Feedback Trading and Bubbles

Brandon Yueyang Han[1] and Igor Makarov[*2]

[1]University of Maryland
[2]London School of Economics

May 1, 2022

ABSTRACT

The paper develops a model of bubbles that can be taken to the data and explain the behavior of asset prices and their statistics. We depart from the rational expectations framework and assume that investors are only boundedly rational. They observe the price process, but do not fully understand how its volatility and expected returns are determined in equilibrium. Investors learn about the market by looking at past prices. When they observe unexpectedly high returns, they infer that the asset must currently have a high Sharpe ratio, and therefore, allocate a higher share of their wealth to the asset, further increasing the asset price. The interaction of this feedback effect with investors' wealth effect determines the price dynamics and evolution of investors' beliefs in the model. We fit the model to cryptocurrency markets and show that it can successfully explain many empirical facts in these markets.

*Brandon Yueyang Han: Robert H. Smith School of Business, College Park, MD 20742, USA. Email: yhan1@umd.edu. Igor Makarov: Houghton Street, London WC2A 2AE, UK. Email: i.makarov@lse.ac.uk. We thank Georgy Chabakauri, Leonid Kogan, Peter Kondor, Pete Kyle, Kai Li, Mark Loewenstein, Cameron Peng, Oleg Rytchkov, Antoinette Schoar, Jiang Wang, and seminar participants at LSE, University of Maryland, University of Reading, and CICF for helpful comments. We also thank Andrew Patton for sharing his code to estimate realized volatility with us.

"I can calculate the movement of stars, but not the madness of men."
— Isaac Newton

There are few events in financial markets that capture public attention more than financial bubbles. Bubbles are a persistent and prevalent phenomenon. They are not confined to a particular country or time period. Yet, despite extensive research, our knowledge of bubbles is still incomplete. This paper aims to develop a realistic model of bubbles that can be taken to the data and explain the behavior of asset prices and their statistics.

The core of almost any model of bubbles is the dynamics of investors' beliefs about expected returns. The rational expectations framework postulates that investors have correct beliefs at every moment in time without making specific predictions about the source of this knowledge. In practice, the knowledge can be either an outcome of learning or a deep understanding of equilibrium. Neither of these two channels is likely to deliver complete knowledge in the case of bubbles since the majority of bubbles are short-lived, with investors displaying notoriously irrational behavior.

In this paper, we depart from the rational expectations framework and assume that investors are only boundedly rational. In our model, investors have CRRA preferences and can trade a risk-free asset and a risky asset. The risky asset does not offer any payoff, so investors trade it for purely speculative reasons. The main sources of uncertainty are random fluctuations in investors' wealth.

Investors observe the price process, but do not fully understand how its volatility and expected returns are determined in equilibrium. They attempt to learn about the market by looking at past prices.[1] They presume that the Sharpe ratio follows a mean-reverting process and try to infer its value from the data. When investors observe unexpectedly high returns, they infer that the asset must currently have a high Sharpe ratio and update their beliefs upwards. A higher estimate of the Sharpe ratio makes investors allocate a higher share of their wealth to the speculative asset, leading to a further increase in the price. This feedback effect amplifies investors' wealth shocks. Its interaction with investors' wealth effect determines the price dynamics and evolution of investors' beliefs in the model.

Investors' estimate of the Sharpe ratio is correct on average, but without the exact knowledge of the dynamics of the true Sharpe ratio, they make systematic mistakes. When investors are optimistic, i.e., their estimate of the Sharpe ratio is high, they invest a large share of their wealth in the risky asset. Therefore, the possibility that

[1]This is consistent with prior research which finds that investors' expectations of future prices are influenced by the history of past returns, Smith et al. (1988), Vissing-Jorgensen (2003), Haruvy et al. (2007), Amromin and Sharpe (2008), Greenwood and Shleifer (2014), Liao et al. (2021).

even more wealth is attracted to the asset is small, and as a consequence, the true risk premium and Sharpe ratio are low. In contrast, when investors are pessimistic, the wealth invested in the speculative asset is expected to increase, resulting in the true risk premium and the Sharpe ratio being high. This negative link stabilizes the dynamics of investors' beliefs and leads to the stationary dynamics of beliefs and asset returns.

