# Learning from Limited History: Recency Bias and Risk Aversion to Bitcoin*

**Yongjin Kim**
Anderson School of Management
University of New Mexico†

**Lars-Alexander Kuehn**
Tepper School of Business
Carnegie Mellon University‡

July 30, 2026

**Abstract**

We study how investors form risk attitudes toward novel assets, such as Bitcoin, when historical information is limited. Our portfolio-choice framework shows that risk aversion evolves endogenously as investors update perceived expected returns and adjust portfolios accordingly. Using Bitcoin option prices, we recover the time series of Bitcoin-specific risk aversion and show that its fluctuations are strongly linked to perceived expected returns shaped by past realizations. A key feature of belief formation is recency bias: investors place disproportionate weight on recent returns when forming expectations. This bias is markedly stronger in the Bitcoin market than in equities, reflecting its novelty and structural uncertainty.

**JEL Codes:** G11, G41

**Keywords**: Risk Aversion; Bitcoin; Belief Formation; Recency Bias; Portfolio Choice.

*We thank seminar participants at Carnegie Mellon University and the University of New Mexico, as well as conference participants, for valuable feedback. We also thank the Secure Blockchain Initiative at CMU's Cylab for financial support. Lars-A. Kuehn acknowledges financial support from the Jason Loh Fellowship at the Tepper School of Business. Contacts: yjkim99@unm.edu and kuehn@cmu.edu

†1 University of New Mexico, Albuquerque, NM 87131

‡5000 Forbes Avenue, Pittsburgh, PA 15213

# Contents


**1 Introduction** **1**

**2 Model** **7**
2.1 Portfolio Choice . . . 7
2.2 Absolute Risk Aversion . . . 9
2.3 Recovering Risk Aversion . . . 11
2.4 Learning with Recency Bias . . . 12

**3 Estimation** **14**
3.1 Data . . . 15
3.2 Estimation of Risk-Neutral Probability . . . 17
3.3 Estimation of Physical Probability . . . 18
3.4 Summary Statistics . . . 19

**4 Empirical Evidence** **20**
4.1 Perception of Expected Returns . . . 20
4.2 Estimates of Security-Specific Risk Aversion . . . 21
4.3 Determinats of Security-Specific Risk Aversion . . . 22
4.3.1 Bitcoin-Specific Risk Aversion . . . 23
4.3.2 Equity-Specific Risk Aversion . . . 24
4.4 Determinats of Expected Risk Aversion . . . 25
4.5 Estimation of Recency Bias in Belief Formation . . . 26
4.6 Robustness Test: Alternative Measures of Risk Aversion . . . 29
4.6.1 Risk-Aversion Estimation without Recency Bias . . . 29
4.6.2 Alternative Data Frequency and Sample Period . . . 31

**5 Conclusion** **31**

**A Proofs of the Portfolio Choice Propositions** **48**
A.1 Proof of Proposition 1 . . . 48
A.2 Proof of Proposition 2 . . . 48
A.3 Proof of Proposition 3 . . . 49
A.4 Proof of Proposition 4 . . . 49
A.5 Proof of Proposition 5 . . . 50

**B Learning with Recency Bias** **50**

**C Nonlinear Least Squares Estimation** **51**

**D Robustness Test - Alternative Data Frequency and Sample Period** **52**

# 1 Introduction

The arrival of a new financial asset, such as Bitcoin, expands investors' opportunity sets but also exposes them to unfamiliar risks and limited historical evidence. In this setting, belief formation and portfolio choice are tightly linked: beliefs about the asset's return distribution shape how much investors hold of it, and portfolio exposure in turn affects how averse they become to losses. This feedback between beliefs, portfolio demand, and risk attitudes is the central mechanism of this paper. A key feature of the belief formation is recency bias: investors place disproportionate weight on recent returns when updating their expectations.[1] We find that this bias is markedly stronger in the Bitcoin market than in equities, reflecting its novelty and perceived structural uncertainty.

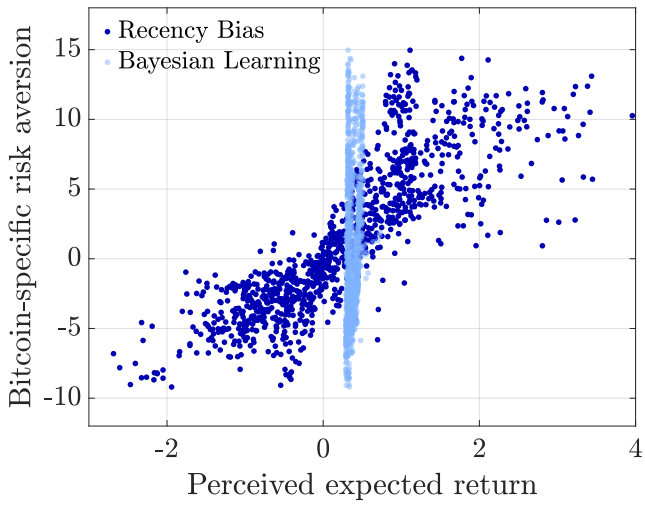
Our analysis starts with a simple portfolio-choice model in which investors' risk attitudes toward a security are shaped by the exposure they choose to hold. In the model, an investor chooses asset allocations to maximize expected utility. The optimal portfolio depends on the curvature of the utility function and the return distribution. When investors perceive an asset to offer higher expected returns or lower risk, they allocate more wealth to it. In turn, portfolio exposure shapes absolute risk aversion. When a larger share of wealth is allocated to a security, the investor becomes more averse to its future loss.

A key feature of our analysis is non-Bayesian learning about asset-specific expected returns. We adopt a flexible learning framework to model how investors update their beliefs in response to new information. Specifically, investors may assign different weights to past observations based on their temporal distance, as they would in heuristic learning. This formulation nests various learning modes, such as Bayesian updating, recency bias, and primacy bias, in which earlier observations exert greater influence. By estimating the model, we let the data reveal which belief-formation process best characterizes investors' perception.

To identify investors' latent risk attitudes, we use option prices, which reflect how investors value payoffs across future states of the economy. Option prices pin down the risk-neutral distribution of a security's return. By comparing the risk-neutral distribution to the physical distribution of returns, which is shaped by investors' subjective belief formation, we can recover investors' projected marginal utility across future return states. We define asset-specific risk aversion as the local rate at which this projected marginal utility declines with the security return. A higher value means that marginal utility falls more steeply as returns

[1] Closely related papers include Malmendier and Nagel (2011); Greenwood and Shleifer (2014); Barberis et al. (2015); Nagel and Xu (2022)

**Figure 1: Bitcoin-Specific Risk Aversion and Belief Formation**




This figure shows a scatter plot of Bitcoin-specific risk aversion against the perceived expected return on Bitcoin under Bayesian learning and recency-biased learning.

rise, so the investor values an extra dollar disproportionately more in adverse states, reflecting stronger aversion to losses on that asset.[2]

Figure 1 summarizes our main empirical findings. The figure shows a scatter plot of Bitcoin-specific risk aversion against the perceived expected return on Bitcoin under Bayesian learning and recency-biased learning. It shows a strong positive association between risk aversion and the perceived expected return formed under recency-biased learning. In contrast, the expected return derived from Bayesian learning remains largely unchanged over the sample, and its relationship with risk aversion is markedly weaker. This stark contrast underscores the central role of heuristic learning – particularly recency bias – in shaping investors' risk attitudes toward Bitcoin.

Our data consist of Bitcoin options, futures, and spot prices from January 2017 to June 2024. We obtain these data from Deribit, a centralized cryptocurrency derivatives exchange operating on a continuous, 24/7 basis. Deribit is the dominant venue for Bitcoin options, consistently accounting for more than 80–85% of global trading volume and open interest

[2]We are not the first to use option prices to estimate risk aversion. Ait-Sahalia and Lo (2000) and Jackwerth (2000) provide the theoretical foundation for this approach, and subsequent studies have applied it to various asset classes. Our contribution is to apply this method to Bitcoin and to link the estimated risk aversion to belief formation and portfolio choice.

in cryptocurrency options. Its central role in the market is reflected in a growing body of academic research that relies on Deribit data to study implied volatility, illiquidity, and trading activity in Bitcoin options (e.g., Alexander et al. (2024); Atanasova et al. (2024); Chen et al. (2024)).

We estimate Bitcoin-specific risk aversion as follows. First, using intraday option trades, we recover the daily risk-neutral distribution of one-month-ahead returns. We then construct the corresponding physical return distribution. Specifically, under a given belief-formation rule – either Bayesian or heuristic learning – we aggregate historical spot returns and infer the conditional perceived probabilities. Under no-arbitrage, the ratio of risk-neutral to physical probabilities identifies investors' marginal utility[3] across future return states and thereby reveals their risk attitudes. In particular, the slope of log marginal utility with respect to future returns yields Bitcoin-specific absolute risk aversion evaluated at a given state. We construct two types of risk-aversion measures. One measure evaluates local risk aversion at two representative future outcomes: 5% and –5% returns over the subsequent month. The second provides a more comprehensive characterization by computing expected risk aversion, defined as the integral of local risk aversion over the support of future returns. Applying this procedure at the daily frequency, we construct the time series of Bitcoin-specific risk aversion.

For comparison, we conduct an analogous analysis for the equity market. We recover equity-specific risk aversion using options on the S&P 500 index from OptionMetrics, together with intraday trades of S&P 500 mini futures obtained from TickData.

Under the physical probabilities influenced by Bayesian learning, estimated Bitcoin-specific risk aversion evaluated at a future return of –5% is 5.55 on average, compared with 35.89 for the equity market. This contrast is consistent with the intuition that investors are relatively less concerned about Bitcoin losses, as it typically constitutes only a small share of their portfolios. Bitcoin-specific risk aversion also exhibits substantial time variation, with a standard deviation of 16.31. These fluctuations are largely independent of equity-specific risk aversion: the correlation coefficient between the two series is only 0.11. Moreover, we show that the specification of the learning process materially affects the dynamics of investors' risk attitudes. Introducing a recency bias generates substantially greater fluctuations in estimated risk aversion.

Empirically, we find that investors' aversion to Bitcoin's risk – whether measured locally or

[3]Formally, our estimate corresponds to the projection of marginal utility onto the space spanned by Bitcoin returns.

as an expectation over the return support – is strongly related to their perceptions of its return distribution formed from past observations. Consistent with our model's prediction, estimated risk aversion increases with the perceived conditional average return and decreases with its volatility. Both relationships are statistically significant at the 1% level and remain robust after controlling for additional determinants of risk preferences identified in the literature. Furthermore, these perceived moments explain between 71.6% and 76.1% of the time-series variation in Bitcoin-specific risk aversion, especially when recency bias is incorporated into the estimation. All the evidence jointly supports our central mechanism: investors' risk attitudes toward Bitcoin evolve as they update their beliefs about return prospects and adjust their portfolio allocations accordingly.

We also document a similar regularity in the equity market. Investors' risk aversion to potential loss on the S&P 500 is strongly tied to both the perceived average return and the conditional volatility of the equity index, with the goodness of fit ranging from 39.5% to 62.2%. Taken together, the evidence from both markets indicates that asset demand arising from portfolio allocation is a primary force shaping investors' risk attitudes across different asset classes. More broadly, our mechanism offers an explanation for fluctuations in risk aversion that are often treated as exogenous in much of the asset-pricing literature (e.g., Campbell and Cochrane (1999)).

Guided by the model's implied link between perceived expected returns and risk attitudes, we identify the belief-formation rule that best explains the observed dynamics of Bitcoin-specific risk aversion. The nonlinear least squares estimates indicate that investors place disproportionately greater weight on recent returns, with the weights decaying exponentially as observations recede into the past. Specifically, a return observed 108 days earlier receives only one-tenth the weight assigned to the most recent return. In contrast, belief formation under a Bayesian framework yields significantly weaker empirical connections between perceived average and risk aversion, suggesting that Bayesian updating alone is insufficient to account for the observed dynamics of Bitcoin-specific risk aversion.

In addition, we find that investors in the equity market also exhibit heuristic learning. Similar to our findings for Bitcoin, the perceived average return constructed with recency bias provides the best fit for explaining fluctuations in equity-specific risk aversion. Yet the degree of recency bias is substantially weaker. In the equity market, the weights assigned to past observations decay much more gradually: a return observed 287 days prior receives one-tenth the weight of the most recent return, compared with only 108 days in the Bitcoin

market. Notably, our estimated rate of weight decay in the equity market is quantitatively consistent with the evidence provided by Greenwood and Shleifer (2014). Using survey-based expected returns, they find that weights on past returns decay exponentially to 9.8% over four quarters in the formation of current expectations. This close alignment lends external validity to our option-based analysis and reinforces the ensuing conclusion regarding the differing magnitudes of recency bias across markets.

Why does the Bitcoin market exhibit a much stronger recency bias, despite its already limited data availability? We conjecture that this short-horizon focus reflects investors' expectations of large and frequent structural changes in the Bitcoin market – stemming from regulatory uncertainty and evolving adoption. In such an environment, investors may heuristically anchor their expectations to the most recent data to capture rapidly changing fundamentals.

### Related Literature

This paper connects three strands of the literature: the recovery of risk attitudes from option prices, behavioral models of belief formation, and the emerging literature on cryptocurrency asset pricing.

A large body of work recovers risk preferences from option-implied information. Ait-Sahalia and Lo (1998), Ait-Sahalia and Lo (2000), Jackwerth (2000), Rosenberg and Engle (2002), Bliss and Panigirtzoglou (2004), Cuesdeanu and Jackwerth (2018), Beason and Schreindorfer (2022), and Schreindorfer and Sichert (2025) estimate the stochastic discount factor or risk aversion implied by options on the stock market index. Beyond the stock market, Christoffersen et al. (2021) analyze options on crude oil futures and document the commodity's unique risk characteristics. In contrast to these studies, our focus is on the Bitcoin option market and its implications for the newly introduced cryptocurrency, which recent studies, including Alexander and Imeraj (2021) and Winkel and Härdle (2023), have begun to explore. We contribute to the literature by providing the dynamics of risk aversion toward Bitcoin, along with a theoretical framework that explains its dynamics.

Our paper is also related to Martin (2017) and Martin and Wagner (2019), who derive lower bounds on expected returns using option prices. In a model-free framework that relies only on a weak preference restriction – the negative correlation condition – they construct a lower bound solely from option prices and establish it as a strong conditional measure of expected returns. Our analysis is complementary in two respects. First, rather than bounding

expected returns, we directly estimate projected marginal utility and recover measures of risk aversion implicit in option prices, allowing us to study the evolution of investors' risk attitudes. Second, we focus on a newly emerging asset that represents only a small fraction of the market portfolio, a setting in which the negative correlation condition is less likely to hold.[4] Nevertheless, we find that our risk-aversion measure – constructed using both option prices and historical spot returns – retains strong empirical validity.

Our paper also builds on the behavioral learning literature. In their foundational work, Tversky and Kahneman (1974) document that decision-makers facing complex problems rely on simplifying heuristics, notably the representativeness and availability heuristics. Gennaioli and Shleifer (2010) formalize the representativeness heuristic, while Bordalo et al. (2018) and Bordalo et al. (2019) develop diagnostic expectations, in which salient or representative news makes certain future states more cognitively accessible, leading agents to overweight those states. This mechanism rationalizes several stylized facts about credit cycles and stock market fluctuations.

In contrast, our learning process is more closely related to the availability heuristic, which leads people to overweight information that is easily accessible – in particular, recent observations. Malmendier and Nagel (2011) and Malmendier and Nagel (2016) provide evidence that individuals form expectations about market returns and inflation based on their lifetime experiences, with disproportionate weight placed on more recent experiences. Nagel and Xu (2022) formalize this idea through a fading-memory model and study its implications for aggregate asset pricing. Building on their flexible framework, we model investors' belief formation about Bitcoin's average return and provide new empirical evidence of recency bias in cryptocurrency markets. A comprehensive overview of economic research on subjective memory and belief formation is provided by Malmendier and Wachter (2024).

Closely related is the literature on return extrapolation. Greenwood and Shleifer (2014) document that survey-based expectations of stock returns are positively correlated with past returns, consistent with investors extrapolating from recent experience. Barberis et al. (2015) develop a model in which extrapolative investors generate asset price booms and busts, while Barberis et al. (2018) show that such behavior can account for boom-bust patterns in asset prices more broadly. Da et al. (2021) show return extrapolation in the cross section of stocks

[4]The negative correlation condition is $\text{cov}_t\left(M_T R_T, R_T\right) < 0$, where $M_T$ is the stochastic discount factor and $R_T$ is an asset's gross return. Martin (2017) notes that this condition would fail for hedge assets such as gold. In addition, when asset returns are independent of or weakly negatively related to the stochastic discount factor, the condition does not hold.

by using a crowdsourcing platform that ranks stocks based on participants' perception of future returns. Cassella and Huber (2020) link extrapolation bias to the predictability of stock returns. On the theoretical side, Adam et al. (2017) show that learning about expected capital gains generates realistic boom-bust dynamics, and Collin-Dufresne et al. (2016) study how parameter learning amplifies stock market volatility. Our contribution relative to this literature is twofold. First, rather than assuming a particular degree of extrapolation, we estimate the recency-bias parameter directly from option-implied risk aversion. Second, by comparing Bitcoin and equity markets, we document that recency bias is substantially stronger for novel assets with limited history, providing new cross-asset evidence on the determinants of heuristic learning.

Finally, our paper relates to the growing literature on cryptocurrency markets. Liu and Tsyvinski (2020) identify risk factors specific to cryptocurrency returns, which we employ as controls. Makarov and Schoar (2020) document trading frictions and arbitrage patterns in crypto markets, while Cong et al. (2021) and Sockin and Xiong (2023) study the role of cryptocurrencies in decentralized platforms. We complement these studies by analyzing the risk-aversion dynamics specific to Bitcoin and showing that they are shaped by heuristic learning about the asset's return distribution.

# 2 Model

In this section, we present a model in which risk aversion toward Bitcoin arises endogenously from the investor's portfolio allocation decision. We begin by deriving the investor's optimal allocation that maximizes expected utility. This optimal allocation, in turn, shapes the stochastic discount factor, which captures the investor's risk attitudes toward the future returns of all securities, including Bitcoin. A central prediction of the model is that the investor's risk aversion to Bitcoin depends on her perception of its conditional return distribution. We then introduce a learning framework that characterizes how these perceptions are formed and updated over time.

## 2.1 Portfolio Choice

We consider an investor who has a unit of wealth at time $t$ and can invest in both risky assets and a risk-free asset. The portfolio return is given by

$$R_{p,t+1} = \mathbf{w}_t^\top \mathbf{R}_{t+1} + R_{f,t} \tag{1}$$

where $\mathbf{w}_t$ denotes the vector of portfolio weights among risky assets, $\mathbf{R}_{t+1}$ the vector of excess returns, and $R_{f,t}$ the risk-free rate of return.

