures.

Table A10. Monthly Correlation Matrix.

Variable	BTC Return (%)	HICP Inflation YoY (%)	EURO STOXX 50 Return (%)	Gold Futures Return (USD, %)	ECB CISS (Monthly Mean)	ECB Deposit Rate (%)
BTC return (%)	1.0000	−0.2008	0.2832	0.0458	−0.1462	0.0244
HICP inflation YoY (%)	−0.2008	1.0000	0.0024	−0.0793	0.6496	0.3169
EURO STOXX 50 return (%)	0.2832	0.0024	1.0000	−0.0705	0.0084	0.0715
Gold futures return (USD, %)	0.0458	−0.0793	−0.0705	1.0000	0.0643	0.1339
ECB CISS (monthly mean)	−0.1462	0.6496	0.0084	0.0643	1.0000	0.1274
ECB deposit rate (%)	0.0244	0.3169	0.0715	0.1339	0.1274	1.0000

Notes: Pearson correlations among the archived monthly variables; the inflation series is observed HICP inflation, not the recursive H1 innovation. Gold is USD-denominated GC=F futures.

Table A11. Additional descriptive distribution diagnostics (archived aligned data).

Frequency	Variable	N	Skewness	Excess Kurtosis	Jarque–Bera
Daily	r_btc_pct	2513	−0.620	9.153	8894.37
Daily	r_sx5e_pct	2513	−0.855	11.103	13,156.64
Daily	r_gold_pct	2448	−0.128	3.626	1340.65
Daily	r_fx_pct	2513	−0.105	4.022	1689.51
Daily	ciss	2514	1.455	1.586	1148.06
Monthly	r_btc_pct	119	−0.120	0.101	0.29
Monthly	inflation_yoy_pct	108	1.430	1.109	40.31
Monthly	r_gold_pct	119	0.139	−0.308	0.97
Monthly	ciss_mean	120	1.388	1.222	43.86

Notes: Descriptive calculations use non-missing observations for each series separately and therefore do not necessarily have identical N to regression complete-case samples. Excess kurtosis equals kurtosis minus three. Jarque–Bera statistics test normality. The monthly table is descriptive and is not a residual-diagnostic table for H1.

References

Acikgoz, T. (2025). Gold and Bitcoin as hedgers and safe havens: Perspective from nonlinear dynamics. *Resources Policy*, *102*, 105489. [CrossRef]

Aguir, A. (2013). *L'indépendance réelle de la banque centrale et le biais inflationniste: Une validation empirique* [Working paper]. HAL. Available online: https://shs.hal.science/halshs-00800059v1 (accessed on 7 August 2026).

Aguir, A. (2025). Post-COVID monetary policy challenges in emerging economies: Revisiting the effectiveness of inflation targeting. *Pakistan Journal of Life and Social Sciences*, *23*(1), 9014–9021. [CrossRef]

Aguir, A., & Fragny, B. (2026). The crypto real estate nexus in the Eurozone: Understanding financial transmission channels. *Bankers, Markets & Investors*, *183*, 15–27. [CrossRef]

Aguir, A., Smida, M., & Ftiti, Z. (2017). Régime de ciblage d'inflation dans les économies émergentes ou en développement: Quels enseignements après la crise? *Management & Prospective*, *34*(4), 77–94. [CrossRef]

Bahloul, S., Mroua, M., & Naifar, N. (2023). Re-evaluating the hedge and safe-haven properties of Islamic indexes, gold and Bitcoin: Evidence from DCC–GARCH and quantile models. *Journal of Islamic Accounting and Business Research*, *14*(8), 1167–1181. [CrossRef]

Baur, D. G., & Lucey, B. M. (2010). Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. *Financial Review*, *45*(2), 217–229. [CrossRef]

Baur, D. G., & McDermott, T. K. (2010). Is gold a safe haven? International evidence. *Journal of Banking & Finance*, *34*(8), 1886–1898. [CrossRef]

Blau, B. M., Griffith, T. G., & Whitby, R. J. (2021). Inflation and Bitcoin: A descriptive time-series analysis. *Economics Letters*, *203*, 109848. [CrossRef]

Bouri, E., Molnár, P., Azzi, G., Roubaud, D., & Hagfors, L. I. (2017). On the hedge and safe haven properties of Bitcoin: Is it really more than a diversifier? *Finance Research Letters*, *20*, 192–198. [CrossRef]

Choi, S., & Shin, J. (2022). Bitcoin: An inflation hedge but not a safe haven. *Finance Research Letters*, *46*, 102379. [CrossRef] [PubMed]

Conlon, T., & McGee, R. (2020). Safe haven or risky hazard? Bitcoin during the COVID-19 bear market. *Finance Research Letters*, *35*, 101607. [CrossRef] [PubMed]

Dardouri, N., Aguir, A., & Smida, M. (2023). The effect of COVID-19 transmission on cryptocurrencies. *Risks*, *11*(8), 139. [CrossRef]

Dyhrberg, A. H. (2016). Hedging capabilities of bitcoin. Is it the virtual gold? *Finance Research Letters*, *16*, 139–144. [CrossRef]

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*(3), 339–350. [CrossRef]

Holló, D., Kremer, M., & Lo Duca, M. (2012). *CISS-A composite indicator of systemic stress in the financial system* (ECB Working Paper No. 1426). Available online: https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1426.pdf (accessed on 12 October 2026).

Jin, C., & Tian, X. (2024). Enhanced safe-haven status of Bitcoin: Evidence from the Silicon Valley Bank collapse. *Finance Research Letters*, *59*, 104689. [CrossRef]

Koenker, R., & Bassett, G., Jr. (1978). Regression quantiles. *Econometrica*, *46*(1), 33–50. [CrossRef]

Liu, J., & Valcarcel, V. J. (2024). Hedging inflation expectations in the cryptocurrency futures market. *Journal of Financial Stability*, *70*, 101205. [CrossRef]

Nakamoto, S. (2008, October). *Bitcoin: A peer-to-peer electronic cash system*. Available online: https://bitcoin.org/bitcoin.pdf (accessed on 8 October 2026).

Pastén-Henríquez, B., Tapia-Griñen, P., & Sepúlveda-Velásquez, J. (2025). Gold and cryptocurrencies as safe-havens: Lessons from wartime. *Finance Research Letters*, *79*, 107230. [CrossRef]

Phochanachan, P., Pirabun, N., Leurcharusmee, S., & Yamaka, W. (2022). Do bitcoin and traditional financial assets act as an inflation hedge during stable and turbulent markets? Evidence from high cryptocurrency adoption countries. *Axioms*, *11*(7), 339. [CrossRef]

Rodriguez, H., & Colombo, J. (2025). Is bitcoin an inflation hedge? *Journal of Economics and Business*, *133*, 106218. [CrossRef]

Smales, L. A. (2024). Cryptocurrency as an alternative inflation hedge? *Accounting & Finance*, *64*(2), 1589–1611. [CrossRef]

Tarchella, S., Khalfaoui, R., & Hammoudeh, S. (2024). The safe haven, hedging, and diversification properties of oil, gold, and cryptocurrency for the G7 equity markets: Evidence from the pre-and post-COVID-19 periods. *Research in International Business and Finance*, *67*, 102125. [CrossRef]

Wu, S., Tong, M., Yang, Z., & Derbali, A. (2019). Does gold or Bitcoin hedge economic policy uncertainty? *Finance Research Letters*, *31*, 171–178. [CrossRef]

Yae, J., & Tian, G. Z. (2024). Volatile safe-haven asset: Evidence from Bitcoin. *Journal of Financial Stability*, *73*, 101285. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.