In equilibrium, the price displays the typical bubble dynamics: it rises in the beginning fueled by the increasing optimism of investors, only to fall later. In steady state, the price and the aggregate investors' wealth are expected to decrease over time because the asset is unproductive, and investors consume over time, but the convergence is not monotone. The market goes through waves of optimism and pessimism.

Our model is set in continuous time and has an explicit solution, which allows us to compute many objects of interest such as volatility and the Sharpe ratio in closed form. The model predicts a novel pattern of volatility of assets affected by strong speculative demand. Typically, in equity markets, volatility is strongly negatively correlated with prices: volatility is high when prices are low and vice versa. In contrast, in our model, volatility is a U-shaped function of investors' beliefs and prices. To the best of our knowledge, this U-shaped pattern of volatility (and as we show later, volume as well) over bubble cycles is a unique and novel prediction of our model not shared with the existing literature.

The U-shape pattern of volatility in our model is a result of interaction of the wealth and feedback effects. When investors are pessimistic they invest a small fraction of their wealth in the asset, so the price is low. In this case, even a small wealth shock leads to a large proportional change in portfolio holdings, and as a consequence, to a large change in the price. Volatility is also high when investors are very optimistic and allocate a large fraction of their wealth to the speculative asset. A a result, the investors' wealth become volatile, which makes the asset price volatile.

Our model is also able to explain time-series momentum often observed during bubble periods. In the model, expected returns are positively autocorrelated because they are a function of persistent state-variables. Realized returns and expected returns tend to be negatively correlated because the true Sharpe ratio and investors' estimate of the Sharpe ratio are inversely related. The momentum is determined by the interplay of these two effects.

When the expected returns are high or low, the realized returns tend to be also high or low and continue to be high or low for some time in the future, leading to momentum in returns. When the expected returns are around zero, the market moves from slight positive expected returns to slight negative expected returns following positive realized returns and vice versa. This creates reversals in returns. In steady state, the second effect dominates. As a result, returns are slightly negatively correlated. But in the beginning of bubble periods and when the bubble crashes prices exhibit momentum.

Another prediction of our model is that assets affected by strong speculative demand can experience large price impact of trades. This prediction is consistent with the results of Makarov and Schoar (2020) who document large price impact in the bitcoin market, and unusually high price impact experienced by stocks like Yahoo and Tesla upon their inclusion in the S&P 500 index. In the model, a large price impact occurs because investors believe that prices are informative about investment opportunities and as a result, follow feedback strategies.

Because the model has only boundedly rational investors, a potential concern could be that the resulting price dynamics leads to many arbitrage opportunities, which in turn would question the plausibility of the model. To address this concern, we study return predictability and show that in steady state prices follow nearly a random walk, and opportunities to make profits by trading against boundedly rational investors for realistic parameter values are limited in our model.

We test the predictions of our model by applying them to cryptocurrency markets. Consistent with the assumption of our model, cryptocurrencies do not offer any payoff to their investors. Further in support of model assumptions, Makarov and Schoar (2020) show that there is a strong positive relation between investor flows and prices in cryptocurrency markets. The net order flow explains 80% of return variation — a substantially greater proportion than in traditional capital markets. Liu and Tsyvinski (2020) show that cryptocurrency returns do not significantly load on macroeconomic factors or cryptocurrency production factors but correlate with measures of investor sentiment. We consider the three most liquid and largest cryptocurrency markets: bitcoin, ethereum, and ripple, and show that consistent with the model's predictions, in each of these markets, volatility is indeed higher at the height and the bottom of cryptocurrency price cycles. We conjecture that the U-shaped behavior of volatility is not limited to cryptocurrency markets but holds more generally for all assets affected by strong speculative demand.