The investor maximizes expected utility, $\mathbb{E}_t[u(1+R_{p,t+1})]$, where the utility is increasing and concave. To derive analytical solutions, we approximate the utility function using a second-order Taylor expansion[5] around $1+R_{f,t}$

$$u(1+R_{p,t+1}) \approx u(1+R_{f,t}) + u'(1+R_{f,t})\mathbf{w}_t^\top \mathbf{R}_{t+1} + \frac{u''(1+R_{f,t})}{2}\mathbf{w}_t^\top \mathbf{R}_{t+1}\mathbf{R}_{t+1}^\top \mathbf{w}_t. \quad (2)$$

**Proposition 1** *Given the second-order Taylor expansion of the utility function, the optimal asset allocation is given by*

$$\mathbf{w}_t^* = -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})}\left(\mathbf{\Sigma}_t + \boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)^{-1}\boldsymbol{\mu}_t \quad (3)$$

*where $\boldsymbol{\mu}_t$ denotes the vector of perceived conditional expected excess returns and $\mathbf{\Sigma}_t$ the perceived conditional covariance matrix of excess returns.*

All proofs can be found in Appendix A. The optimal allocation depends on the average and variance of security returns, as in the mean-variance optimization.[6] In addition, when the investor is more risk-averse, as indicated by a higher curvature of the utility function, $-u''(1+R_{f,t})/u'(1+R_{f,t})$, the investor allocates less to the risky assets, consistent with intuition.

In the following, we consider an investor with access to two risky assets: Bitcoin and the stock market. Their excess returns are denoted by $R_{i,t+1}$, where $i \in \{B, S\}$, with $B$ representing Bitcoin and $S$ representing the stock index. We denote the fraction of wealth invested in risky assets by $\mathbf{w}_t = [w_{B,t}, w_{S,t}]^\top$, and the conditional moments of the perceived distribution of $\mathbf{R}_{t+1} = [R_{B,t+1}, R_{S,t+1}]^\top$ are given by

$$\boldsymbol{\mu}_t = \begin{bmatrix}\mu_{B,t}\\ \mu_{S,t}\end{bmatrix} \text{ and } \mathbf{\Sigma}_t = \begin{bmatrix}\sigma_{B,t}^2 & 0\\ 0 & \sigma_{S,t}^2\end{bmatrix}. \quad (4)$$

For simplicity, we assume that the two risky assets are conditionally uncorrelated. The following proposition provides the optimal allocation in this setting.

[5] Note that $1+R_{f,t}$ denotes the gross risk-free return, which is strictly greater than one. In contrast, the term $R_{p,t+1} - R_{f,t} = \mathbf{w}_t^\top \mathbf{R}_{t+1}$ represents a weighted average of excess returns, which is typically much smaller than one in practice. This disparity in magnitudes justifies using a second-order Taylor expansion as a reasonable approximation of the utility function.

[6] The optimal allocation is slightly different from the mean-variance solution $w_t^{\text{MV}} \propto \mathbf{\Sigma}_t^{-1}\boldsymbol{\mu}_t$. The presence of the term $\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top$ reflects that the second-order approximation of expected utility is taken around the risk-free payoff rather than the mean portfolio payoff. This choice simplifies our analysis of the allocation dependence on return distribution under the general form of the utility. When expected utility is expanded around the mean payoff, the optimal allocation converges to the standard mean-variance solution.

**Proposition 2** *Consider an economy with two risky assets, $i \in \{B, S\}$, and a risk-free asset. The utility-maximizing investor chooses the following optimal portfolio weights*

$$w_{i,t}^* = -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})} \frac{\mu_{i,t}}{\sigma_{i,t}^2 + \mu_{i,t}^2 + \mu_{-i,t}^2 \sigma_{i,t}^2 / \sigma_{-i,t}^2} \tag{5}$$

*where $-i$ indicates the other risky asset.*

These analytic expressions enable us to examine how the optimal allocation changes in response to shifts in the conditional return distribution. The following proposition describes the relationship.

**Proposition 3** *The optimal portfolio weight for risky asset $i \in \{B, S\}$ increases with the conditional expected excess return, i.e.,*

$$\frac{\partial w_{i,t}^*}{\partial \mu_{i,t}} > 0 \tag{6}$$

*if $\sigma_{i,t} > \mu_{i,t}$. In addition, the portfolio weight decreases with the conditional volatility, i.e.,*

$$\frac{\partial w_{i,t}^*}{\partial \sigma_{i,t}^2} < 0 \tag{7}$$

*if $\mu_{i,t} > 0$.*

Under realistic parameter values, the optimal allocation aligns with intuition: when a security return is higher on average or less volatile, the investor optimally allocates more wealth in that security.

## 2.2 Absolute Risk Aversion

Having obtained the optimal allocation, we next characterize the investor's risk attitudes. The Arrow-Pratt measure of absolute risk aversion is defined as $\mathcal{R} = -\frac{u''(1+R)}{u'(1+R)}$. It is a measure of how strongly an investor dislikes risk, expressed as the curvature of their utility function relative to its slope. Alternatively, it can be expressed as the slope of log marginal utility, i.e., $\mathcal{R} = -\frac{\partial \log u'(1+R)}{\partial R}$.

To obtain a security-level measure of risk aversion, we project marginal utility onto the space spanned by the return of security $i$ and denote the resulting projection by $u'_{i,t}$:

$$u'_{i,t}(R_{i,t+1}) = \mathbb{E}_t \left[ u'(1 + \mathbf{w}_t^\top \mathbf{R}_{t+1} + R_{f,t}) \middle| R_{i,t+1} \right]. \tag{8}$$

The projected marginal utility is the conditional expectation of marginal utility given $R_{i,t+1}$, integrating out the realizations of all other securities. It correctly prices any payoff measurable

with respect to $R_{i,t+1}$ – in particular, options written on security $i$ – by the absence of arbitrage.[7]

We define the measure of absolute risk aversion specific to security $i$ as the local rate at which the investor's projected marginal utility declines with the security's return:

$$\mathcal{R}_{i,t}(R_{i,t+1}) = -\frac{\partial \log u'_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}}. \tag{9}$$

This measure shows how strongly the investor's marginal valuation responds to fluctuations in security $i$'s return – equivalently, the curvature of the projected log pricing kernel – with a higher $\mathcal{R}_i$ indicating stronger aversion to losses on security $i$.

It is worth clarifying that security-specific risk aversion $\mathcal{R}_i$ is a derived object, not a primitive preference parameter. The investor has a single utility function $u$. The key implication is that $\mathcal{R}_i$ varies across securities held by the same investor – not because preferences differ, but because each security occupies a different position in the portfolio. We make this dependence explicit in the next proposition.

**Proposition 4** *Given the optimal weights $w^*_{i,t}$ for $i \in \{B, S\}$, absolute risk aversion contingent on the future return $R_{i,t+1}$ equals*

$$\mathcal{R}_{i,t}(R_{i,t+1}) = -\frac{u''(1+R_{f,t})w^*_{i,t}}{u'(1+R_{f,t}) + u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t} + w^*_{i,t}R_{i,t+1})} \tag{10}$$

*where $-i$ indicates the other risky asset. It follows that the risk aversion increases with the portfolio weight, i.e.,*

$$\frac{\partial \mathcal{R}_{i,t}(R_{i,t+1})}{\partial w^*_{i,t}} > 0. \tag{11}$$

Intuitively, when more wealth is allocated to a security, the investor is more risk-averse to its potential loss in subsequent periods. Next, building on Propositions 3 and 4, we predict how changes in the conditional return distribution should influence risk aversion.

**Proposition 5** *For the utility-maximizing investor, absolute risk aversion to security $i \in \{B, S\}$ increases with the conditional expcected excess return, i.e.,*

$$\frac{\partial \mathcal{R}_{i,t}(R_{i,t+1})}{\partial \mu_{i,t}} > 0 \tag{12}$$

[7] Alternatively, projected marginal utility could be defined as the projected pricing kernel

$$u'_{i,t}(R_{i,t+1}) \propto \mathbb{E}_t\left[M(R_{p,t+1})|\, R_{i,t+1}\right]$$

because the pricing kernel is proportional to marginal utility, i.e., $M(R_{p,t+1}) \propto u'(1+R_{p,t+1})$.

*if* $\sigma_{i,t} > \mu_{i,t}$. *In addition, absolute risk aversion to security* $i \in \{B, S\}$ *decreases with the conditional volatility, i.e.,*

$$\frac{\partial \mathcal{R}_{i,t}(R_{i,t+1})}{\partial \sigma_{i,t}^2} < 0 \tag{13}$$

*if* $\mu_{i,t} > 0$.

This result follows directly from combining the preceding two propositions. When a security is perceived to offer a higher expected excess return or lower volatility, the investor optimally increases her allocation to that asset. Portfolio exposure, in turn, shapes subsequent risk aversion: having allocated a larger share of wealth to a security, the investor becomes more averse to potential losses in that security. The model prediction in Proposition 5 serves as the primary hypothesis we test empirically.

### 2.3 Recovering Risk Aversion

The security-specific risk aversion $\mathcal{R}_{i,t}$ characterized in Proposition 4 is defined through the projected marginal utility $u'_{i,t}(R_{i,t+1})$ and therefore appears, at first sight, to require knowledge of the investor's utility primitive. In this section, we show that $\mathcal{R}_{i,t}$ can in fact be identified from two observable distributions: the risk-neutral distribution of $R_{i,t+1}$ implied by option prices, and the physical distribution perceived by the investor.

By construction, projected marginal utility prices any payoff measurable with respect to $R_{i,t+1}$. In particular, applying the law of iterated expectations, the investor's Euler equation for the excess return satisfies $\mathbb{E}_t\big[R_{i,t+1}\, u'_{i,t}(R_{i,t+1})\big] = 0$. More generally, for any contingent claim $X(R_{i,t+1})$ on security $i$, the investor's Euler equation can be stated as

$$\mathbb{E}_t\big[X(R_{i,t+1})\, u'_{i,t}(R_{i,t+1})\big] = \frac{1}{1+R_{f,t}}\, \tilde{\mathbb{E}}_t[X(R_{i,t+1})] = 0, \tag{14}$$

where $\tilde{\mathbb{E}}_t$ denotes the expectation under the risk-neutral measure. Writing both expectations as integrals against the corresponding densities and matching integrands, equation (14) implies that the projected marginal utility is proportional to the likelihood ratio of the risk-neutral and the physical density:

$$u'_{i,t}(R_{i,t+1}) = \frac{1}{1+R_{f,t}}\, \frac{\tilde{\mathbb{P}}_{i,t}(R_{i,t+1})}{\mathbb{P}_{i,t}(R_{i,t+1})}, \tag{15}$$

where $\tilde{\mathbb{P}}_{i,t}(\cdot)$ and $\mathbb{P}_{i,t}(\cdot)$ denote, respectively, the risk-neutral and physical conditional densities of $R_{i,t+1}$. The proportionality constant $1/(1+R_{f,t})$ ensures that $u'_{i,t}$ correctly prices the risk-free payoff. Although this projection onto the space spanned by security $i$'s return cannot

recapture the full stochastic discount factor, by construction it preserves all information relevant for pricing claims on security $i$ and, consequently, for measuring the investor's risk aversion to that security.

Combining (15) with the definition of security-specific absolute risk aversion in (9) yields the central identification result of the paper:

$$\mathcal{R}_{i,t}(R_{i,t+1}) \;=\; -\frac{\partial \log u'_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}} \;=\; \frac{\partial \log \mathbb{P}_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}} \;-\; \frac{\partial \log \tilde{\mathbb{P}}_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}}. \tag{16}$$

The investor's absolute risk aversion toward security $i$ is therefore the gap between the score of the physical and risk-neutral density, evaluated locally at the future return $R_{i,t+1}$. Intuitively, the risk-neutral density is a probability-weighted reweighting of the physical density by marginal utility; the steeper that reweighting is in $R_{i,t+1}$, the more strongly the investor dislikes losses on security $i$, and the larger is $\mathcal{R}_{i,t}$.

Equation (16) suggests a natural two-step empirical strategy, which we implement in Section 3. First, we recover $\tilde{\mathbb{P}}_{i,t}$ from a panel of option prices using the Breeden and Litzenberger (1978) result. Second, we recover $\mathbb{P}_{i,t}$ from the time series of realized spot returns, taking seriously the fact that what enters (16) is not the econometrician's best estimate of the return-generating process but the distribution *perceived* by the investor. The next section is devoted to specifying how this perceived distribution is formed.

### 2.4 Learning with Recency Bias

The investor's wealth allocation to a security and subsequent risk aversion depend on her perception of the return distribution. In this section, we explicitly describe how the perceived average is formed and updated from realized data. We focus on the expected return as a quantity to be learned because it has long been recognized that determining the conditional average is significantly more challenging than the conditional volatility, as noted by Merton (1980). In contrast, volatility tends to be more persistent, and models such as ARCH and GARCH have effectively characterized it.

Specifically, we formulate the learning about the average return, using the non-Bayesian framework suggested by Nagel and Xu (2022). The reason for this modeling choice is twofold. First, a growing number of studies testify to the presence of heuristic belief formation in financial markets, including Gennaioli and Shleifer (2010), Malmendier and Nagel (2011), Malmendier and Nagel (2016), Bordalo et al. (2018), Bordalo et al. (2019), and Nagel and Xu (2022). Second, this approach offers the flexibility to incorporate rational Bayesian updates

as well, with a simple change to a model parameter.

Similar to Collin-Dufresne et al. (2016), we assume that asset $i$'s excess return is given by

$$R_{i,t} = m_i + \sigma_{i,t-1}\epsilon_{i,t} \tag{17}$$

where $m_i$ denotes the unconditional true average, $\sigma_{i,t-1}$ the conditional volatility, and $\epsilon_{i,t}$ is an independent standard normal innovation. A key assumption here is that the parameter $m_i$ is constant over time but unknown. Facing this uncertainty, the investor learns about $m_i$ from data on realized excess returns.

Initially, the investor has prior beliefs about the true average return, denoted by $\pi(m_i)$, which are normally distributed with mean $\mu_{i,0}$ and standard deviation $\varphi_{i,0}$. As realized returns are observed over time, the investor updates her beliefs considering the likelihood of these returns. Let $p(R_{i,s}|m_i)$ denote the sampling probability of date-$s$ return conditional on the average return $m_i$. We allow belief updating to deviate from standard Bayesian learning by incorporating the potential influence of heuristic bias. Specifically, we assume a flexible learning framework in which past observations may receive different weights depending on their temporal distance from the belief-formation date:

$$\pi(m_i|R_{i,1},..,R_{i,t};\nu_i) \propto \prod_{s=1}^{t} p(R_{i,s}|m_i)^{(1-\nu_i)^{t-s}} \times \pi(m_i)^{(1-\nu_i)^t}, \tag{18}$$

where the exponent $(1-\nu_i)^{t-s}$ determines the weight on the return observed $t-s$ days before the belief formation on date $t$. If we set $0 < \nu_i < 1$, the weights on more distant past observations decay exponentially, reflecting recency bias in the perceived average. On the other hand, by setting $\nu_i$ equal to zero, we replicate the Bayesian learning, where all observations receive equal weights irrespective of observation dates.

Our learning framework is closely related to the fading-memory model of Nagel and Xu (2022), in which older observations receive exponentially declining weights. They characterize beliefs in the limit of an infinite history and a fixed weight prior. Relative to their formulation, we explicitly model learning from a finite and growing sample of observations and allow the influence of prior beliefs to decay over time. These modifications are particularly relevant in the context of Bitcoin, where investors face limited historical information and may gradually discard both early observations and initial priors as new information arrives.

In the following proposition, we summarize the dynamics of the posterior.

**Proposition 6** *Given a normally distributed posterior belief with mean $\mu_{i,t}$ and standard deviation $\varphi_{i,t}$, with newly realized return $R_{i,t+1}$, the conditional mean and standard deviation are updated in the following recursive way:*

$$\mu_{i,t+1} = \varphi^2_{i,t+1}\left[\frac{R_{i,t+1}}{\sigma^2_{i,t}} + (1-\nu_i)\frac{\mu_{i,t}}{\varphi^2_{i,t}}\right] \tag{19}$$

$$\frac{1}{\varphi^2_{i,t+1}} = \frac{1}{\sigma^2_{i,t}} + (1-\nu_i)\frac{1}{\varphi^2_{i,t}}. \tag{20}$$

The proof can be found in Appendix B. In belief updates, the posterior mean is a weighted average of the prior mean $\mu_{i,t}$ and the new observation $R_{i,t+1}$. The weights are determined by the recency-bias parameter $\nu_i$, the parameter uncertainty $\varphi_{i,t}$, and the volatility of noise $\sigma^2_{i,t}$. Intuitively, the belief update is more sensitive to the new observation when investors display a stronger recency bias (i.e., $\nu_i$ is higher) or when the latest signal is more informative (i.e., $\sigma_{i,t}$ is relatively lower compared to $\varphi_{i,t}$).

It is noteworthy that the recency bias results in perpetual updates of beliefs, distinct from the Bayesian process. In the latter, learning becomes complete after infinitely many observations; if $\nu_i = 0$, the precision of beliefs $1/\varphi^2_{i,t}$ increases monotonically with every new observation and eventually becomes infinity after a long process of learning. At that point, a new observation has almost zero impact on belief revisions. In contrast, when the recency bias is in effect with $\nu_i > 0$, old data becomes obsolete in memory, and the precision of beliefs never approaches infinity, according to equation (20).[8] Consequently, the posterior mean continues to be revised in response to new observations, even when past data is plentiful. In the empirical analysis section, we illustrate this distinction in the perceived average of Bitcoin returns between the two forms of learning.

Given the belief dynamics in Proposition 6, the investor's risk aversion toward security $i$ depends on the recency-bias parameter $\nu_i$ through the investor's subjective belief formation

$$\mathcal{R}_{i,t}(R_{i,t+1};\nu_i) = \frac{\partial \log \mathbb{P}_{i,t}(R_{i,t+1};\nu_i)}{\partial R_{i,t+1}} - \frac{\partial \log \tilde{\mathbb{P}}_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}}. \tag{21}$$

## 3 Estimation

In this section, we first describe the data on Bitcoin, equity index options, and other relevant variables. Next, we outline how to estimate the projection of the marginal utilities onto the space spanned by security returns and recover the implied risk aversion.

[8] If the noise volatility is constant at $\sigma_i$, the precision of beliefs $1/\varphi^2_{i,t}$ converges to $1/\left(\nu_i\sigma_i^2\right)$ and does not become infinitely high even when the number observations approaches infinity.

### 3.1 Data

Our data consists of option prices for two underlying securities – Bitcoin and the S&P 500 index – along with their respective spot and futures prices. To ensure consistency and reliability, we preprocess the raw data by filtering out anomalies and normalizing it before estimation.