Our model has few parameters and can be efficiently fitted to the data. The model parameters and the price history uniquely determine the evolution of investors' beliefs about the Sharpe ratio, which in turn govern all statistics of asset returns. We fit the model to bitcoin data and show that the model can successfully match the time-series of moments returns.

Finally, in our main setting, investors have homogeneous beliefs. To study trading volume, we extend the model to allow investors to have heterogeneous priors and update them at different speed. In this paper, we focus on the case where there are two groups of investors. One group has constant beliefs about the Sharpe ratio. The other group, as before, uses the history of past prices to update their beliefs.

The main economic mechanism in the heterogeneous beliefs model is similar to that in the main setup, but prices display considerably more complicated dynamics. The asset pricing moments are no longer just functions of investors' beliefs but depend on the ratio of the aggregate wealth of the constant-belief and feedback investors. Following positive wealth shocks, the feedback investors increase their holdings of the speculative asset relative to the constant-belief investors. We show that in our extended setup trading volume and volatility are strongly positively correlated. In particular, similar to volatility, trading volume has a U-shaped pattern: It is high at the bottom and the height of price cycles. We show that these predictions hold in the bitcoin market.

Our paper is related to several strands of literature. First, it is connected to large literature that studies bubbles. We refer the reader to excellent surveys by Xiong (2013) and Brunnermeier and Oehmke (2013). In this literature, our paper is closest to the research that emphasizes investors' bounded rationality and the resulting use of feedback trading strategies. As in Fuster et al. (2012), investors in our paper adopt a parsimonious model to infer quantities of interest. Since the model does not nest the true model investors in our model tend to over extrapolate their beliefs. Similar to Cutler et al. (1990), De Long et al. (1990), Hong and Stein (1999), Barberis et al. (2015, 2018), past prices influence investors' trading decisions in our model. However, prior research abstracted from the wealth effect of investors by assuming that investors either have CARA preferences or follow exogenously specified trading strategies. These assumptions lead to the linear dynamics for the price which is too restrictive to fit these models successfully to the data. For example, these models have constant volatility, which is strongly rejected in the data. In contrast, investors in our model have CRRA preferences. We show that the interaction of extrapolative beliefs and wealth effect help explain many stylized facts associated with bubble periods in a single parsimonious model.

Our paper is also linked to literature that studies the role of extrapolative beliefs in asset prices. Choi and Mertens (2013), Jin and Sui (2019), Hirshleifer et al. (2015), Li and Liu (2019) study Lucas type economies and show that extrapolative beliefs help explain excess volatility and high equity premium in the stock market. In contrast, we focus on the price dynamics during bubble episodes and apply our framework to cryptocurrency markets.

Our paper is also linked to large literature that studies the interaction of differences in beliefs and trading volume, Harris and Raviv (1993), Kandel and Pearson (1995), Scheinkman and Xiong (2003). The main source of heterogeneity in our model is how strongly investors update their beliefs in response to past prices. We show that the resulting joint volatility and trading volume dynamics are consistent with that observed in cryptocurrency markets.

Finally, our paper contributes to the emerging literature on cryptocurrencies. Similar to our paper, Biais et al. (2019) develop and estimate an equilibrium model of bitcoin pricing. In their model, investors are risk-neutral and trade off the transactional benefits against the risk of hacking. Cong et al. (2020), Pagnotta and Buraschi (2018), Pagnotta (2020), and Sockin and Xiong (2020) link cryptocurrency prices to the degree of bitcoin adoption and network properties of bitcoin users. Different from these papers, we do not assume that cryptocurrencies have any intrinsic value. The demand for a cryptocurrency in our model is speculative and comes from boundedly rational investors who form their expectations based on the history of past prices.

The rest of the paper is organized as follows. We describe the model in the next section. Section 2 provides analysis of the model. Section 3 applies the model to cryptocurrency markets. In Section 4, we discuss some of our assumptions and extend the model to allow investors to have heterogeneous priors. Section 5 provides our conclusions.