Data on Bitcoin options, futures, and spot prices are obtained from Deribit. Deribit is a cryptocurrency derivatives exchange specializing in futures and options trading and became part of Coinbase following its acquisition in May 2025. Originally founded in the Netherlands in June 2016, Deribit has since evolved into a global trading venue, with its headquarters now located in Dubai, United Arab Emirates.[9] Deribit is widely recognized as the dominant platform for Bitcoin and Ethereum options trading. Industry reports and market-data providers[10] consistently indicate that it accounts for approximately 80–85% of global trading volume and open interest in cryptocurrency options markets.

Panel A of Figure 2 compares Bitcoin option trading volumes on Deribit and the CME, showing that Deribit handles a substantially larger volume of contracts despite the presence of regulated venues. Reflecting its central role in cryptocurrency derivatives markets, a growing body of academic research relies on Deribit data to study implied volatility, illiquidity, and trading activity in Bitcoin options (e.g., Alexander et al. (2024); Atanasova et al. (2024); Chen et al. (2024)). We obtain historical option, futures, and spot data from CoinAPI.

Bitcoin options traded on Deribit are European-style and cash-settled. The contract multiplier is one, which means that each BTC call (put) option grants the right to buy (sell) one Bitcoin at the specified strike price. Both option prices and payoffs are denominated and settled in Bitcoin, the base currency. For example, consider a trader who purchases a call option with a strike price of \$100,000 for a premium of 0.05 BTC. If the Bitcoin spot price at expiration is \$125,000, the option payoff would be 0.2 BTC, calculated as $(125,000 - 100,000)/125,000$. These option prices can be converted to U.S. dollars by multiplying the Bitcoin-denominated price by Bitcoin's current spot price.

Bitcoin option data is collected at the transaction level from January 2017 to June 2024. To mitigate potential bias from outliers, we apply the following sampling criteria. First, we exclude data points with a time to expiration of less than 10 days or more than 180 days.

[9] According to Deribit support, the platform is not accessible to individuals and institutions located in 26 jurisdictions, including the United States.

[10] Sources include the 2025 industry report by CfC St. Moritz, a leading institutional digital-asset conference, and market analyses by OKX, a major participant in cryptocurrency derivatives markets.

Additionally, observations with extreme values of implied volatility, defined as those lower than 2% or higher than 200%, are excluded. Lastly, for each option contract, we calculate its volatility-adjusted moneyness[11] and exclude observations with moneyness below –2 or above 2. Figure 2 Panel B illustrates the daily trading volume of Bitcoin option contracts that satisfy all of these sampling criteria. The trading volume has grown substantially in recent years, and the median volume is 2,106 contracts per day.

The calculation of the implied volatility and moneyness for each option requires the futures price corresponding to its expiration date. For each option transaction, we collect futures trades that have occurred within 15 minutes before the option trade and fit this data to a second-order polynomial function of time to expiration in order to estimate the futures price.

The Bitcoin spot price is provided by Deribit. It collects real-time prices from multiple spot exchanges and computes an equally weighted average to construct the price index. We aggregate this tick-by-tick index into one-minute intervals, and these minute observations serve as intraday spot prices.

With respect to the S&P 500 index, we obtain its options data from OptionMetrics and apply the same sampling criteria used for Bitcoin options. The stock index's intraday spot returns are measured based on high-frequency data on S&P 500 mini futures provided by TickData. The tick-by-tick futures prices are aggregated into one-minute intervals, and we calculate the return for each minute.

As control variables, we obtain macroeconomic data from the Federal Reserve Bank of St. Louis (FRED). These include the growth rates of industrial production, personal consumption expenditures, and the consumer price index, all observed at the monthly frequency and expressed in annualized terms. For the indicator of overall financial market conditions, we also include the Chicago Fed National Financial Conditions Index. In addition, we incorporate Bitcoin-specific risk factors following the methodology of Liu and Tsyvinski (2020). Specifically, the attention indicator is constructed as the logarithm of one plus the Google search volume for "Bitcoin" and related terms. The network effect indicator is defined as the first principal component extracted from proxies such as the number of active accounts and transaction counts from Blockchain.info.

[11] For an option contract with time to expiration $\tau_{i,t}$, strike price $K_{i,t}$ and forward price $F_{i,t}$, the volatility-adjusted moneyness is

$$\frac{\log\left(K_{i,t}/F_{i,t}\right)}{\sigma_t\sqrt{\tau_{i,t}}}$$

where $\sigma_t$ is the median value of implied volatility observed on day $t$.

### 3.2 Estimation of Risk-Neutral Probability

In this section, we describe the estimation of risk-neutral probability $\tilde{\mathbb{P}}_{i,t}(R_{i,t+\tau})$, one of the key components required to measure the security-specific risk aversion in equation (16). We then turn to estimating the physical probability in the next section. In both cases, we focus on the probability of one-month-ahead returns with $\tau = 30$, in line with the VIX index.

Consider a European call option on security $i$ with strike price $K$ and expiration date $t+\tau$, of which the price at time $t$ is $C_{i,t}(K, t+\tau)$. The security's current spot price is $S_{i,t}$. According to Breeden and Litzenberger (1978), the second derivative of the option price determines the risk-neutral probability $\tilde{\mathbb{P}}_{i,t}$ that the underlying asset price becomes $S_{i,t+1} = S_{i,t}(1 + R_{f,t} + R_{i,t+1})$ in the future:[12]

$$\tilde{\mathbb{P}}_{i,t}(R_{i,t+\tau}) = \left(1 + R_{f,t}^{(\tau)}\right) S_{i,t} \left.\frac{\partial^2 C_{i,t}(K, t+\tau)}{\partial K^2}\right|_{K=S_{i,t+\tau}}. \tag{22}$$

To calculate the derivative from data at sparsely observed strike prices, we map the option price of contract $j$ into Black-Scholes implied volatility units, denoted by $\sigma_{t,j}^{\text{BS}}$, which vary with moneyness and time to maturity. We parameterize implied volatilities following the stochastic volatility inspired (SVI) model of Gatheral and Jacquier (2011). Despite its parsimony, the SVI is known to capture key features of implied volatility well, and it has been adopted in a growing number of studies, including Berger et al. (2019) and Beason and Schreindorfer (2022).

Specifically, we employ the modified SVI formulation suggested by Beason and Schreindorfer (2022), which allows us to pool option contracts with different maturities and use more data points than would otherwise be available in daily estimation. To formulate[13], we define the standardized moneyness, $\kappa_{t,\tau} = \frac{\log(K/F_{t,\tau})}{v_t\sqrt{\tau}}$, where $F_{t,\tau}$ is the forward price with the time to expiration $\tau$ and $v_t$ the conditional volatility measured from observed option prices.[14] The implied variance is modeled as a function of the moneyness and time to maturity as follows:

$$\sigma_t^{\text{SVI}}(\kappa, \tau; \Theta_t)^2 = a_{t,\tau} + b_{t,\tau}\left(\rho_{t,\tau}(\kappa_{t,\tau} - m_{t,\tau}) + \sqrt{(\kappa_{t,\tau} - m_{t,\tau})^2 + s_{t,\tau}^2}\right), \tag{23}$$

[12] In the original result of Breeden and Litzenberger (1978), the risk-neutral probability is defined with respect to the future spot price: $\tilde{\mathbb{P}}_{i,t}(S_{i,t+\tau}) = \left(1 + R_{f,t}^{(\tau)}\right) \cdot \left.\frac{\partial^2 C_{i,t}}{\partial K^2}\right|_{K=S_{i,t+\tau}}$. In our formulation, we define the risk-neutral probability in terms of future returns. With the change of variables, the probability becomes

$$\tilde{\mathbb{P}}_{i,t}(S_{i,t+\tau})dS_{i,t+\tau} = \tilde{\mathbb{P}}_{i,t}(S_{i,t+\tau})S_{i,t}dR_{i,t+\tau} = \left(1 + R_{f,t}^{(\tau)}\right) \cdot S_{i,t} \cdot \left.\frac{\partial^2 C_{i,t}}{\partial K^2}\right|_{K=S_{i,t+\tau}} dR_{i,t+\tau}.$$

Therefore, $\tilde{\mathbb{P}}_i(R_{i,t+\tau}) = \left(1 + R_{f,t}^{(\tau)}\right) \cdot S_{i,t} \cdot \left.\frac{\partial^2 C_{i,t}}{\partial K^2}\right|_{K=S_{i,t+\tau}}$.

[13] For notational simplicity, we omit the security subscript $i$.

[14] We measure $v_t$ as the median value of the implied volatility from option price observations on date $t$.

where $\Theta_t = (a_{t,\tau}, b_{t,\tau}, \rho_{t,\tau}, m_{t,\tau}, s_{t,\tau})$ is the set of model parameters that control the shape of the implied volatility surface, including asymmetry and slope across strike prices. Given that the volatility surface likely varies with maturity, we define each parameter as a linear function of $\tau$.[15]

For each day in the sample, we fit the SVI model to the implied volatility from the observed option prices, $\left\{\sigma_{t,j}^{\mathrm{BS}}\right\}_{j=1}^{N_t}$, where $N_t$ is the number of observations on date $t$. The parameter set is chosen to minimize the sum of squared errors:

$$\hat{\Theta}_t = \underset{\Theta_t}{\arg\min} \quad \sum_{j=1}^{N_t} \left[\sigma_{t,j}^{\mathrm{BS}} - \sigma_t^{\mathrm{SVI}}(\kappa_j, \tau_j; \Theta_t)\right]^2, \tag{24}$$

where $\sigma_t^{\mathrm{SVI}}(\kappa_j, \tau_j; \Theta_t)$ is the SVI prediction of volatility at moneyness $\kappa_j$ and time to maturity $\tau_j$. This optimization is conducted under the no-arbitrage constraint, i.e., the resulting option prices must satisfy $-\frac{1}{1+R_{f,t}^{(\tau)}} \leq \frac{\partial C_{i,t}}{\partial K} \leq 0$ and $\frac{\partial^2 C_{i,t}}{\partial K^2} \geq 0$ as in Aıt-Sahalia and Duarte (2003). Plugging the implied volatility into the Black-Scholes call option formula, we obtain the fitted prices of one-month call options across a detailed grid of strike prices and numerically compute the second derivative of the call prices as in equation (22).

### 3.3 Estimation of Physical Probability

We estimate the physical probability using historical spot prices of the underlying securities. Our objective here is to construct the date $t$-conditional *perceived* probability of 1-month-ahead returns, consistent with the horizon of the risk-neutral probability. To this end, we explicitly model how the recency bias affects the formation of the perceived distribution by applying the learning framework developed in Section 2.4.

In this estimation, we first determine the conditional volatility for each day, following Christoffersen et al. (2021). Next, we compute standardized innovations from daily returns and obtain their empirical distribution. Finally, using the estimated date-$t$ volatility and the empirical distribution, we simulate multiple paths of daily returns over a month and construct the conditional distribution of future returns implied by the investor's heuristic beliefs.

We determine the conditional volatility by utilizing high-frequency intraday spot prices. Tick-by-tick trades are aggregated into one-minute intervals and averaged for each minute. We then apply the estimation approach by Zhang et al. (2005) to the intraday prices and

[15] For example,

$$a_{t,\tau} = a_{t,0} + a_{t,1}\tau,$$

and similarly for other parameters. For each day, we estimate the set of parameters, $\Theta_t = (a_{t,0}, a_{t,1}, b_{t,0}, b_{t,1}, \rho_{t,0}, \rho_{t,1}, m_{t,0}, m_{t,1}, s_{t,0}, s_{t,1})$, by minimizing the sum of squared errors.

compute daily realized spot volatilities, denoted by $\sigma_{i,t}$. Given the spot volatilities, we model their dynamics using the heterogeneous autoregressive (HAR) model suggested by Corsi (2009). This framework captures key empirical features of financial returns in a tractable way. Specifically, we formulate the next-day volatility as a linear function of current volatilities with different frequencies:

$$\sigma_{i,t+1}^{\text{HAR}} = a_i + b_{i,d}\sigma_{i,t} + b_{i,w}\sigma_{i,t-6:t} + b_{i,m}\sigma_{i,t-29:t} + e_{i,t+1} \tag{25}$$

where $\sigma_{i,t}$ is the realized spot volatility, $\sigma_{i,t-6:t}$ is the average daily volatility in the past week and $\sigma_{i,t-29:t}$ is the average in the past month. The coefficients are estimated via rolling-window regressions on the past 365 days of data. Consequently, this estimation avoids any look-ahead bias.

We normalize daily returns and obtain standardized innovations as

$$\epsilon_{i,t} = \frac{R_{i,t} - \mu_{i,t}(\nu_i)}{\sigma_{i,t-1}} \tag{26}$$

where $\mu_{i,t}$ is the perceived average that is constantly updated according to equation (19) under the recency-bias parameter $\nu_i$. Consistent with the return dynamics specified in (17), we assume that the standardized innovations $\epsilon_{i,t}$ are independently and identically distributed and drawn from a time-invariant distribution $F_i(\epsilon_{i,t}; \nu_i)$. We obtain its empirical distribution by pooling observations across the entire sample period.

Having determined the distribution $F_i(\epsilon_i; \nu_i)$, we simulate future returns over one month beginning on date $t$. For example, the next-day return is given by $\tilde{R}_{i,t+1} = \mu_{i,t}(\nu_i) + \hat{\sigma}_{i,t}^{\text{HAR}}\tilde{\epsilon}_{i,t+1}$, where $\hat{\sigma}_{i,t}^{\text{HAR}}$ is the volatility predicted by the HAR in equation (25) and $\tilde{\epsilon}_{i,t+1}$ is a random draw from $F_i(\epsilon_i; \nu_i)$. In this way, we simulate 1,000 paths of 30-day returns. Applying a kernel smoothing function to these generated monthly returns as in Jackwerth (2000), we estimate the perceived physical probability $\mathbb{P}_{i,t}\left(R_{t+30}; \nu_i\right)$ given the recency-bias parameter $\nu_i$.

### 3.4 Summary Statistics

Table 1 presents summary statistics. Several points are worth noting. Bitcoin's implied volatility is 73.8% on average, noticeably higher than the S&P 500 index's average of 20.1%. Next, the estimates of the security-specific risk aversion indicate that investors are noticeably less averse to future losses in Bitcoin than those in the stock market. Under the physical probabilities implied by Bayesian learning ($\nu_i = 0$), the Bitcoin-specific risk aversion evaluated

at –5% future return is 5.55 on average, which is considerably lower than 35.89 for the S&P 500 index. This observation aligns with the intuition that investors are relatively less concerned about Bitcoin's poor performance, as it typically constitutes only a small fraction of investment portfolios. Meanwhile, the risk aversion for the market index evaluated at 5% return is occasionally negative. These negative values reflect the non-monotonicity of the empirical pricing kernel as documented by Jackwerth (2000) and Cuesdeanu and Jackwerth (2018).

# 4 Empirical Evidence

In this section, we present empirical evidence that supports the model's main prediction with respect to security-specific risk aversion. We then estimate the model to identify the belief-formation rule that best fits the data, focusing on the extent to which recency bias shapes the perceived return distribution.

## 4.1 Perception of Expected Returns

We now describe the estimates of perceived expected returns $\mu_{i,t}$ that investors form from the history of realized returns. As outlined in section 2.4, we model this perception as evolving continuously as new observations arrive, with the parameter $\nu_i$ governing the weighting scheme applied to past data. When $\nu_i = 0$, belief updating follows the Bayesian benchmark and assigns equal weights to all observations. In contrast, a positive value of $\nu_i$ reflects recency bias, in which more recent returns exert greater influence on current beliefs.

As an illustration, Figure 3 depicts the weighting schemes that different learning frameworks apply to the past 250 observations. When the belief-formation parameter $\nu_i$ is positive, the weights assigned to past returns decline exponentially with temporal distance from the present. A higher value of $\nu_i$ corresponds to a steeper decay in weights, indicating a stronger degree of recency bias.

We consider alternative values of $\nu_i = \{0, 0.01, 0.02\}$ and examine how the choice of $\nu_i(\nu_i)$ affects the dynamics of perceived average return. Figure 4 illustrates the trajectories of beliefs about Bitcoin's average return for various levels of recency bias. All of these paths of belief begin with the same diffuse priors that have a mean of zero ($\mu_{B,0} = 0$ and $\varphi_{B,0} = 10^{10}$) and are subsequently updated based on the same realized daily returns on Bitcoin, presented in panel A.

Although derived from the same sequence of realized returns, the belief trajectories in Panel B differ markedly depending on the degree of recency bias. When $\nu_B = 0$, corresponding to Bayesian updating, the perceived average is highly persistent, as illustrated by the dashed line. In particular, toward the latter part of the sample, the Bayesian belief remains largely unchanged. This persistence arises because equal weighting of past and present observations causes the influence of new data to diminish as the sample grows. Once a sufficiently large number of observations has accumulated, learning is effectively complete, and additional data induce only negligible revisions to the perceived average.

In contrast, when investors are subject to recency bias ($\nu_B > 0$), and recent observations carry greater weight in belief formation, the perceived drift fluctuates more markedly along with the latest return. In this heuristic learning, the influence of distant-past observations fades away, making belief updates constantly sensitive to new information – even when past observations are plentiful. Consequently, the belief update is never complete, unlike in Bayesian learning. A stronger recency bias, indicated by a higher value of $\nu_B$, amplifies this sensitivity and produces more pronounced swings in the perceived average, as illustrated by the solid line.

### 4.2 Estimates of Security-Specific Risk Aversion

Next, we present our estimates of security-specific risk aversion. Figure 5 displays the daily estimates of Bitcoin-specific risk aversion under different degrees of recency bias. At this point, we are agnostic about the magnitude of the bias and instead illustrate how the choice of $\nu_B$ affects the dynamics of risk aversion. Specifically, we consider three benchmark cases in which $\nu_B = \{0, 0.01, 0.02\}$ – corresponding to increasing levels of recency bias.

Panels A and B plot risk aversion $\mathcal{R}_{B,t}(R_{B,t+30}; \nu_B)$ evaluated at two future outcomes, a 5% return and a –5% return over the next month, when $\nu_B$ equals zero. This case corresponds to the absence of a heuristic bias in the perceived physical distribution, as implied by Bayesian learning. We notice that Bitcoin risk aversion exhibits substantial time variation, underscoring its conditional nature. Focusing on the estimate at –5% return, its standard deviation is 16.31%, almost three times as large as its average of 5.55%, as reported in Table 1.

Interestingly, the choice of the recency bias parameter has a noticeable impact on the risk aversion estimates. Panels C and D show that introducing a recency bias with $\nu_B = 0.01$ leads to substantially greater volatility in Bitcoin risk aversion. This observation aligns with

the learning framework, in which the agent's heuristic belief updating causes the perceived distribution to respond more strongly to the latest realization of security returns. According to equation (21), larger shifts in the perceived distribution translate into wider fluctuations in risk aversion, while holding the risk-neutral distribution fixed. This sensitivity is even more pronounced under a stronger recency bias, with $\nu_B = 0.02$ in Panels E and F.

For comparison, Figure 6 plots equity-specific risk aversion, estimated from options on the S&P 500 index, under alternative learning frameworks. Similar to Bitcoin, equity-specific risk aversion exhibits substantial time variation. When $\nu_S = 0$, the risk-aversion estimate evaluated at a –5% return has a standard deviation of 37.13% and a mean of 35.89%. Meanwhile, equity and Bitcoin risk aversion evolve largely independently. The correlation between the two series at the –5% return is only 0.11. This weak association suggests that investors' risk attitudes toward these two securities are distinct, consistent with the evidence in Liu and Tsyvinski (2020), who find that the risk factors driving cryptocurrency returns differ from those in the equity market.

Across Panels A through F, we find that the choice of belief formation rules – Bayesian versus recency-biased learning – also has a noticeable effect on the magnitude of equity-specific risk aversion. This dependence underscores the importance of identifying the underlying belief formation, a task we undertake in later sections.

### 4.3 Determinats of Security-Specific Risk Aversion

We now test the model's prediction by examining the empirical response of security-specific risk aversion to perceived expected returns. As shown in the previous sections, the estimation of risk aversion and perceived conditional returns depends on the learning parameter $\nu_i$, which governs the degree of recency bias. In this section, we use preset values of $\nu_i$ and examine whether the model predictions hold under different assumptions about belief updating. In the following section, we identify the recency bias parameter that best fits the data.

For each value of $\nu_i$, we construct the dynamics of risk aversion $\mathcal{R}_{i,t}(R_{i,t+30}; \nu_i)$ and the perceived conditional mean $\mu_{i,t}(\nu_i)$ for security $i \in \{B, S\}$. Using these estimates, we run time series regressions of the form

$$\mathcal{R}_{i,t}(R_{i,t+30}; \nu_i) = \alpha + \beta \times \mu_{i,t}(\nu_i) + \gamma \times \text{Controls}_t + \epsilon_{i,t} \tag{27}$$

where the dependent variable is the daily estimate of security-specific risk aversion evaluated at future returns of 5% or –5%. As described in Proposition 4, we expect risk aversion to

respond positively to the perceived conditional mean return – when investors perceive a higher return on a security, they allocate more wealth to the asset and therefore become more averse to its potential losses in subsequent periods.

Controls include the conditional volatility $\sigma_{i,t}$, described in Section 3.3, and macroeconomic indicators, including growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), the Chicago Fed National Financial Condition Index (NFCI), and the inflation rate ($\Delta$CPI). These macroeconomic variables are available at the monthly frequency. We therefore assign each monthly observation to all trading days within the corresponding month, following standard practice in the literature. In regressions of Bitcoin-specific risk aversion, we also include the cryptocurrency's characteristics – network effect ($\text{Network}_B$) and attention ($\text{Attention}_B$) indicators – which the recent study by Liu and Tsyvinski (2020) finds predict its return and potentially reflect the risk profile. The standard errors for the coefficient estimates are adjusted for serial correlation using the Newey-West method.

### 4.3.1 Bitcoin-Specific Risk Aversion

Table 2 presents the regression results for Bitcoin-specific risk aversion. Panel A shows that the perceived conditional average return $\mu_{B,t}$ is a strong and positive predictor of risk aversion evaluated at a future return of –5%, regardless of the degree of recency bias. This robust positive link indicates that a higher average return perceived by investors from history leads them to become more averse to Bitcoin's downside risk in the future, consistent with the model's prediction.

We next examine the alternative measure of risk aversion evaluated at a future return of 5% in Table 2 Panel B. The perceived expected return continues to positively predict risk aversion at the 1% significance level. However, the goodness of fit varies markedly across learning frameworks. In regressions onto the set of conditional moments $\mu_{B,t}$ and $\sigma_{B,t}$, the adjusted $R^2$ ranges from 13.3% to 72.8% across the considered levels of recency bias. In the following section, we will use these differences in model fit to identify the parameter $\nu_B$ that underlies belief formation.

In addition to the average return, conditional volatility also shapes risk aversion. Proposition 5 states that a rise in Bitcoin volatility should reduce investors' risk aversion going forward. While this prediction may appear counterintuitive at first, it arises naturally from portfolio choice: in times of higher Bitcoin volatility, investors optimally reduce their exposure to the asset and are therefore less sensitive to its future downside risk. The regression

results in Table 2 support this prediction. Across all specifications, risk aversion responds negatively to conditional volatility. Relatedly, Bliss and Panigirtzoglou (2004) and Schreindorfer and Sichert (2025) document that risk aversion implicit in stock options is elevated during periods of low volatility. Our findings reveal that a similar regularity holds in the Bitcoin market. Moreover, our model provides a unified explanation for this stylized fact: the dependence of risk aversion on volatility emerges endogenously through investors' portfolio allocation decisions.

It is noteworthy that the combination of perceived conditional average and volatility accounts for a significant share of fluctuations in risk aversion, especially when recency bias is present with $\nu_B > 0$. In regressions using only the combination, the adjusted $R^2$ values range from 72.3% to 76.1%. Meanwhile, adding additional controls only marginally improves the goodness of fit. We therefore establish the perceived return distribution as a primary driver of investors' portfolio allocation and risk attitude, consistent with the model.

#### 4.3.2 Equity-Specific Risk Aversion

Having confirmed the model prediction in the Bitcoin market, we turn to the S&P 500 index and examine equity-specific risk aversion. Given that our model's predictions are based on portfolio choices, they should be applicable to any security. To test this idea, we estimate the perceived conditional average $\mu_{S,t}(\nu_S)$ using the index's daily returns and evaluate its predictive ability for equity-specific risk aversion in the specification (27).

Table 3 presents the results for the risk aversion dynamics in the equity market. Overall, the patterns closely mirror those observed in the Bitcoin market. The perceived average return $\mu_{S,t}$ formed under recency-biased learning is a significant and positive predictor of risk aversion, while the conditional volatility $\sigma_{S,t}$ is negatively related. These findings suggest that investors become more sensitive to equity's potential downside risk when they perceive stronger past performance and less uncertainty in historical returns.

The choice of learning framework is likewise central to explaining the dynamics of equity-specific risk aversion. Under Bayesian updating ($\nu_S = 0$), the perceived average return is negatively associated with the risk aversion evaluated at a future return of –5%, which is the opposite of the model prediction. In contrast, introducing a recency bias ($\nu_S > 0$) restores the positive relationship consistent with the model. Notably, this sensitivity to belief formation suggests that even small departures from Bayesian learning can generate markedly different empirical implications. This result is consistent with the findings of Malmendier

and Nagel (2011), who show – using consumer finance survey data – that ignoring heuristic biases in stock return expectations leads to counterfactual predictions about investors' market participation decisions.

### 4.4 Determinats of Expected Risk Aversion

In the analysis so far, security-specific risk aversion has been defined as the local derivative of projected marginal utility. While this measure is directly motivated by the theory, it focuses only on specific future return realizations, such as ±5%. To capture the agent's overall risk attitude over a broader range of future returns, we consider an alternative measure. Specifically, we define expected risk aversion as

$$\overline{\mathcal{R}}_{i,t}(\nu_i) = \int_{-10\%}^{10\%} \mathcal{R}_i(R_{i,t+30};\nu_i)\mathbb{P}_{i,t}(R_{i,t+30};\nu_i)dR_{i,t+30}, \tag{28}$$

where the local curvature of marginal utility is integrated over the return support $[-10\%, 10\%]$ under the perceived physical distribution $\mathbb{P}_{i,t}(R_{i,t+30};\nu_i)$. This formulation aggregates local information about risk aversion in a manner consistent with the agent's perceived probabilities. Figure 7 plots the time series of expected risk aversion in the Bitcoin and equity markets.

Under Bayesian updating, expected risk aversion for Bitcoin averages 1.43, as shown in Table 1. As the aggregate measure of risk preferences across return states, its magnitude falls between those of the local risk-aversion measures evaluated at –5% and 5% returns. Additionally, the expected risk aversion measure exhibits more moderate time-series fluctuations than the local measures; its standard deviation is 4.02, approximately one-quarter of the standard deviation of risk aversion at a –5% return. A similar pattern holds in the equity market, where the expected risk aversion measure lies between the local measures and is less volatile over time.[16]

Using expected risk aversion $\overline{\mathcal{R}}_{i,t}(\nu_i)$ and the perceived expected return $\mu_{i,t}(\nu_i)$ for a given value of $\nu_i$, we run the time-series regression of the form

$$\overline{\mathcal{R}}_{i,t}(\nu_i) = \alpha + \beta \times \mu_{i,t}(\nu_i) + \gamma \times \text{Controls}_t + \epsilon_{i,t}, \tag{29}$$

[16] Under Bayesian updating, expected risk aversion for equity averages 0.49, below the corresponding Bitcoin estimate of 1.43. This comparison is difficult to reconcile with the premise that investors typically allocate substantially more wealth to equities and should therefore be more averse to equity downside risk than to adverse returns in Bitcoin. We interpret this counterintuitive result as additional evidence supporting recency-biased learning. When expected risk aversion is instead computed using the estimated recency bias parameters in Section 4.5, $\nu_B = 0.021$ and $\nu_S = 0.008$, equity's expected risk aversion becomes 6.62 on average, whereas the corresponding Bitcoin estimate is 1.49.

where controls are the same as in the specification (27). This analysis allows us to examine how the overall risk attitude evolves in relation to the perceived return distribution.

Table 4 presents the regression results. Our baseline findings continue to hold for the dynamics of expected risk aversion. In both Bitcoin and equity markets, a rise in the perceived average implied by historical returns is associated with greater expected risk aversion. At the same time, fluctuations in conditional volatility remain a negative predictor of expected risk aversion. Furthermore, the higher adjusted $R^2$ values with positive $\nu_i$ indicate that incorporating recency bias into the estimation yields a superior fit to the data compared to Bayesian learning.

To summarize, our analysis confirms the model's prediction that investors' updates of their perceived return distributions cause security-specific risk aversion to vary over time in both the Bitcoin and equity markets. This pattern persists whether risk aversion is measured using the local or average curvature of marginal utility. The positive association is most pronounced when beliefs are formed under recency bias, whereas Bayesian updating yields weaker and less consistent results across alternative measures of risk aversion. This contrast highlights the importance of identifying the bias parameter $\nu_i$ in better understanding the belief-formation rule that shapes investors' empirical risk attitudes – a task we undertake in the next section.

### 4.5 Estimation of Recency Bias in Belief Formation

The objective of this section is to identify the parameter $\nu_i$ that governs recency bias in belief formation. Our preceding analysis shows that different values of $\nu_i$ – particularly those distinguishing recency-biased learning from Bayesian updating – yield distinct paths for the perceived conditional mean return and security-specific risk aversion. Given this sensitivity, the degree of recency bias cannot be imposed a priori and must instead be identified from the data.

To identify the belief-formation parameter, we use estimates of expected risk aversion. We focus on this measure rather than the local estimates of risk aversion because we aim to capture the dynamics of overall risk attitudes rather than the behavior at a particular return outcome. Equation (29) provides the model prediction for expected risk aversion, and we examine its goodness of fit across different choices of $\nu_i$.[17]

[17] Our estimation approach is motivated by Malmendier and Nagel (2011), who estimate a learning parameter from stock market participation survey data.

For each asset, we estimate the regression model (29), including the recency bias parameter $\nu_i$. Specifically, we search for the set of parameters, $\Theta_i = (\nu_i, \alpha, \beta, \gamma)$ that minimize the sum of normalized squared residuals:

$$\hat{\Theta}_i = \arg\min_{\Theta_i} \sum_{t=1}^{T} \varepsilon_{i,t}(\Theta_i)^2 = \arg\min_{\Theta_i} \sum_{t=1}^{T} \left[ \frac{\overline{\mathcal{R}}_{i,t}(\nu_i) - \widehat{\mathcal{R}}_{i,t}(\Theta_i)}{\sqrt{\mathrm{SST}(\nu_i)}} \right]^2 \tag{30}$$

where $\widehat{\mathcal{R}}_{i,t}(\Theta_i)$ are the fitted values of equation (29) and SST denotes the total sum of squares, defined as

$$\mathrm{SST}(\nu_i) = \sum_{t=1}^{T} \left( \overline{\mathcal{R}}_{i,t}(\nu_i) - \frac{1}{T} \sum_{t=1}^{T} \overline{\mathcal{R}}_{i,t}(\nu_i) \right)^2 . \tag{31}$$

This normalization reflects the fact that the dependent variable $\overline{\mathcal{R}}_{i,t}(\nu_i)$ itself varies with the choice of $\nu_i$, and therefore, non-normalized residuals do not permit a meaningful comparison of goodness of fit across specifications. Essentially, minimizing the sum of normalized squared residuals is equivalent to maximizing the regression $R^2$ across alternative values of $\nu_i$.[18] Details of this estimation are provided in Appendix C.

Figure 8 illustrates our identification strategy by plotting the sum of squared normalized residuals from regression (29) for alternative values of the recency-bias parameter $\nu_i$ over the range [0,0.05]. In both the Bitcoin and equity markets, the objective function exhibits a pronounced U-shaped pattern with a well-defined interior minimum, indicating that the data contain sufficient information to identify $\nu_i$. Moreover, the minimum occurs at a substantially higher value of $\nu_i$ in the Bitcoin market than in the equity market, providing preliminary evidence that investors place greater weight on recent observations when forming beliefs about Bitcoin returns. We quantify this difference formally in the estimation results below.

We emphasize that no constraints on $\nu_i$ are imposed in the estimation. Instead, we let data reveal which learning framework best explains the dynamics of risk aversion. The estimate of $\nu_i$ being not significantly different from zero would point to Bayesian learning. Alternatively, a positive (negative) estimate of $\nu_i$ would indicate recency (primacy) bias in learning, where more recent (more distant) observations in the security history have greater impact on current beliefs.

Table 5 reports the estimation results. In the Bitcoin market, the null hypothesis of no

[18] Formally,

$$R^2\text{statistics} = 1 - \frac{\sum_{t=1}^{T} \left[ \overline{\mathcal{R}}_{i,t}(\nu_i) - \widehat{\mathcal{R}}_{i,t}(\Theta_i) \right]^2}{\mathrm{SST}(\nu_i)} = 1 - \sum_{t=1}^{T} \left[ \frac{\overline{\mathcal{R}}_{i,t}(\nu_i) - \widehat{\mathcal{R}}_{i,t}(\Theta_i)}{\sqrt{\mathrm{SST}(\nu_i)}} \right]^2 = 1 - \sum_{t=1}^{T} \varepsilon_{i,t}(\Theta_i)^2$$

recency bias ($\nu_B = 0$) is rejected. Instead, the parameter estimates are positive and highly significant, with a $t$-statistic of 26.57, confirming the presence of recency bias in the perceived distribution of Bitcoin returns. Quantitatively, the estimate of 0.021 implies that in belief formation the weights on past observations decay exponentially, such that a return observed 108 days ago receives only one-tenth the weight of today's return.[19] This rapid decay in the influence of past observations stands in sharp contrast to the equal weighting implied by Bayesian learning.

Simultaneously, we confirm that Bitcoin-specific risk aversion is positively associated with the perceived average return $\mu_{B,t}$ formed under recency-biased learning. This relationship is economically significant. The coefficient estimate of 3.810 implies that a one-standard-deviation increase in the perceived average raises risk aversion by 62.5% of its own standard deviation, representing the strongest effect among the determinants considered. The second most influential factor is volatility. Based on the coefficients on $\sigma_{B,t}$, a one-standard-deviation increase in volatility lowers risk aversion by 16.2% of its standard deviation. These jointly significant relationships establish that investors' continuous updates of perceived return distributions – and the resulting portfolio reallocations – constitute a fundamental force driving the dynamics of risk aversion.

Next, we examine the results for the equity market index. The estimate of recency bias $\nu_S$ implicit in the equity market is positive, and statistical significance is even stronger with a $t$-statistic of 54.83. We therefore confirm that investors' perception of stock market returns is also strongly influenced by recency bias.

Importantly, the point estimate of $\nu_S$ is approximately half the magnitude of its Bitcoin counterpart. According to the estimated recency bias parameter of 0.008, investors place one-tenth the weight on a stock market return observed 287 days earlier relative to the most recent return. This rate of decay is markedly lower than in the Bitcoin market, where the same reduction in weight occurs after 108 days. This comparison suggests that participants in the Bitcoin market process information in a more short-sighted manner: their perceptions of the return distribution are substantially more influenced by recent data than those formed in the equity market.

[19] According to equation (18), the weight on returns observed $k$ days earlier than the current date is proportional to $(1-\nu_i)^k$. Thus, the weight decay from today to $k$ days earlier is $\frac{(1-\nu_i)^k}{(1-\nu_i)^0}$. Given the estimate of $\nu_B$, we solve for $k$ such that

$$\frac{(1-\nu_B)^k}{(1-\nu_B)^0} = 0.1.$$

Solving for the unknown, we find $k = 108.49$.

Interestingly, our estimate of the weight decay in the equity market is quantitatively similar to the evidence in the literature. Greenwood and Shleifer (2014) find that survey-based expected returns correlate strongly with an exponentially weighted moving average of past returns. In particular, their estimated weights suggest that returns from four quarters earlier are 9.8% as influential as returns from the most recent quarter in shaping current expectations.[20] This close alignment lends external validity to our option-based analysis and reinforces the ensuing conclusion that the magnitudes of recency bias differ across markets.

Why does the Bitcoin market exhibit a stronger recency bias than the equity market? One interpretation is that recency-biased learning serves as a heuristic approximation to an optimally chosen estimation window. Econometric studies, such as Pesaran and Timmermann (2007), emphasize that the estimation window should be adjusted to minimize forecast errors when structural breaks in the data-generating process are anticipated. In environments characterized by larger or more frequent structural shifts, shorter estimation windows are preferred to reduce bias arising from outdated information. In this regard, Bitcoin's regulatory outlook and long-run adoption path remain highly ambiguous, leading investors to anticipate more frequent fundamental changes. As a result, investors may heuristically anchor their expectations to recent observations as a way to capture rapidly changing fundamentals.

### 4.6 Robustness Test: Alternative Measures of Risk Aversion

In this section, we examine the robustness of our findings using alternative measures of risk aversion. First, we construct security-specific risk aversion based on the rational physical distribution, which is not subject to recency bias. Second, we employ weekly, rather than daily, estimates of our baseline risk aversion measures.