1. The Model

The model studies equilibrium of an asset market, which is small in size compared with the size of the aggregate economy. It takes the interest rate as exogenous and assumes that all market participants can borrow and lend their capital at a constant risk-free rate r.

The financial market consists of one infinitely lived asset, which is in constant supply normalized to one share. The asset does not pay any dividend. The demand for the asset is speculative and comes from a mass one of boundedly rational utility maximizing investors. Each investor i maximizes her utility over consumption,

$$\max_{\pi_{it}} \quad E_0 \left[\int_0^\infty e^{-\rho t} \log(c_{it}) dt \right].$$

The assumption of logarithmic utility is for simplicity.

The economy evolves in continuous time. Denote the price of the asset at time t by P_t, the wealth of investor i by W_{it}, and the proportion of investor i's wealth invested in the asset by π_{it}. Each investor i faces the budget constraint

$$\frac{dW_{it}}{W_{it}} = \left((1-\pi_{it})r - \frac{c_{it}}{W_{it}} \right) dt + \pi_{it}\frac{dP_t}{P_t} + \mu dt + \sigma dB_t + \beta_i dB_{it}, \tag{1}$$

where B_t and B_{it}, $t \geq 0$ are one-dimensional standard Brownian motions. Brownian motions B_{it}, $i \in [0,1]$ are mutually independent and independent from B_t. The process

$$\mu dt + \sigma dB_t + \beta_i dB_{it}$$

represents random fluctuations in investor's i wealth which occur because of unmodeled labor shocks or unexpected returns on other investments, and are the main source of uncertainty in the model.

We assume that B_t is defined on a complete probability space $(\Omega, \mathcal{F}, P)$. We also assume that the exact law of large numbers holds in this economy and focus on equilibria where the asset price evolve according to an Ito process:

$$\frac{dP_t}{P_t} = \mu_t^P dt + \sigma_t^P dB_t. \tag{2}$$

Without loss of generality we can rewrite equation (2) as

$$\frac{dP_t}{P_t} = (r + \sigma_t^P h_t)dt + \sigma_t^P dB_t, \tag{3}$$

where

$$h_t = \frac{\mu_t^P - r}{\sigma_t^P} \tag{4}$$

is the asset's Sharpe ratio.

1.1. Bounded Rationality

Investors' knowledge about the state of the world is affected by behavioral limitations. They observe the price process but do not understand how volatility and Sharpe ratio are determined in equilibrium. Instead, they believe that the current Sharpe ratio follows an Ornstein-Uhlenbeck process

$$dh_t = \gamma(\bar{h} - h_t)dt + \hat{\sigma} d\hat{B}_t \tag{5}$$

and use the history of normalized returns

$$dz_t = \frac{1}{\sigma_t^P}\left(\frac{dP_t}{P_t} - rdt\right) = h_t dt + dB_t \tag{6}$$

to infer them. The parameter γ represents the speed of mean reversion and $\bar{h}$ is the long-run mean. Investors assume that Brownian motions B_t and $\hat{B}_t$ have constant correlation. Denote investors' estimate of the Sharpe ratio by $\theta_t = E_t[h_t|\mathcal{F}_t]$. Lemma 1 describes the evolution of investor beliefs θ_t.

Lemma 1: *Suppose investors believe that the dynamics of the Sharpe ratio follows (5) and observe the history of normalized returns (6). Then investors' estimates of the Sharpe ratio evolves according to*

$$d\theta_t = \big(\gamma(\bar{h} - \theta_t) + k(h_t - \theta_t)\big)dt + kdB_t, \tag{7}$$

where k is given in the Appendix.

Proof. See the Appendix.

The dynamics of investors' beliefs (7) can be written as the sum of predictable and unpredictable components:

$$d\theta_t = \underbrace{\gamma(\bar{h} - \theta_t)dt}_{\text{predictable change}} + k \underbrace{(dz_t - E_t(dz_t))}_{\text{unpredictable change}}. \tag{8}$$

When investors observe unexpectedly high returns, they infer that the asset must currently have a high Sharpe ratio and update their estimate of the Sharpe ratio upwards. Inspecting (8) we can see that investors' beliefs are a function of past return innovations. A fully rational investor would also use the history of past returns to make predictions but would use the correct functional form for the dynamics of beliefs.