#### 4.6.1 Risk-Aversion Estimation without Recency Bias

In our baseline estimation, recency bias shapes both the physical probability $\mathbb{P}_{i,t}\left(R_{i,t+30};\nu_i\right)$ – and hence the option-implied risk aversion – and the perceived average return $\mu_{i,t}(\nu_i)$. While this formulation ensures internal consistency, the estimated relationship between risk

[20] Greenwood and Shleifer (2014) conduct time-series regressions of survey expectation ($\text{Exp}_t$) of stock market returns for the next year on the exponentially weighted average of past market returns

$$\text{Exp}_t = \alpha + \beta \sum_{j=0}^{k} w_j R_{t-j} + \epsilon_t,$$

where the weight on return from $j$ quarters earlier is $w_j = \frac{\lambda^j}{\sum_{i=0}^{k} \lambda^i}$. They estimate $\lambda = 0.56$ on average. The resulting decay in weight over four quarters is $\frac{w_4}{w_0} = \frac{\lambda^4}{\lambda^0} \approx 0.098$.

aversion and perceived average might partly reflect a mechanical relation; both the dependent and independent variables in equation (29) depend on the same parameter $\nu_i$. To address this concern and isolate the economic force of portfolio choice as the driver of the relationship, we consider an alternative measure of risk aversion that is constructed independently of $\nu_i$.

Specifically, we define the rational physical probability $\mathbb{P}_{i,t}\left(R_{i,t+30}; \nu_i = 0\right)$, under which all past observations up to date $t$ receive equal weight as in Bayesian learning. The resulting expected risk aversion is

$$\overline{\mathcal{R}}_{i,t}(0) = \int_{-10\%}^{10\%} \mathcal{R}_{i,t}(R_{i,t+30}; \nu_i = 0)\mathbb{P}_{i,t}(R_{i,t+30}; \nu_i = 0)dR_{i,t+30}. \tag{32}$$

Using this alternative measure of risk aversion as the dependent variable, we re-estimate equation (29) including the recency bias parameter. Importantly, the parameter $\nu_i$ enters only through the construction of the perceived expected return $\mu_{i,t}(\nu_i)$ and does not affect $\overline{\mathcal{R}}_{i,t}(0)$. As a result, this specification eliminates the mechanical relation between the dependent and independent variables.

Table 6 reports the results. The estimates of $\nu_i$ remain positive and statistically significant in both the Bitcoin and equity markets. These consistently positive estimates confirm that our main findings are not spuriously driven by statistical artifacts, and that recency bias indeed characterizes the underlying belief formation. Furthermore, we continue to observe differential degrees of heuristic bias across markets; $\nu_B$ is again more than twice as large as $\nu_S$. This observation echoes our baseline results and further supports the conclusion that recency bias is considerably more pronounced in the Bitcoin market than in the equity market.

Meanwhile, this alternative parameter estimate differs in magnitude from our baseline. For example, in the equity market, the alternative estimate $\nu_s = 0.015$ indicates that the weights on past observations decline to one-tenth after 152 days, compared with 287 days under baseline specification. We interpret that this discrepancy reflects a biased characterization of belief formation arising from model misspecification. In the correct specification, where both risk aversion and perceived average are constructed under the same heuristic beliefs, the resulting weighting scheme closely aligns with the evidence in Greenwood and Shleifer (2014). On the other hand, the alternative specification helps us establish the robust role of heuristics in the agent's perception, but it overstates the degree of recency bias due to misspecification.

### 4.6.2 Alternative Data Frequency and Sample Period

In our main analysis, we examine the daily estimates of risk aversion and apply Newey-West adjustments in the regressions to address potential bias arising from serial correlation. As an alternative approach, we re-estimate the risk aversion at the weekly frequency and estimate the recency bias parameter $\nu_i$ that best explains its weekly variation. The results, presented in Appendix D, show that $\nu_B$ is 0.020 for Bitcoin, and $\nu_S$ is 0.007 for the equity market. These estimates are quantitatively close to those obtained from daily data, and therefore testify to the robustness of our finding regarding the magnitude of recency bias in both markets.

In addition, we consider an alternative sample period for the stock-market estimation. Our baseline estimates use the full sample of S&P 500 index options from January 1997 to June 2024, which is substantially longer than the Bitcoin option sample beginning in January 2017. Given this disparity in data coverage, one might argue that our finding of stronger recency bias in the Bitcoin market simply reflects investors' general behavior over more recent years rather than their security-specific attitudes. To address this concern, we sub-sample the data on stock index options to match the availability of Bitcoin option data and re-estimate the model. The result, reported in Appendix D, is that $\nu_S$ from this matched sample is 0.009, very close to the baseline estimate obtained from the full sample.

## 5 Conclusion

This paper examines how investors form beliefs about the return distribution and adjust their risk attitudes when facing a new financial asset, such as Bitcoin. Using a theoretical framework and option-implied measures of absolute risk aversion, we show that investors' risk aversion toward securities evolves endogenously, as beliefs formed from growing observations lead them to adjust portfolio exposures. Empirically, we find that option-implied risk aversion toward Bitcoin is lower on average than equity risk aversion, and that investors exhibit a pronounced recency bias, placing excessive weight on recent information when forming return expectations. This bias is stronger in the Bitcoin market, consistent with investors' reliance on heuristics when historical data are limited and structural changes are anticipated. Overall, our findings emphasize the central role of belief formation and learning dynamics in shaping demand, risk attitudes, and asset prices in emerging financial markets.

## References


Adam, Klaus, Albert Marcet, and Johannes Beutel, 2017, Stock price booms and expected capital gains, *American Economic Review* 107, 2352–2408.

Ait-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Ait-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Alexander, Carol, Jung Deng, Jianfen Feng, and Huning Wan, 2024, Net buying pressue and the information in bitcoin option trades, *Working Paper* .

Alexander, Carol, and Arben Imeraj, 2021, The bitcin vix and its variance risk premium, *Journal of Alternative Investments* 23, 84–109.

Atanasova, Christina, Terrel Miao, Ignacio Segarra, Tony (Tong) Sha, and Frederick Willeboordse, 2024, Illiquidity premium and crypto option returns, *Working Paper* .

Aıt-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9–47, Frontiers of financial econometrics and financial engineering.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-capm: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.

Beason, Tyler, and David Schreindorfer, 2022, Dissecting the equity premium, *Journal of Political Economy* 130, 2203–2222.

Berger, David, Ian Dew-Becker, and Stefano Giglio, 2019, Uncertainty shocks as second-moment news shocks, *The Review of Economic Studies* 87, 40–76.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal of Finance* 59, 407–446.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *The Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *The Journal of Finance* 73, 199–227.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: a consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Cassella, Stefano, and Juergen Huber, 2020, Extrapolation bias and the predictability of stock returns by price-scaled variables, *Review of Financial Studies* 33, 4345–4397.

Chen, Tian, Jun Deng, and Jing Nie, 2024, Implied volatility slopes and jumps in bitcoin options market, *Operations Research Letters* 55, 107135.

Christoffersen, Peter, Kris Jacobs, and Xuhui (Nick) Pan, 2021, The state price density implied by crude oil futures and option prices, *The Review of Financial Studies* 35, 1064–1103.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664–698.

Cong, Lin William, Zhiguo He, and Jiasun Li, 2021, Decentralized mining in centralized pools, *Review of Financial Studies* 34, 1191–1235.

Corsi, Fulvio, 2009, A simple approximate long-memory model or realized volatility, *Journal of Financial Econometrics* 7, 174–196.

Cuesdeanu, Haratio, and Jens Carsten Jackwerth, 2018, The pricing kernel puzzle: Survey and outlook, *Annals of Finance* 14, 289–329.

Da, Zhi, Xing Huang, and Lawrence J. Jin, 2021, Extrapolative beliefs in the cross-section: What can we learn from the crowds?, *Journal of Financial Economics* 140, 175–196.

Gatheral, Jim, and Antoine Jacquier, 2011, Convergence of heston to svi, *Quantitative Finance* 11, 1129–1132.

Gennaioli, Nicola, and Andrei Shleifer, 2010, What comes to mind*, *The Quarterly Journal of Economics* 125, 1399–1433.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.

Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *The Review of Financial Studies* 13, 433–451.

Liu, Yukun, and Aleh Tsyvinski, 2020, Risks and returns of cryptocurrency, *The Review of Financial Studies* 34, 2689–2727.

Makarov, Igor, and Antoinette Schoar, 2020, Trading and arbitrage in cryptocurrency markets, *Journal of Financial Economics* 135, 293–319.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?, *The Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Stefan Nagel, 2016, Learning from inflation experiences, *The Quarterly Journal of Economics* 131, 53–87.

Malmendier, Ulrike, and Jessica A. Wachter, 2024, Memory of past experiences and economic decisions, in *The Oxford Handbook of Human Memory, Two Volume Pack: Foundations and Applications* (Oxford University Press).

Martin, Ian, 2017, What is the expected return on the market?, *Quarterly Journal of Economics* 132, 367–433.

Martin, Ian, and Christian Wagner, 2019, What is the expected return on a stock?, *Journal of Finance* 74, 1887–1929.

Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *The Review of Financial Studies* 35, 2190–2245.

Pesaran, M. Hashem, and Allan Timmermann, 2007, Selection of estimation window in the presence of breaks, *Journal of Econometrics* 137, 134–161.

Rosenberg, Joshua V., and Robert F. Engle, 2002, Emprical pricing kernels, *Journal of Financial Economics* 64, 341–372.

Schreindorfer, David, and Tobias Sichert, 2025, Conditional risk and the pricing kernel, *Journal of Financial Economics* 171, 104106.

Sockin, Michael, and Wei Xiong, 2023, A model of cryptocurrencies, *Management Science* 69, 6735–6754.

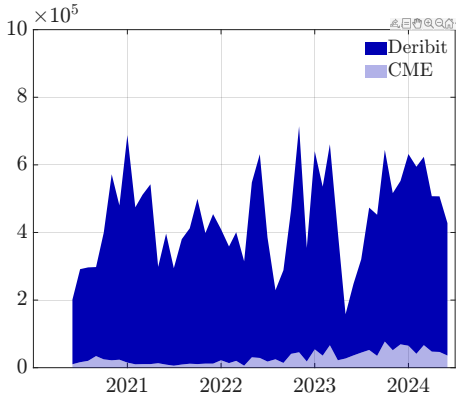
Tversky, Amos, and Daniel Kahneman, 1974, Judgment under uncertainty: Heuristics and biases, *Science* 185, 1124–1131.

Winkel, Julian, and Wolfgang Karl Härdle, 2023, Pricing kernels and risk premia implied in bitcoin options, *Risks* 11.

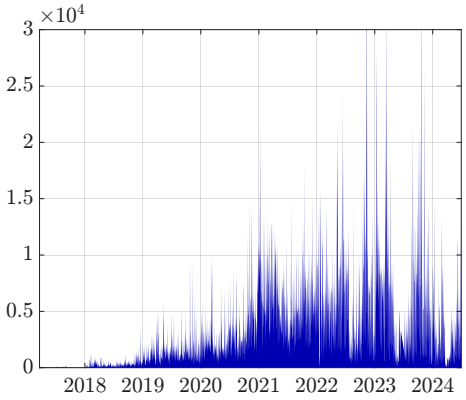
Zhang, Lan, Per A Mykland, and Yacine Aït-Sahalia, 2005, A tale of two time scales: Determining integrated volatility with noisy high frequency data, *Journal of the American Statistical Association* 100, 1394–1411.

**Figure 2: Bitcoin Options Trading Volume on Deribit**

Panel A: Monthly Trading Volume

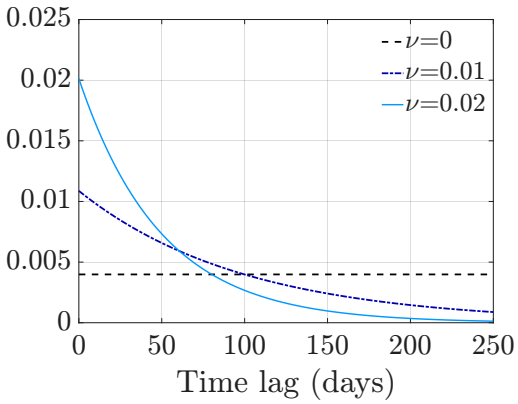



Panel B: Daily Trading Volume



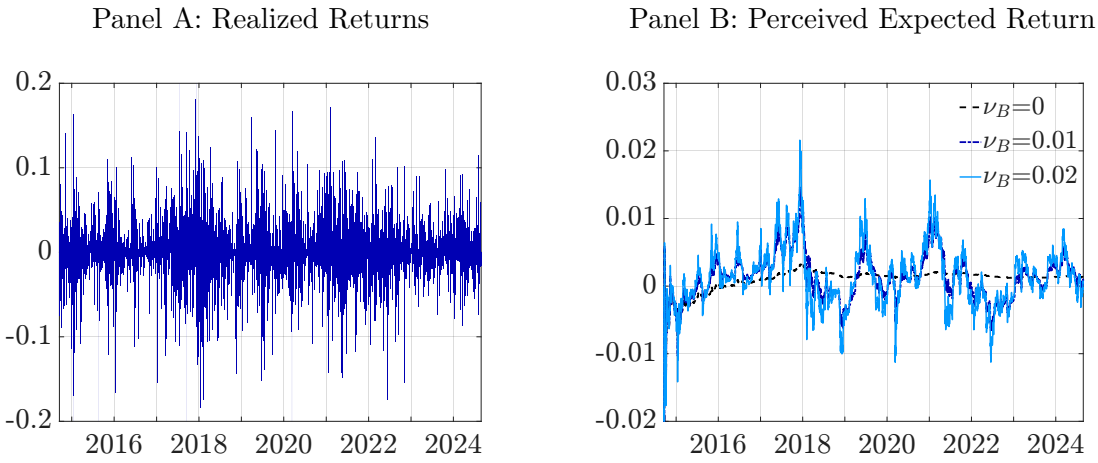

This figure summarizes the trading activity in Bitcoin options. Panel A compares monthly Bitcoin option trading volumes on Deribit and CME. Panel B displays the daily contract volume traded on Deribit. Historical CME data are obtained from The Block.

**Figure 3: Weighting Functions under Bayesian and Recency-Biased Beliefs**




This figure shows the weighting schemes implied by Bayesian ($\nu = 0$) and recency-biased learning ($\nu > 0$) when the perceived average is constructed from the past 250 observations.

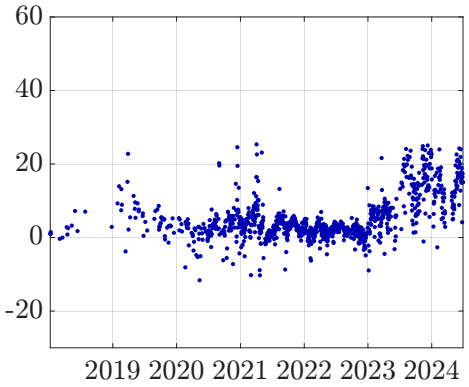
**Figure 4: Perceived Expected Return on Bitcoin under Bayesian and Recency-Biased Beliefs**



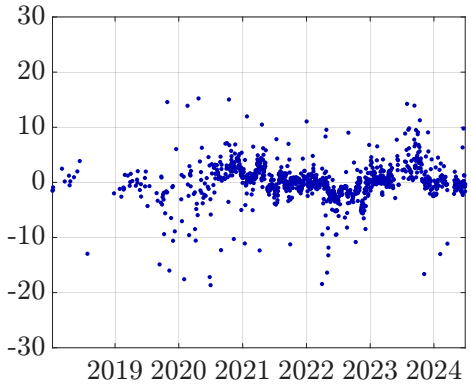

This figure shows the daily realized returns on Bitcoin (Panel A) and the corresponding perceived average return on Bitcoin under different degrees of recency bias with parameter values $\nu_B = \{0, 0.01, 0.02\}$ in Panel B.

**Figure 5: Bitcoin-Specific Risk Aversion**

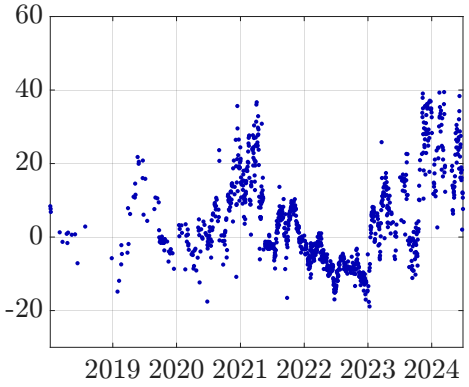
Panel A: $\mathcal{R}_{B,t}\left(R_{B,t+30} = -5\%; \nu_B = 0\right)$



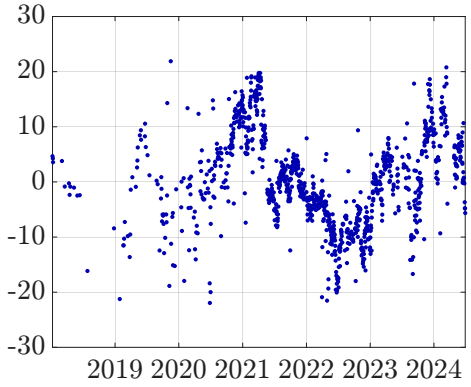

Panel B: $\mathcal{R}_{B,t}\left(R_{B,t+30} = 5\%; \nu_B = 0\right)$



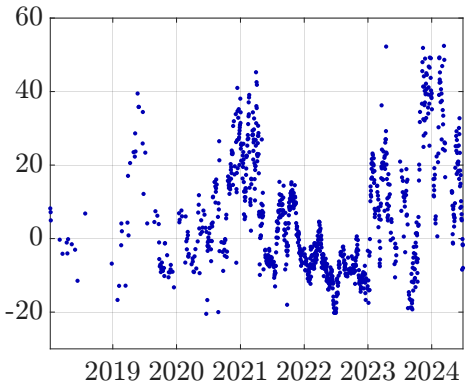

Panel C: $\mathcal{R}_{B,t}\left(R_{B,t+30} = -5\%; \nu_B = 0.01\right)$



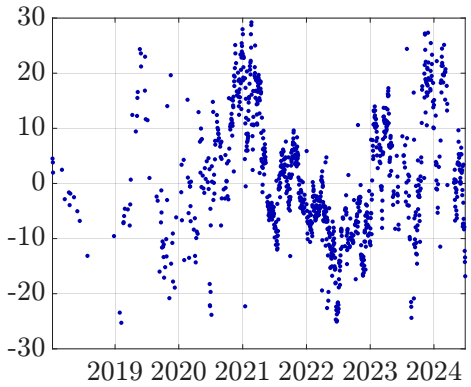

Panel D: $\mathcal{R}_{B,t}\left(R_{B,t+30} = 5\%; \nu_B = 0.01\right)$




Panel E: $\mathcal{R}_{B,t}\left(R_{B,t+30} = -5\%; \nu_B = 0.02\right)$



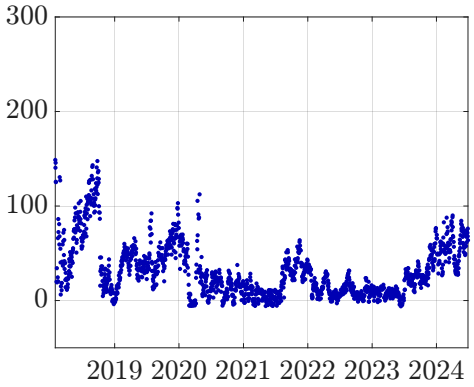

Panel F: $\mathcal{R}_{B,t}\left(R_{B,t+30} = 5\%; \nu_B = 0.02\right)$



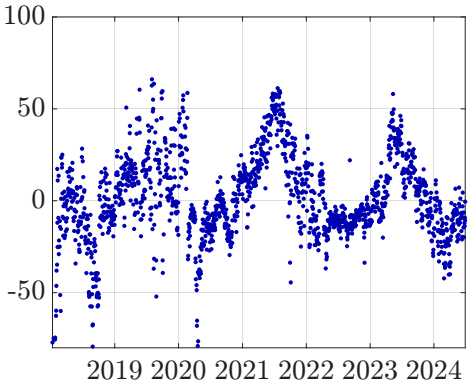

This figure shows daily estimates of Bitcoin-specific risk aversion, $\mathcal{R}_{B,t}(R_{B,t+30}; \nu_B)$, evaluated at Bitcoin returns of –5% and 5% over the subsequent month. Panels A and B report estimates for $\nu_B = 0$, Panels C and D for $\nu_B = 0.01$, and Panels E and F for $\nu_B = 0.02$.

**Figure 6: Equity-Specific Risk Aversion**

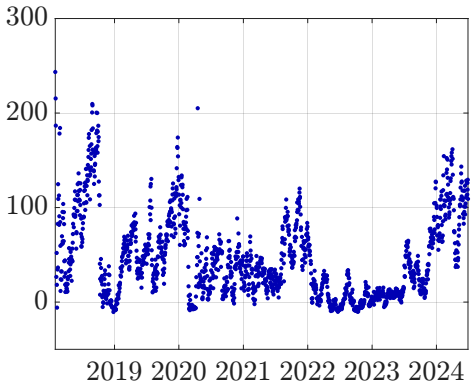
Panel A: $\mathcal{R}_{S,t}\left(R_{S,t+30} = -5\%; \nu_S = 0\right)$



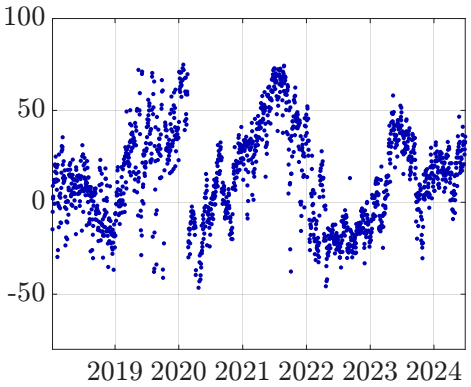

Panel B: $\mathcal{R}_{S,t}\left(R_{S,t+30} = 5\%; \nu_S = 0\right)$



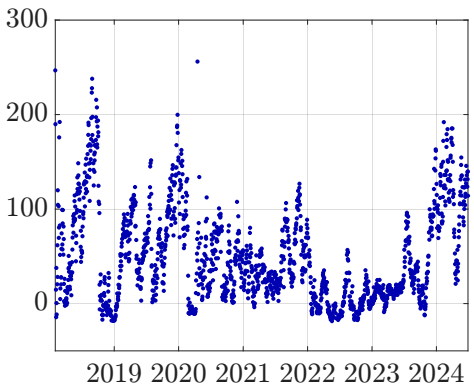

Panel C: $\mathcal{R}_{S,t}\left(R_{S,t+30} = -5\%; \nu_S = 0.01\right)$



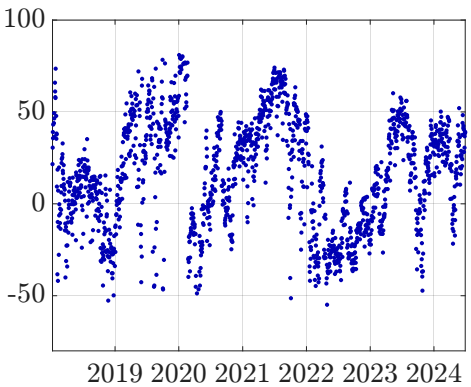

Panel D: $\mathcal{R}_{S,t}\left(R_{S,t+30} = 5\%; \nu_S = 0.01\right)$




Panel E: $\mathcal{R}_{S,t}\left(R_{S,t+30} = -5\%; \nu_S = 0.02\right)$

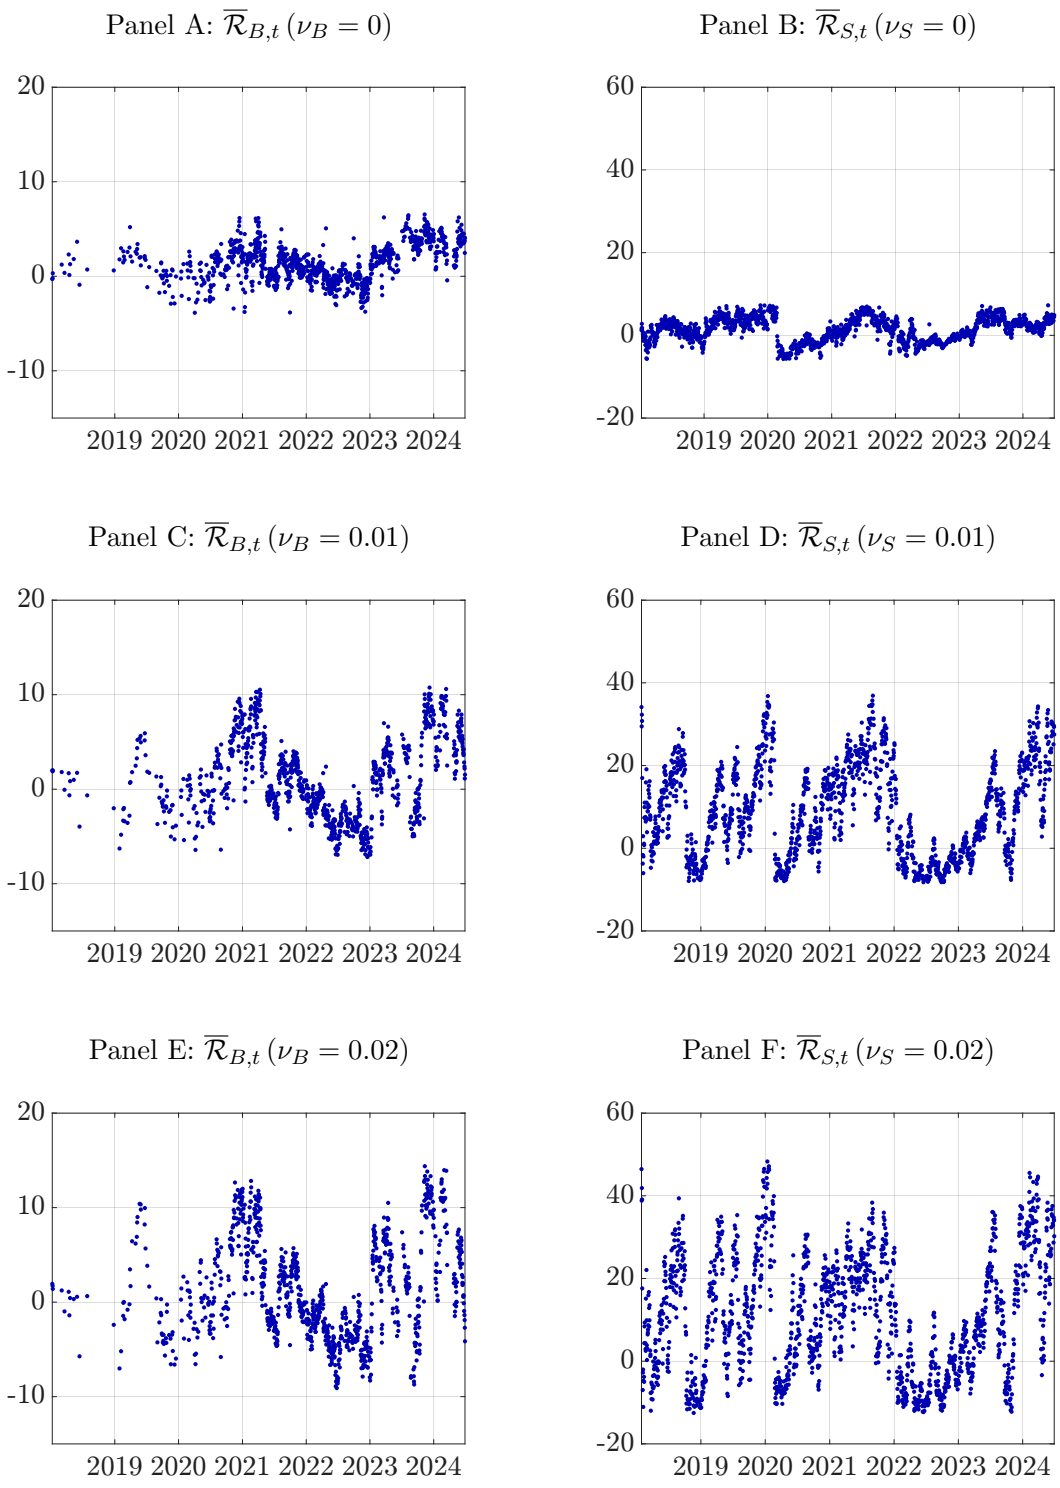



Panel F: $\mathcal{R}_{S,t}\left(R_{S,t+30} = 5\%; \nu_S = 0.02\right)$



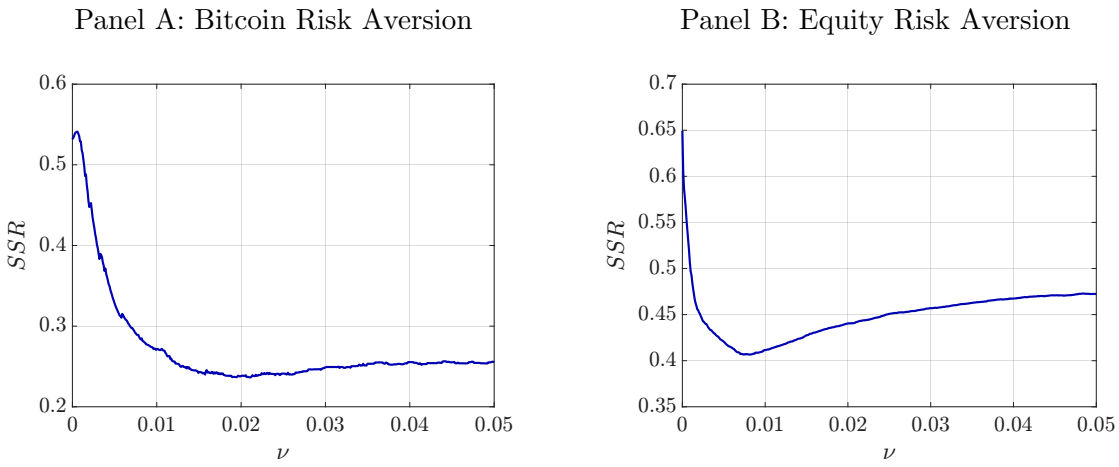

This figure shows daily estimates of equity-specific risk aversion, $\mathcal{R}_{S,t}(R_{S,t+30}; \nu_S)$, evaluated at market returns of –5% and 5% over the subsequent month. Panels A and B report estimates for $\nu_S = 0$, Panels C and D for $\nu_S = 0.01$, and Panels E and F for $\nu_S = 0.02$.

**Figure 7: Expected Risk Aversion**




This figure shows daily estimates of expected risk aversion, $\overline{\mathcal{R}}_{i,t}(\nu_i)$, where the security $i \in \{B, S\}$ is either Bitcoin or the equity market index. Panels A and B report the estimates for $\nu_i = 0$, Panels C and D for $\nu_i = 0.01$, and Panels E and F for $\nu_i = 0.02$.

**Figure 8: Goodness of Fit of Security-Specific Risk Aversion Regressions**




This figure shows the goodness of fit in the regression of the option-implied risk aversion across alternative belief formation parameters. In both panels, $y$-axis indicates the sum of squared normalized residuals from the following regression specification:

$$\overline{\mathcal{R}}_{i,t}(\nu_i) = \alpha + \beta \times \mu_{i,t}(\nu_i) + \gamma \times \text{Controls}_t + \epsilon_{i,t}$$

where the security $i \in \{B, S\}$ is either Bitcoin or the S&P 500 index. $\overline{\mathcal{R}}_{i,t}(R_{i,t+1})$ is the security-specific expected risk aversion defined in equation (28) . Controls are the same as in equation (27). For the belief formation parameter $\nu_i$, we use 500 different values between 0 and 0.05. For each value of $\nu$, we estimate the perceived average $\mu_{i,t}(\nu_i)$ and conduct the time-series regression. Panel A presents the sum of residuals in the Bitcoin market, and Panel B is for the equity market.

**Table 1: Summary Statistics**

| | Mean | Std.Dev. | 5% | 25% | 50% | 75% | 95% |
|---|---|---|---|---|---|---|---|
| Bitcoin Options | | | | | | | |
| Implied Volatility | 0.7375 | 0.2277 | 0.4205 | 0.5772 | 0.7103 | 0.8607 | 1.1511 |
| Moneyness ($K/F$) | 1.0544 | 0.2623 | 0.7148 | 0.9004 | 1.0218 | 1.1582 | 1.5119 |
| Time to Maturity (days) | 41.3033 | 36.4555 | 10 | 15 | 28 | 54 | 127 |
| $\mathcal{R}_{B,t}(R_{B,t+1} = -5\%; \nu_B = 0)$ | 5.5450 | 16.3123 | -1.7651 | 1.4823 | 3.4121 | 7.1145 | 19.6223 |
| $\mathcal{R}_{B,t}(R_{B,t+1} = 5\%; \nu_B = 0)$ | 0.0529 | 5.0947 | -5.2429 | -1.2326 | 0.1413 | 1.6177 | 5.8315 |
| $\overline{\mathcal{R}}_{B,t}(\nu_B = 0)$ | 1.4290 | 4.0262 | -1.8077 | -0.0477 | 1.3026 | 2.7084 | 4.9703 |
| $\overline{\mathcal{R}}_{B,t}(\nu_B = 0.021)$ | 1.4892 | 6.0911 | -6.2531 | -2.7677 | 0.4847 | 5.3075 | 11.2949 |
| Equity Options | | | | | | | |
| Implied Volatility | 0.2006 | 0.0857 | 0.0937 | 0.1379 | 0.1888 | 0.2449 | 0.3482 |
| Moneyness ($K/F$) | 0.9846 | 0.0999 | 0.8273 | 0.9210 | 0.9828 | 1.0437 | 1.1460 |
| Time to Maturity (days) | 66.4823 | 45.1741 | 14 | 28 | 56 | 98 | 156 |
| $\mathcal{R}_{S,t}(R_{S,t+1} = -5\%; \nu_S = 0)$ | 35.8881 | 37.1280 | -1.0676 | 84701 | 24.7519 | 56.4483 | 93.9251 |
| $\mathcal{R}_{S,t}(R_{S,t+1} = 5\%; \nu_S = 0)$ | -5.1753 | 26.8353 | -36.6753 | -14.5075 | -6.9149 | 3.5761 | 30.6537 |
| $\overline{\mathcal{R}}_{S,t}(\nu_S = 0)$ | 0.4874 | 2.7843 | -3.5738 | -1.6749 | 0.2855 | 2.4335 | 5.2964 |
| $\overline{\mathcal{R}}_{S,t}(\nu_S = 0.008)$ | 6.6190 | 10.8186 | -7.4745 | -2.6984 | 5.2001 | 14.0657 | 26.1328 |
| Underlying Securities | | | | | | | |
| Daily Excess Returns of Bitcoin | 0.0013 | 0.0368 | -0.0589 | -0.0129 | 0.0013 | 0.0168 | 0.0577 |
| Daily Excess Returns of S&P 500 | $1.72e^{-4}$ | 0.0103 | -0.0155 | -0.0045 | 0.0003 | 0.0051 | 0.0152 |
| Macroeconomic Controls | | | | | | | |
| Industrial Production Growth | 0.0081 | 0.0510 | -0.0481 | -0.0106 | 0.0114 | 0.0316 | 0.0657 |
| Consumption Growth | 0.0275 | 0.0458 | -0.0080 | 0.0128 | 0.0244 | 0.0432 | 0.0739 |
| Consumer Price Index Growth | 0.0164 | 0.0128 | 0.0028 | 0.0086 | 0.0129 | 0.0204 | 0.0421 |
| NFCI | 0.0025 | 0.9998 | -0.8043 | -0.6117 | -0.3698 | 0.2182 | 2.3230 |
| Bitcoin Controls | | | | | | | |
| Network | 18.3824 | 0.2988 | 17.8826 | 18.2001 | 18.3860 | 18.5749 | 18.8744 |
| Attention | 1.3401 | 0.2239 | 1.0021 | 1.1761 | 1.3222 | 1.4914 | 1.6981 |

The table presents summary statistics for the estimates and control variables. For security $i \in \{B, S\}$, $\mathcal{R}_i\left(R_{i,t+1}; \nu_i = 0\right)$ is the security-specific risk aversion evaluated at the future return $R_{i,t+1}$ of 5% or –5%, and $\overline{\mathcal{R}}_{i,t}(\nu_i = 0)$ is the security-specific expected risk aversion. Both measures are constructed under the assumption that the recency bias parameter $\nu_i$ is zero. $\overline{\mathcal{R}}_{B,t}(\nu_B = 0.021)$ and $\overline{\mathcal{R}}_{S,t}(\nu_S = 0.008)$ are alternative measures of expected risk aversion under the estimated recency bias parameter for each security. Macroeconomic controls include industrial production growth ($\Delta$IP), consumption growth ($\Delta$PCE), consumer price index growth ($\Delta$CPI), and the Chicago Fed National Financial Conditions Index (NFCI). The growth rates are annualized. Bitcoin controls include the attention indicator, constructed from Google search volume, and the network indicator, calculated as the first principal component extracted from the numbers of transactions and active accounts in the Bitcoin network.