As suggested in the Introduction, the knowledge of the correct functional form for the dynamics of beliefs might be an unrealistically strong assumption in the case of bubbles. In practice, this knowledge can be either an outcome of learning or a deep understanding of equilibrium. But the majority of bubbles are short-lived, with investors displaying notoriously irrational behavior. It seems natural to assume that investors trading in such complicated markets do not have perfect foresight of future prices. Therefore, they use some simplified models to describe the evolution of prices.

The dynamics (8) can be viewed as a continuous-time analog of the VAR method for discrete systems. It is commonly viewed as a first-order approximation of many economies and markets, and for this reason, is widely used in economics and finance.

Specification (8) can also fit many forms of technical analysis, which is popular among retail investors, and cryptocurrency investors in particular.

In specification (5), we assume that investors form their beliefs about the Sharpe ratio. Since volatility is observable, there is a one-to-one map between beliefs about the Sharpe ratio and beliefs about the expected returns. The specification where investors learn about the expected returns following an Ornstein-Uhlenbeck process leads to a qualitatively similar but less tractable dynamics.

Finally, we assume that investors neglect the correlation of investors' wealth shocks and asset returns when making their portfolio decisions. Our results are not driven by this assumption, but this assumption simplifies the analysis. It is consistent with the notion that investors have bounded rationality. The correlation of investors' wealth shocks and the asset returns at the individual level is likely to be low: $\beta_i \gg \sigma$, and might be unstable over time. Thus, as a first-order approximation, investors may disregard this correlation in their decision-making process.

1.2. Portfolio Choice

Investors' estimate of the Sharpe ratio determines their portfolio allocation choices. By the well-known property of the logarithmic preferences, Merton (1973), the optimal portfolio of investors is myopic and takes a simple form:

$$\pi_{it} = \pi_t = \frac{E_t[h_t|\mathcal{F}_t]}{\sigma_t^P} = \frac{\theta_t}{\sigma_t^P}. \tag{9}$$

Consumption is a constant fraction ρ of wealth, $c_{it} = \rho W_{it}$. Therefore, investor i's wealth evolves according to

$$\frac{dW_{it}}{W_{it}} = (r + \mu - \rho)dt + \theta_t h_t dt + (\theta_t + \sigma)dB_t + \beta_i dB_{it}. \tag{10}$$

Standard aggregation results imply that the aggregate investors' wealth obeys the following dynamics:

$$\frac{dW_t}{W_t} = (r + \mu - \rho)dt + \theta_t h_t dt + (\theta_t + \sigma)dB_t. \tag{11}$$

2. The Equilibrium

An equilibrium is a set of asset prices, aggregate wealth, and investors' portfolio choices $\{P_t, W_t, \pi_t\}$ that clears the asset market:

$$P_t = \pi_t \cdot W_t. \tag{12}$$

In what follows, we show that investors' beliefs θ_t is the only state variable and characterize the equilibrium in terms of θ_t.

The market clearing (12) and Ito's lemma imply that return volatility σ_t^P is the sum of two components:

$$\sigma_t^P = \underbrace{k\frac{\pi'(\theta_t)}{\pi(\theta_t)}}_{\text{changes in portfolio } \pi_t} + \underbrace{(\theta_t + \sigma)}_{\text{changes in wealth } W_t}. \tag{13}$$

The percentage change in the price is the sum of percentage changes in the portfolio policy and aggregate wealth. From equation (11), the percentage change in investors' total wealth at time t is equal to $(\theta_t + \sigma)dB_t$. The percentage change in investors' portfolio choice is equal to $\pi'(\theta_t)/\pi(\theta_t) \cdot kdB_t$.