## Table 2: Determinants of Bitcoin-Specific Risk Aversion

| Panel A: Risk Aversion at $R_{B,t+30} = -5\%$ | | | | | | |
|---|---|---|---|---|---|---|
| | $\nu_B = 0$ | | $\nu_B = 0.01$ | | $\nu_B = 0.02$ | |
| $\mu_{B,t}$ | 27.687*** | 22.025** | 10.456*** | 9.491*** | 10.259*** | 10.207*** |
| | (4.754) | (2.542) | (15.149) | (10.910) | (15.903) | (15.916) |
| $\sigma_{B,t}$ | −38.160*** | −31.511*** | −30.936*** | −29.386*** | −27.502*** | −31.718*** |
| | (-9.042) | (-6.605) | (-4.632) | (-4.827) | (-3.928) | (-4.626) |
| ΔIP | | 0.117 | | 1.140 | | 1.633 |
| | | (0.204) | | (1.113) | | (1.313) |
| ΔPCE | | -0.284 | | −1.329* | | −2.038** |
| | | (-0.689) | | (-1.690) | | (-2.098) |
| NFCI | | −3.135* | | −4.953* | | −4.626* |
| | | (-1.954) | | (-1.763) | | (-1.717) |
| ΔCPI | | -0.685 | | -0.141 | | 8.524 |
| | | (-0.184) | | (-0.020) | | (1.095) |
| $\text{Network}_B$ | | 2.324** | | 2.985 | | 0.132 |
| | | (2.243) | | (1.129) | | (0.049) |
| $\text{Attention}_B$ | | -1.720 | | 0.004 | | 0.356 |
| | | (-1.574) | | (0.002) | | (0.199) |
| $N$ | 982 | 982 | 982 | 982 | 982 | 982 |
| adj-$R^2$ | 0.496 | 0.525 | 0.732 | 0.742 | 0.761 | 0.775 |
| Panel B: Risk Aversion at $R_{B,t+30} = 5\%$ | | | | | | |
| | $\nu_B = 0$ | | $\nu_B = 0.01$ | | $\nu_B = 0.02$ | |
| $\mu_{B,t}$ | 14.128*** | 16.888** | 7.887*** | 7.920*** | 8.029*** | 8.293*** |
| | (3.902) | (2.285) | (15.189) | (16.498) | (15.443) | (15.309) |
| $\sigma_{B,t}$ | −8.430*** | −7.264** | −9.842*** | −9.189*** | −8.412* | −11.366*** |
| | (-3.772) | (-2.278) | (-3.874) | (-3.425) | (-1.880) | (-2.736) |
| ΔIP | | 0.536 | | 1.488** | | 0.989 |
| | | (1.583) | | (1.964) | | (1.157) |
| ΔPCE | | -0.158 | | −1.255** | | -0.884 |
| | | (-0.515) | | (-2.131) | | (-1.268) |
| NFCI | | -0.369 | | -0.444 | | -0.636 |
| | | (-0.330) | | (-0.314) | | (-0.317) |
| ΔCPI | | 0.199 | | 2.284 | | 11.228*** |
| | | (0.113) | | (0.641) | | (2.353) |
| $\text{Network}_B$ | | 0.595 | | 0.465 | | -0.377 |
| | | (0.772) | | (0.343) | | (-0.193) |
| $\text{Attention}_B$ | | −1.181* | | -0.257 | | -0.017 |
| | | (-1.701) | | (-0.238) | | (-0.014) |
| $N$ | 982 | 982 | 982 | 982 | 982 | 982 |
| adj-$R^2$ | 0.133 | 0.187 | 0.723 | 0.732 | 0.728 | 0.743 |

The table reports regressions of Bitcoin-specific risk aversion $\mathcal{R}_{B,t}(R_{B,t+30}; \nu_B)$ on the perceived expected return $\mu_{B,t}$ and controls, for various levels of recency bias, $\nu_B = \{0, 0.01, 0.02\}$. Panel A presents the result for the risk aversion evaluated at $R_{B,t+30} = -5\%$, and Panel B is for the risk aversion at $R_{B,t+30} = 5\%$. Controls include the Bitcoin's conditional volatility ($\sigma_{B,t}$), growth in industrial production (ΔIP), consumption growth (ΔPCE), Chicago Fed National Financial Condition Index (NFCI), inflation (ΔCPI), Bitcoin's network factor ($\text{Network}_B$) and attention indicator ($\text{Attention}_B$). Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

## Table 3: Determinants of Equity-Specific Risk Aversion

| Panel A: Risk Aversion at $R_{S,t+30} = -5\%$ | | | | | | |
|---|---|---|---|---|---|---|
| | $\nu_S = 0$ | | $\nu_S = 0.01$ | | $\nu_S = 0.02$ | |
| $\mu_{S,t}$ | $-318.417^{**}$ | -116.988 | $137.515^{***}$ | $154.384^{***}$ | $114.967^{***}$ | $136.519^{***}$ |
| | (-2.434) | (-0.819) | (8.150) | (16.291) | (7.392) | (14.833) |
| $\sigma_{S,t}$ | $-444.974^{***}$ | $-402.083^{***}$ | $-489.204^{***}$ | $-526.026^{***}$ | $-505.216^{***}$ | $-583.818^{***}$ |
| | (-14.079) | (-13.046) | (-12.761) | (-16.053) | (-12.425) | (-15.692) |
| $\Delta$IP | | 0.232 | | 2.887 | | 6.574 |
| | | (0.105) | | (0.801) | | (1.506) |
| $\Delta$PCE | | -2.502 | | $-6.467^{**}$ | | $-11.649^{***}$ |
| | | (-1.208) | | (-2.062) | | (-3.625) |
| NFCI | | $-11.582^{***}$ | | 7.437 | | $13.721^{**}$ |
| | | (-3.168) | | (1.123) | | (2.292) |
| $\Delta$CPI | | $-56.421^{***}$ | | $-88.535^{***}$ | | $-92.081^{***}$ |
| | | (-4.047) | | (-5.604) | | (-5.203) |
| $N$ | 5,739 | 5,739 | 5,739 | 5,739 | 5,739 | 5,739 |
| adj-$R^2$ | 0.542 | 0.582 | 0.622 | 0.652 | 0.596 | 0.637 |
| Panel B: Risk Aversion at $R_{S,t+30} = 5\%$ | | | | | | |
| | $\nu_S = 0$ | | $\nu_S = 0.01$ | | $\nu_S = 0.02$ | |
| $\mu_{S,t}$ | $408.606^{***}$ | $367.976^{***}$ | $67.910^{***}$ | $77.934^{***}$ | $61.711^{***}$ | $68.306^{***}$ |
| | (4.201) | (3.504) | (10.204) | (14.348) | (9.307) | (12.923) |
| $\sigma_{S,t}$ | $45.878^{***}$ | $65.750^{***}$ | -20.808 | $-34.413^{*}$ | $-41.405^{**}$ | $-58.310^{***}$ |
| | (4.245) | (5.725) | (-1.482) | (-1.917) | (-2.373) | (-2.669) |
| $\Delta$IP | | -2.201 | | -3.042 | | -0.001 |
| | | (-1.276) | | (-1.175) | | (-0.000) |
| $\Delta$PCE | | 0.357 | | -2.831 | | $-6.066^{***}$ |
| | | (0.243) | | (-1.509) | | (-2.803) |
| NFCI | | $-4.353^{***}$ | | 3.730 | | 3.675 |
| | | (-3.424) | | (1.312) | | (1.044) |
| $\Delta$CPI | | -1.009 | | -2.060 | | 1.536 |
| | | (-0.083) | | (-0.163) | | (0.105) |
| $N$ | 5,739 | 5,739 | 5,739 | 5,739 | 5,739 | 5,739 |
| adj-$R^2$ | 0.048 | 0.068 | 0.395 | 0.413 | 0.442 | 0.457 |

The table reports regressions of equity-specific risk aversion $\mathcal{R}_{S,t}(R_{S,t+30}; \nu_S)$ on the perceived expected return $\mu_{S,t}$ and controls, for various levels of recency bias, $\nu_S = \{0, 0.01, 0.02\}$. Panel A presents the result for the risk aversion evaluated at $R_{S,t+30} = -5\%$, and Panel B is for the risk aversion at $R_{S,t+1} = 5\%$. Controls include the stock market conditional volatility ($\sigma_{S,t}$), growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), Chicago Fed National Financial Condition Index (NFCI), and inflation ($\Delta$CPI). Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

## Table 4: Determinants of Expected Risk Aversion

| Panel A: Bitcoin | | | | | | |
|---|---|---|---|---|---|---|
| | $\nu_B = 0$ | | $\nu_B = 0.01$ | | $\nu_B = 0.02$ | |
| $\mu_{B,t}$ | 14.715*** | 11.582*** | 3.953*** | 3.772*** | 3.781*** | 3.794*** |
| | (6.350) | (2.607) | (12.505) | (11.452) | (13.058) | (14.620) |
| $\sigma_{B,t}$ | −13.208*** | −9.909*** | −7.985*** | −8.034*** | −6.452*** | −8.429*** |
| | (-9.471) | (-4.951) | (-3.657) | (-4.130) | (-2.790) | (-3.688) |
| ΔIP | | 0.370 | | 0.781 | | 0.771 |
| | | (1.480) | | (1.589) | | (1.332) |
| ΔPCE | | -0.217 | | −0.776** | | −0.816* |
| | | (-1.487) | | (-2.269) | | (-1.873) |
| NFCI | | -1.035 | | -0.985 | | -1.381 |
| | | (-1.385) | | (-0.984) | | (-1.326) |
| ΔCPI | | -1.584 | | 0.786 | | 3.949 |
| | | (-1.252) | | (0.300) | | (1.398) |
| $\text{Network}_B$ | | 1.143*** | | 0.590 | | -0.067 |
| | | (2.834) | | (0.617) | | (-0.069) |
| $\text{Attention}_B$ | | -0.456 | | 0.206 | | 0.322 |
| | | (-0.972) | | (0.299) | | (0.424) |
| $N$ | 982 | 982 | 982 | 982 | 982 | 982 |
| adj-$R^2$ | 0.442 | 0.490 | 0.716 | 0.729 | 0.742 | 0.762 |
| Panel B: Stock Market | | | | | | |
| | $\nu_S = 0$ | | $\nu_S = 0.01$ | | $\nu_S = 0.02$ | |
| $\mu_{S,t}$ | 75.122*** | 85.935*** | 40.720*** | 47.921*** | 36.314*** | 41.229*** |
| | (4.381) | (4.834) | (8.853) | (12.917) | (8.434) | (12.296) |
| $\sigma_{S,t}$ | −19.907*** | −14.535*** | −45.581*** | −60.310*** | −58.577*** | −76.248*** |
| | (-8.900) | (-6.242) | (-5.753) | (-6.657) | (-5.733) | (-6.814) |
| ΔIP | | −0.459* | | 0.423 | | 1.724 |
| | | (-1.805) | | (0.335) | | (1.106) |
| ΔPCE | | -0.432 | | -1.417 | | −2.956*** |
| | | (-1.580) | | (-1.613) | | (-2.604) |
| NFCI | | −1.156*** | | 3.535** | | 3.529** |
| | | (-5.140) | | (2.271) | | (2.025) |
| ΔCPI | | -3.414 | | -1.394 | | -3.706 |
| | | (-1.542) | | (-0.249) | | (-0.539) |
| $N$ | 5739 | 5739 | 5739 | 5739 | 5739 | 5739 |
| adj-$R^2$ | 0.302 | 0.351 | 0.569 | 0.589 | 0.541 | 0.560 |

The table reports regressions of security-specific expected risk aversion $\overline{\mathcal{R}}_{i,t}(\nu_i)$ on the perceived expected return $\mu_{i,t}(\nu_i)$ and controls, for various levels of recency bias, $\nu_i = \{0, 0.01, 0.02\}$. Panel A presents the results for Bitcoin, and Panel B presents the results for the stock market. Controls include the security's conditional volatility ($\sigma_{i,t}$), growth in industrial production (ΔIP), consumption growth (ΔPCE), Chicago Fed National Financial Condition Index (NFCI), and inflation (ΔCPI). For Bitcoin, the network factor ($\text{Network}_B$) and attention indicator ($\text{Attention}_B$) are included as additional controls. Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

**Table 5: Estimation of Recency Bias Parameter**

| | Bitcoin | Equity |
|---|---|---|
| Parameters | | |
| $\nu_B$ | $0.021^{***}$ | |
| | (26.573) | |
| $\nu_S$ | | $0.008^{***}$ |
| | | (54.834) |
| Regression Coefficients | | |
| $\mu_{B,t}$ | $3.810^{***}$ | |
| | (14.603) | |
| $\sigma_{B,t}$ | $-8.550^{***}$ | |
| | (-3.813) | |
| $\mu_{S,t}$ | | $48.734^{***}$ |
| | | (13.018) |
| $\sigma_{S,t}$ | | $-54.560^{***}$ |
| | | (-6.688) |
| $\Delta$IP | 0.736 | 0.254 |
| | (1.265) | (0.217) |
| $\Delta$PCE | $-0.801^{*}$ | -1.227 |
| | (-1.808) | (-1.510) |
| NFCI | -1.448 | $2.910^{**}$ |
| | (-1.371) | (1.982) |
| $\Delta$CPI | 4.470 | -2.450 |
| | (1.562) | (-0.468) |
| $\text{Network}_B$ | -0.198 | |
| | (-0.202) | |
| $\text{Attention}_B$ | 0.290 | |
| | (0.377) | |
| $N$ | 982 | 5,739 |
| adj-$R^2$ | 0.764 | 0.594 |

The table presents the recency bias parameter estimates for Bitcoin and the stock market by regressing security-specific expected risk aversion $\overline{\mathcal{R}}_{i,t}(\nu_i)$ on the perceived expected return $\mu_{i,t}(\nu_i)$ and controls. Controls include the security's conditional volatility ($\sigma_{i,t}$), growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), Chicago Fed National Financial Condition Index (NFCI), and inflation ($\Delta$CPI). For Bitcoin, the network factor ($\text{Network}_B$) and attention indicator ($\text{Attention}_B$) are included as additional controls. Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

**Table 6: Alternative Estimation of Recency Bias Parameter**

| | Bitcoin | Equity |
|---|---|---|
| Parameters | | |
| $\nu_B$ | $0.050^{***}$ | |
| | (4.017) | |
| $\nu_S$ | | $0.015^{***}$ |
| | | (5.206) |
| Regression Coefficients | | |
| $\mu_{B,t}$ | $0.348^{***}$ | |
| | (4.016) | |
| $\sigma_{B,t}$ | $-6.836^{***}$ | |
| | (-5.691) | |
| $\mu_{S,t}$ | | $5.400^{***}$ |
| | | (7.129) |
| $\sigma_{S,t}$ | | $-26.666^{***}$ |
| | | (-10.490) |
| $\Delta$IP | 0.288 | $-0.842^{**}$ |
| | (1.116) | (-2.567) |
| $\Delta$PCE | $-0.293^{*}$ | $-0.613^{**}$ |
| | (-1.662) | (-2.342) |
| NFCI | $-1.875^{***}$ | -0.519 |
| | (-3.536) | (-1.569) |
| $\Delta$CPI | 0.074 | -0.410 |
| | (0.066) | (-0.221) |
| $\text{Network}_B$ | $0.960^{***}$ | |
| | (2.791) | |
| $\text{Attention}_B$ | -0.172 | |
| | (-0.493) | |
| $N$ | 982 | 5,739 |
| adj-$R^2$ | 0.548 | 0.482 |

The table presents the recency bias parameter estimates for Bitcoin and the stock market by regressing security-specific expected risk aversion $\overline{\mathcal{R}}_{i,t}(\nu_i = 0)$ on the perceived expected return $\mu_{i,t}(\nu_i)$ and controls. Controls include the security's conditional volatility ($\sigma_{i,t}$), growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), Chicago Fed National Financial Condition Index (NFCI), and inflation ($\Delta$CPI). For Bitcoin, the network factor ($\text{Network}_B$) and attention indicator ($\text{Attention}_B$) are included as additional controls. Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

# Internet Appendix

## A Proofs of the Portfolio Choice Propositions

### A.1 Proof of Proposition 1

In this section, we derive the portfolio allocation that maximizes the agent's utility. The expectation of the approximated utility is

$$\begin{aligned}\mathbb{E}_t\left[u(W_{t+1})\right] &= \mathbb{E}_t\left[u(1+R_{f,t}+\mathbf{w}_t^\top \mathbf{R}_{t+1})\right] \\ &\approx \mathbb{E}_t\left[u(1+R_{f,t})+u'(1+R_{f,t})\mathbf{w}_t^\top \mathbf{R}_{t+1}+\frac{u''(1+R_{f,t})}{2}\mathbf{w}_t^\top \mathbf{R}_{t+1}\mathbf{R}_{t+1}^\top \mathbf{w}\right] \\ &= u(1+R_{f,t})+u'(1+R_{f,t})\mathbf{w}_t^\top \boldsymbol{\mu}_t+\frac{u''(1+R_{f,t})}{2}\mathbf{w}_t^\top\left(\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)\mathbf{w}_t.\end{aligned} \tag{33}$$

To find the optimal allocation, we differentiate the approximated expected utility with respect to the portfolio weight:

$$\frac{\partial \mathbb{E}_t\left[u(W_{t+1})\right]}{\partial \mathbf{w}_t} \approx u'(1+R_{f,t})\boldsymbol{\mu}_t+u''(1+R_{f,t})\left(\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)\mathbf{w}_t=0. \tag{34}$$

The FOC leads to the optimal allocation:

$$\mathbf{w}_t^*=-\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})}\left(\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)^{-1}\boldsymbol{\mu}_t. \tag{35}$$

### A.2 Proof of Proposition 2

We consider a simple case where there are two risky assets, equity and Bitcoin. We first obtain the inverse matrix in the solution of the optimal allocation:

$$\begin{aligned}\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top &= \begin{bmatrix}\sigma_{B,t}^2+\mu_{B,t}^2 & \mu_{B,t}\mu_{S,t}\\ \mu_{B,t}\mu_{S,t} & \sigma_{S,t}^2+\mu_{S,t}^2\end{bmatrix} \\ \left(\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)^{-1} &= \frac{1}{\sigma_{S,t}^2\sigma_{B,t}^2+\mu_{S,t}^2\sigma_{B,t}^2+\mu_{B,t}^2\sigma_{S,t}^2}\begin{bmatrix}\sigma_{S,t}^2+\mu_{S,t}^2 & -\mu_{B,t}\mu_{S,t}\\ -\mu_{B,t}\mu_{S,t} & \sigma_{B,t}^2+\mu_{B,t}^2\end{bmatrix}.\end{aligned} \tag{36}$$

Then, the weights for Bitcoin and equity are

$$\begin{aligned}\begin{bmatrix}w_{B,t}^*\\ w_{S,t}^*\end{bmatrix} &= -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})}\left(\boldsymbol{\Sigma}_t+\boldsymbol{\mu}_t\boldsymbol{\mu}_t^\top\right)^{-1}\boldsymbol{\mu}_t \\ &= -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})}\cdot\frac{1}{\sigma_{S,t}^2\sigma_{B,t}^2+\mu_{S,t}^2\sigma_{B,t}^2+\mu_{B,t}^2\sigma_{S,t}^2}\begin{bmatrix}\sigma_{S,t}^2+\mu_{S,t}^2 & -\mu_{B,t}\mu_{S,t}\\ -\mu_{B,t}\mu_{S,t} & \sigma_{B,t}^2+\mu_{B,t}^2\end{bmatrix}\begin{bmatrix}\mu_{B,t}\\ \mu_{S,t}\end{bmatrix} \\ &= -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})}\begin{bmatrix}\frac{\mu_{B,t}}{\sigma_{B,t}^2+\mu_{S,t}^2\sigma_{B,t}^2/\sigma_{S,t}^2+\mu_{B,t}^2}\\ \frac{\mu_{S,t}}{\sigma_{S,t}^2+\mu_{S,t}^2+\mu_{B,t}^2\sigma_{S,t}^2/\sigma_{B,t}^2}\end{bmatrix}\end{aligned} \tag{37}$$

### A.3 Proof of Proposition 3

Using the analytic expression of the optimal weights provided by Proposition 2, we examine the dependence of the weights on the return statistics. First, to assess the volatility impact, we differentiate the weight respect to the variance:

$$\frac{\partial w^*_{i,t}}{\partial \sigma^2_{i,t}} = \frac{u'(1+R_{f,t})}{u''(1+R_{f,t})} \frac{\mu_{i,t}\left(1+\mu^2_{-i,t}/\sigma^2_{-i,t}\right)}{\left[\sigma^2_{i,t}+\mu^2_{-i,t}\sigma^2_{i,t}/\sigma^2_{-i,t}+\mu^2_{i,t}\right]^2} \tag{38}$$

where $i \in \{B, S\}$, and $-i$ indicates the other risky asset. Note that $u'(1+R_{f,t}) > 0$ and $u''(1+R_{f,t}) < 0$. Thus, as long as $\mu_{i,t} > 0$, $\frac{\partial w^*_{i,t}}{\partial \sigma^2_{i,t}} < 0$.

Next, to examine the impact of the average return, we differentiate as follows:

$$\frac{\partial w^*_{i,t}}{\partial \mu_{i,t}} = -\frac{u'(1+R_{f,t})}{u''(1+R_{f,t})} \frac{\sigma^2_{i,t}-\mu^2_{i,t}+\mu^2_{-i,t}\sigma^2_{i,t}/\sigma^2_{-i,t}}{\left[\sigma^2_{i,t}+\mu^2_{-i,t}\sigma^2_{i,t}/\sigma^2_{-i,t}+\mu^2_{i,t}\right]^2} \tag{39}$$

Under the realistic assumptions that $\sigma_{i,t} > \mu_{i,t}$, it follows that $\frac{\partial w^*_{i,t}}{\partial \mu_{i,t}} > 0$.

### A.4 Proof of Proposition 4

To characterize the absolute risk aversion, we first expand the expected marginal utility:

$$\begin{aligned} u'_{i,t}(R_{i,t+1}) &= \mathbb{E}_t\left[u'(1+\mathbf{w}_t^\top \mathbf{R}_{t+1}+R_{f,t})\middle|\, R_{i,t+1}\right] && (40)\\ &\approx \mathbb{E}_t\left[u'(1+R_{f,t})+u''(1+R_{f,t})(W_{t+1}-1-R_{f,t})\middle|\, R_{i,t+1}\right] \\ &= u'(1+R_{f,t})+u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t}+w^*_{i,t}R_{i,t+1}). && (41)\end{aligned}$$

According to the definition in equation (9), the risk aversion specific to security $i$ is

$$\mathcal{R}_i(R_{i,t+1}) = -\frac{\partial \log u'_{i,t}(R_{i,t+1})}{\partial R_{i,t+1}} = -\frac{u''(1+R_{f,t})w^*_{i,t}}{u'(1+R_{f,t})+u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t}+w^*_{i,t}R_{i,t+1})}. \tag{42}$$

The security-specific risk aversion rises with its weight $w^*_{i,t}$. Formally,

$$\begin{aligned} \frac{\partial \mathcal{R}_i\left(R_{i,t+1}\right)}{\partial w^*_{i,t}} &= \frac{-u''(1+R_{f,t})\left[u'(1+R_{f,t})+u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t}+w^*_{i,t}R_{i,t+1})\right]+u''(1+R_{f,t})w^*_{i,t}u''(1+R_{f,t})R_{i,t+1}}{\left[u'(1+R_{f,t})+u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t}+w^*_{i,t}R_{i,t+1})\right]^2} \\ &= \frac{-u''(1+R_{f,t})(u'(1+R_{f,t})+u''(1+R_{f,t})w^*_{-i,t}\mu_{-i,t})}{\left[u'(1+R_{f,t})+u''(1+R_{f,t})(w^*_{-i,t}\mu_{-i,t}+w^*_{i,t}R_{i,t+1})\right]^2} > 0 \end{aligned}$$

The inequality holds as long as the second term in the numerator, $u'(1+R_{f,t})+u''(1+R_{f,t})w^*_{-i,t}\mu_{-i,t}$ , is positive. This term is the same as marginal utility at $1+R_{f,t}+x$ when $x = w^*_{-i,t}\mu_{-i,t} \ll 1+R_{f,t}$ and the second-order Taylor approximation is employed:

$$\begin{aligned} \frac{\partial u(1+R_{f,t}+x)}{\partial x} &\approx \frac{\partial\left[u(1+R_{f,t})+u'(1+R_{f,t})x+0.5u''(1+R_{f,t})x^2\right]}{\partial x} \\ &= u'(1+R_{f,t})+u''(1+R_{f,t})x. \end{aligned} \tag{43}$$

In addition, according to Lemma 1, $w^*_{-i,t} > 0$ if $\mu_{-i,t} > 0$. As $x = w^*_{-i,t}\mu_{-i,t} > 0$, the marginal utility $\frac{\partial \tilde{u}(1+R_{f,t}+x)}{\partial x} > 0$, leading to $u'(1+R_{f,t}) + u''(1+R_{f,t})w^*_{-i,t}\mu_{-i,t} > 0$.

### A.5 Proof of Proposition 5

Combining the results outlined in Proposition 2 and 3, we derive the absolute risk aversion's dependence on the conditional return distribution.

$$\frac{\partial \mathcal{R}_i\left(R_{i,t+1}\right)}{\partial \mu_{i,t}} = \frac{\partial \mathcal{R}_i\left(R_{i,t+1}\right)}{\partial w^*_{i,t}} \frac{\partial w^*_{i,t}}{\partial \mu_{i,t}} > 0 \tag{44}$$

$$\frac{\partial \mathcal{R}_i\left(R_{i,t+1}\right)}{\partial \sigma^2_{i,t}} = \frac{\partial \mathcal{R}_i\left(R_{i,t+1}\right)}{\partial w^*_{i,t}} \frac{\partial w^*_{i,t}}{\partial \sigma^2_{i,t}} < 0 \tag{45}$$

if $0 < \mu_{i,t} < \sigma_{i,t}$.

## B Learning with Recency Bias

This section presents how beliefs about the average return of security $i$ are updated given the parameter $\nu_i$ that controls heuristic bias. The investors begins with prior beliefs about the true average return $m_i$ given by

$$\pi(m_i) = \frac{1}{\sqrt{2\varphi^2_{i,0}}} \exp\left(-\frac{(m_i - \mu_{i,0})^2}{2\varphi^2_{i,0}}\right). \tag{46}$$

The sampling probability of realized returns up to date $t$ under heuristic bias is

$$\prod_{s=1}^{t} p(R_{i,s}|m_i)^{(1-\nu_i)^{t-s}} = \prod_{s=1}^{t} \left[\frac{1}{\sqrt{2\sigma^2_{i,s-1}}} \exp\left(-\frac{(R_{i,s} - m_i)^2}{2\sigma^2_{i,s-1}}\right)\right]^{(1-\nu_i)^{t-s}}. \tag{47}$$

The posterior beliefs are formed according to Baye's law, and they are normally distributed as follows:

$$\begin{aligned}
&\pi(m_i|R_{i,1},..,R_{i,t};\nu_i) \propto \prod_{s=1}^{t} p(R_{i,s}|m_i)^{(1-\nu_i)^{t-s}} \times \pi(m_i)^{(1-\nu_i)^t} \\
&\propto \prod_{s=1}^{t} \left[\exp\left(-\frac{(R_{i,s} - m_i)^2}{2\sigma^2_{i,s-1}}\right)\right]^{(1-\nu_i)^{t-s}} \times \left[\exp\left(-\frac{(m_i - \mu_{i,0})^2}{2\varphi^2_{i,0}}\right)\right]^{(1-\nu_i)^t} \\
&= \exp\left(-\sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s}(R_{i,s} - m_i)^2}{2\sigma^2_{i,s-1}} - (1-\nu_i)^t \frac{(m_i - \mu_{i,0})^2}{2\varphi^2_{i,0}}\right) \\
&\propto \exp\left(-\frac{1}{2}\left[\sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s}}{\sigma^2_{i,s-1}} + \frac{(1-\nu_i)^t}{\varphi^2_{i,0}}\right] m_i^2 + \left[\sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s} R_{i,s}}{\sigma^2_{i,s-1}} + (1-\nu_i)^t \frac{\mu_{i,0}}{\varphi^2_{i,0}}\right] m_i\right) \\
&\propto \exp\left(-\frac{(m_i - \mu_{i,t})^2}{2\varphi^2_{i,t}}\right),
\end{aligned}$$

where the mean and standard deviation of distributions are

$$\mu_{i,t} = \varphi_{i,t}^2 \left[ \sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s} R_{i,s}}{\sigma_{i,s-1}^2} + \frac{(1-\nu_i)^t \mu_{i,0}}{\varphi_{i,0}^2} \right]$$
$$\frac{1}{\varphi_{i,t}^2} = \sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s}}{\sigma_{i,s-1}^2} + \frac{(1-\nu_i)^t}{\varphi_{i,0}^2}.$$

These formulations imply recursive updates of posterior beliefs:

$$\begin{aligned}
\mu_{i,t+1} &= \varphi_{i,t+1}^2 \left[ \sum_{s=1}^{t+1} \frac{(1-\nu_i)^{t+1-s} R_{i,s}}{\sigma_{i,s-1}^2} + \frac{(1-\nu_i)^{t+1} \mu_{i,0}}{\varphi_{i,0}^2} \right] \\
&= \varphi_{i,t+1}^2 \left[ \frac{R_{i,t+1}}{\sigma_{i,t}^2} + (1-\nu_i) \sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s} R_{i,s}}{\sigma_{i,s-1}^2} + (1-\nu_i) \frac{(1-\nu_i)^t \mu_{i,0}}{\varphi_{i,0}^2} \right] \\
&= \varphi_{i,t+1}^2 \left[ \frac{R_{i,t+1}}{\sigma_{i,t}^2} + (1-\nu_i) \frac{\mu_{i,t}}{\varphi_{i,t}^2} \right] \qquad (48) \\
\frac{1}{\varphi_{i,t+1}^2} &= \sum_{s=1}^{t+1} \frac{(1-\nu_i)^{t+1-s}}{\sigma_{i,s-1}^2} + \frac{(1-\nu_i)^{t+1}}{\varphi_{i,0}^2} \\
&= \frac{1}{\sigma_{i,t}^2} + (1-\nu_i) \sum_{s=1}^{t} \frac{(1-\nu_i)^{t-s}}{\sigma_{i,s-1}^2} + (1-\nu_i) \frac{(1-\nu_i)^t}{\varphi_{i,0}^2} \\
&= \frac{1}{\sigma_{i,t}^2} + (1-\nu_i) \frac{1}{\varphi_{i,t}^2}. \qquad (49)
\end{aligned}$$

# C Nonlinear Least Squares Estimation

In section 4.5, we estimate the nonlinear model with the set of the recency-bias parameter and regression coefficients $\Theta_i = (\nu_i, \alpha, \beta, \gamma)$. We search for the set that minimizes the sum of squared normalized residuals:

$$\hat{\Theta}_i = \arg\min_{\Theta_i} \sum_{t=1}^{T} \varepsilon_{i,t}(\Theta_i)^2 = \arg\min_{\Theta_i} \sum_{t=1}^{T} \left[ \frac{\overline{\mathcal{R}}_{i,t}(\nu_i) - \widehat{\mathcal{R}}_{i,t}(\Theta_i)}{\sqrt{\mathrm{SST}(\nu_i)}} \right]^2 \qquad (50)$$

where $\widehat{\mathcal{R}}_{i,t}(\Theta_i) = \alpha + \beta \times \mu_{i,t}(\nu_i) + \gamma \times \mathrm{Controls}_t$. The estimates' standard errors are obtained from their asymptotic distribution with residuals' serial correlations taken into account. First, we numerically calculate the $T \times k$ Jacobian matrix $J(\hat{\Theta}_i)$ with rows $J_t(\hat{\Theta}_i) = \frac{\partial g_t(\Theta_i)}{\partial \Theta_i'}|_{\Theta_i = \hat{\Theta}_i}$. Next, we obtain the Newey-West estimator of covariance matrix, $J'\hat{\Gamma}J$:

$$\begin{aligned}
J'\hat{\Gamma}J = &\sum_{t=1}^{T} \hat{\epsilon}_t^2 {J_t}^\top(\hat{\Theta}_i) J_t(\hat{\Theta}_i) \qquad (51) \\
&+ \sum_{l=1}^{q} \left[ \left( 1 - \frac{l}{q+1} \right) \sum_{t=l+1}^{T} \left( \hat{\epsilon}_t {J_t}^\top(\hat{\Theta}_i) J_{t-l}(\hat{\Theta}_i) \hat{\epsilon}_{t-l} + \hat{\epsilon}_{t-l} J_{t-l}^\top(\hat{\Theta}_i) J_t(\hat{\Theta}_i) \hat{\epsilon}_t \right) \right],
\end{aligned}$$

where the lag length $q$ is set to be 30 days for daily estimates of risk aversion. Then, the variance-covariance of estimates is

$$\hat{\text{Var}}(\hat{\Theta}_i) = (J'J)^{-1} \left( J'\hat{\Gamma}J \right) (J'J)^{-1}. \tag{52}$$

# D Robustness Test - Alternative Data Frequency and Sample Period

In our main empirical analysis, we examine the daily estimate of risk aversion and its response to the perceived drift and other controls. Considering the persistence of the daily estimates and a potential bias in regression coefficients, we test the robustness of the findings using weekly estimates. For each security, we replace the dependent variable in equation (29) with the weekly estimates of expected risk aversion and estimate the model using the nonlinear least squares as in the main analysis. The results in Table D1 show that the estimates of recency bias derived from weekly data are quantitatively very similar to those obtained from daily data. Moreover, the security-specific risk aversion's empirical link to the perceived return distribution, average, and volatility is consistent with the model prediction.

Next, we use an alternative sample period for the stock-market estimation. Our baseline estimate utilizes the full sample of S&P 500 index options from January 1997 to June 2024. In this alternative estimation, we instead use a matched sample aligned with the availability of Bitcoin options, beginning in January 2017. The results in Table D2 show that the recency bias parameter estimated from this matched sample is quantitatively very similar to the estimate from the full sample.

**Table D1: Estimation of Recency Bias Parameter using Weekly Data**

| | Bitcoin | Equity |
|---|---|---|
| Parameters | | |
| $\nu_B$ | $0.020^{***}$ | |
| | (17.555) | |
| $\nu_S$ | | $0.007^{***}$ |
| | | (35.536) |
| Regression Coefficients | | |
| $\mu_{B,t}$ | $3.584^{***}$ | |
| | (7.832) | |
| $\sigma_{B,t}$ | $-8.674^{***}$ | |
| | (-6.573) | |
| $\mu_{S,t}$ | | $46.899^{***}$ |
| | | (8.217) |
| $\sigma_{S,t}$ | | $-48.338^{***}$ |
| | | (-4.154) |
| $\Delta$IP | $1.180^{**}$ | $-0.313$ |
| | (2.062) | (-0.267) |
| $\Delta$PCE | $-1.107^{***}$ | $-0.737$ |
| | (-2.678) | (-0.953) |
| NFCI | -1.554 | 2.442 |
| | (-1.031) | (1080) |
| $\Delta$CPI | 1.362 | -4.504 |
| | (0.380) | (-0.552) |
| $\text{Network}_B$ | -0.092 | |
| | (-0.060) | |
| $\text{Attention}_B$ | 1.021 | |
| | (1.218) | |
| $N$ | 158 | 1,073 |
| adj-$R^2$ | 0.749 | 0.585 |

The table presents the recency bias parameter estimates for Bitcoin and the stock market by regressing security-specific expected risk aversion $\overline{\mathcal{R}}_{i,t}(\nu_i)$ on the perceived expected return $\mu_{i,t}(\nu_i)$ and controls. Controls include the security's conditional volatility ($\sigma_{i,t}$), growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), Chicago Fed National Financial Condition Index (NFCI), and inflation ($\Delta$CPI). For Bitcoin, the network factor ($\text{Network}_B$) and attention indicator ($\text{Attention}_B$) are included as additional controls. Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

**Table D2: Estimation of Recency Bias Parameter in Stock Market Using Matched Sample**

| | Full Sample | Matched Sample |
|---|---|---|
| Parameter | | |
| $\nu_S$ | $0.008^{***}$ | $0.009^{***}$ |
| | (54.834) | (44.359) |
| Regression Coefficients | | |
| $\mu_{S,t}$ | $48.734^{***}$ | $62.278^{***}$ |
| | (13.018) | (7.360) |
| $\sigma_{S,t}$ | $-54.560^{***}$ | $-70.304^{***}$ |
| | (-6.688) | (-2.938) |
| $\Delta$IP | 0.254 | -0.714 |
| | (0.217) | (-0.504) |
| $\Delta$PCE | -1.227 | $-1.907^{*}$ |
| | (-1.510) | (-1.926) |
| NFCI | $2.910^{**}$ | -0.270 |
| | (1.982) | (-0.056) |
| $\Delta$CPI | -2.450 | 6.712 |
| | (-0.468) | (1.090) |
| $N$ | 5,739 | 1,467 |
| adj-$R^2$ | 0.594 | 0.673 |

The table presents the recency bias parameter estimates for the stock market by regressing security-specific expected risk aversion $\overline{\mathcal{R}}_{S,t}(\nu_S)$ on the perceived expected return $\mu_{S,t}(\nu_S)$ and controls. The first column shows the estimation result using the full sample of S&P 500 index options from January 1997 to June 2024. The second column reports the results using a matched sample aligned with the availability of Bitcoin options, beginning in January 2017. Controls include the security's conditional volatility ($\sigma_{i,t}$), growth in industrial production ($\Delta$IP), consumption growth ($\Delta$PCE), Chicago Fed National Financial Condition Index (NFCI), and inflation ($\Delta$CPI). Parentheses report $t$-statistics based on standard errors estimated with the Newey–West correction for serial correlation. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.