From equation (9), return volatility σ_t^P can also be expressed as the ratio of θ_t and $\pi(\theta_t)$. Substituting it into (13), we arrive at an ordinary differential equation that describes the evolution of investors' portfolio choice as a function of their beliefs:

$$\theta_t = k\pi'(\theta_t) + (\theta_t + \sigma)\pi(\theta_t). \tag{14}$$

Equation (14) contains only parameters k and σ. Volatility and portfolio choice do not directly depend on the belief mean reversion parameters γ and $\bar{h}$ because they affect only predictable changes in the dynamics of investors' beliefs (7). Proposition 1 reports the solution to this equation and characterizes the resulting equilibrium.

Proposition 1: *There is a unique equilibrium, where the price follows an Ito process and investors' beliefs evolve according to (7). In equilibrium, investors' portfolio choice π_t, the Sharpe ratio h_t, and volatility σ_t^P are functions of investors' estimate of the*

Sharpe ratio θ_t *and are given by*

$$\pi(\theta_t) = \frac{1}{k} e^{-\frac{(\theta_t+\sigma)^2}{2k}} \int_0^{\theta_t} x e^{\frac{(x+\sigma)^2}{2k}} dx, \tag{15}$$

$$h(\theta_t) = \frac{\mu-\rho}{\sigma} + \left(1 + \frac{\gamma}{k\sigma}(\bar{h}-\theta_t)\right) k \frac{\pi'(\theta_t)}{\pi(\theta_t)} + \frac{1}{2\sigma} k^2 \frac{\pi''(\theta_t)}{\pi(\theta_t)}, \tag{16}$$

$$\sigma^P(\theta_t) = \frac{\theta_t}{\pi(\theta_t)}. \tag{17}$$

The process θ_t *is stationary and recurrent on* $(0,\infty)$*, and has a density*

$$\psi(\theta) = C \exp\left\{ \frac{2}{k^2} \int_{\theta_0}^{\theta} \Big(\gamma(\bar{h}-x) + k(h(x)-x)\Big) dx \right\}, \quad \theta > 0. \tag{18}$$

Proof. See the Appendix.

2.1. Volatility

Figure 1 depicts volatility and investors' portfolio choice as a function of investors' estimate of the Sharpe ratio θ_t. The parameter values are set at $\sigma = 0.05$, $k = 0.05$. These values are in line with our empirical estimates in Section 3. But it should be stressed that qualitative behavior of the model is robust and does not depend on particular values of parameters.

Looking at Figure 1, we can see that consistent with intuition, investors' portfolio choice π_t is monotonically increasing in θ_t. A higher estimate of the Sharpe ratio makes investors allocate a larger share of their wealth to the speculative asset. We provide a formal prove of this fact in the Appendix. In contrast, volatility is a U–shape function of θ_t. Volatility goes to infinity as θ_t approaches zero or infinity.

To understand these results, notice that volatility decomposition (13) shows that return volatility always exceeds volatility of the aggregate wealth $\theta_t+\sigma$, which is linearly increasing in θ_t and becomes the dominant component of return volatility when θ_t is large. This linear growth stabilizes the portfolio choice π_t, which is equal to the ratio of θ_t and return volatility σ_t^P. As θ_t approaches infinity, π_t tends to one and its percentage changes tend to zero.

Equation (13) also indicates that for $\theta_t \geq 0$ return volatility is bounded below by the volatility of wealth shocks σ. Therefore when θ_t tends to zero, $\pi(\theta_t)$ also tends to zero, and so does the asset price. In this situation, the percentage change in π_t, amplified by the feedback effect, exceeds the percentage change in W_t, and becomes the dominant component of return volatility. The feedback effect arises because investors use past returns to update their beliefs about the Sharpe ratio. A positive systematic shock to

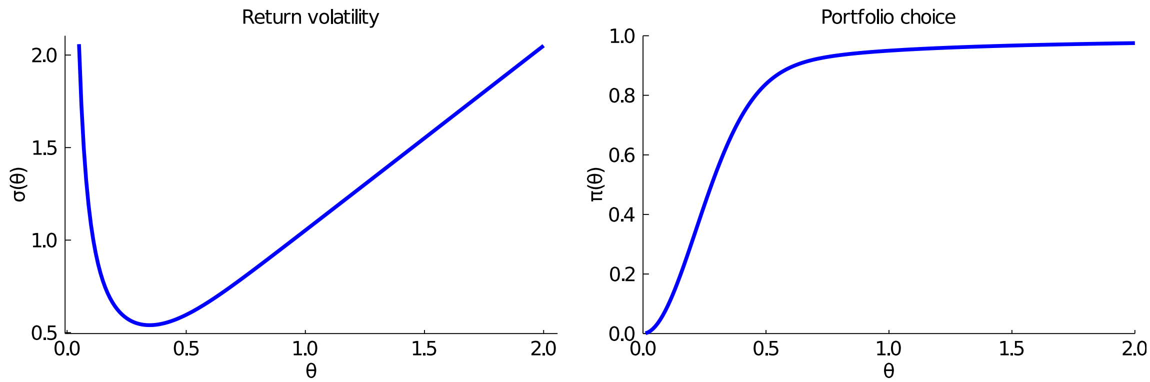

Figure 1: Portfolio choice and return volatility

The left and right panel depict return volatility σ_t^P and the fraction of wealth invested π_t, respectively, as functions of investors' estimate of the Sharpe ratio θ_t. The parameter values are set as follows: $\sigma = 0.05$, $k = 0.05$. Return volatility and portfolio choice do not depend on parameters γ and $\bar{h}$.

their wealth leads to a price increase, which in turn, leads to a more rosy outlook for the speculative asset. As a consequence, investors allocate a large fraction of wealth to the asset, which leads to a further increase in the price. As θ_t approaches zero, the feedback effect gets stronger. As a result, volatility goes to infinity.

For a given level of aggregate wealth, the price P_t is increasing in investors' estimate of the Sharpe ratio θ_t. Thus, our model predicts a novel pattern of volatility that differs from that observed in traditional financial markets. In traditional markets, volatility is strongly negatively correlated with prices: volatility is high when prices are low and vice versa. In contrast, our model predicts that the volatility of assets affected by strong speculative demand should be high when investors are either very pessimistic or optimistic about the asset.

2.2. Sharpe Ratio

A higher estimate of the Sharpe ratio, θ_t, induces investors to allocate a larger fraction of their wealth to the speculative asset and, therefore, leads to a higher price of the asset. The actual Sharpe ratio, h_t, measures the expected change in the asset price and therefore is determined by the dynamics of θ_t, which in turn, depends on h_t through investors' belief updates (7). The equilibrium is determined by this fixed point problem of the Sharpe ratio and its estimate.

Figure 2 plots the Sharpe ratio and risk premium. For large values of θ_t, the Sharpe ratio goes to $(\mu - \rho)/\sigma$. For small values of θ_t, the Sharpe ratio is positive and tends to infinity as θ_t approaches zero. The risk premium exhibits a similar pattern.

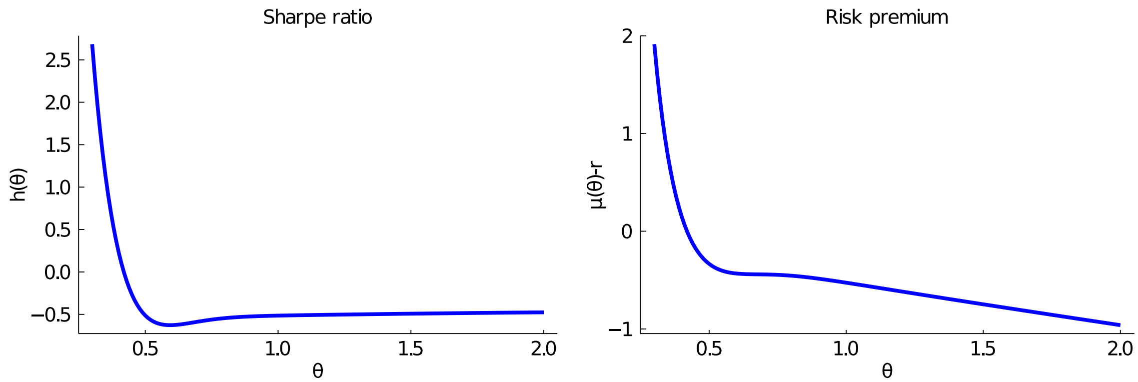

Figure 2: Sharpe ratio and risk premium

The left and right panel depict the Sharpe ratio h_t and risk premium $\mu_t^P - r$, respectively, as functions of investors' estimate of Sharpe ratio θ_t . The parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, $k = 0.05$, $\gamma = 0.2$, $\bar{h} = 0.5$.

To understand these results, notice that the Sharpe ratio and risk premium are high when investors are expected to increase the amount of their wealth invested in the speculative asset. When θ_t is high, investors allocate a large share of their wealth to the risky asset, so the possibility that even more wealth is attracted to the asset is small. Unless the expected rate of wealth shocks μ exceeds the consumption rate ρ the risk premium and Sharpe ratio are negative. In contrast, when θ_t is low, only a small fraction of investors' wealth is invested in the asset, so there is ample room for further price appreciation. As θ_t goes to zero and the feedback affect gets stronger, the Sharpe ratio and risk premium go to infinity.

Keeping other parameters constant, the Sharpe ratio and risk premium are increasing in belief mean reversion parameter γ when θ_t is below $\bar{h}$, and, and are decreasing in γ when θ_t is above $\bar{h}$. When θ_t is below $\bar{h}$, investors are expected to increase their investment in the speculative asset, which has a positive effect on the price, and vice versa, when θ_t is below $\bar{h}$, investors are expected to decrease their investment in the speculative asset, which has a negative impact on the price.

In general, the Sharpe ratio and investors' estimate of the Sharpe ratio are inversely related. While investors' estimate of the Sharpe ratio is correct on average (if $\gamma = 0$ or $\bar{h} = Eh_t$), they make systematic mistakes. When investors are least optimistic the Sharpe ratio is positive and high, and when they are most optimistic it is negative. This negative link stabilizes the dynamics of investors' beliefs and leads to the stationary dynamics of θ_t on $(0, \infty)$.

2.3. Steady State

In steady state, the Sharpe ratio follows a mean-reverting process. When the Sharpe ratio is high, investors' expectation of the Sharpe ratio θ_t is low, and vice versa. The resulting mean-reverting dynamics, however, does not follow an Ornstein-Uhlenbeck process assumed by investors. The discrepancy between the assumed and true dynamics of the Sharpe ratio is the reason for investors' systematic learning mistakes.

Because the Sharpe ratio goes to infinity sufficiently quickly as θ_t goes to zero, if investors' start with strictly positive beliefs about the Sharpe ratio their beliefs stay away from zero for all time. As a result, investors never take a short position in the speculative asset, ensuring that the price is strictly positive at all times. This inability to reach 0 at a positive time is similar to the behavior of the square-root process often used to model the dynamics of interest rates and volatility. We would like to stress, however, that in our model the fact that the dynamics of investors' beliefs is restricted to positive values is not a consequence of a hard-wired assumption, but an equilibrium outcome.

The left panel of Figure 3 plots the stationary distribution $\psi(\theta)$ for the case $\gamma = 0$, in which investors believe that the Sharpe ratio follows a random walk, and $\gamma = 0.25$ and $\bar{h} = 0.5$, in which they believe that the Sharpe ratio follows a mean-reverting process. Mean reversion in investors' beliefs shifts the stationary distribution towards $\bar{h}$, and when γ goes to infinity, the stationary distribution converges to a degenerate distribution. The middle and right panel Figure 3 show volatility and the Sharpe ratio restricted to values of θ_t in the interquantile range $[0.001, 0.999]$. While both volatility and the Sharpe ratio can take arbitrary large values, in practice, the probability of reaching extreme values is very small.

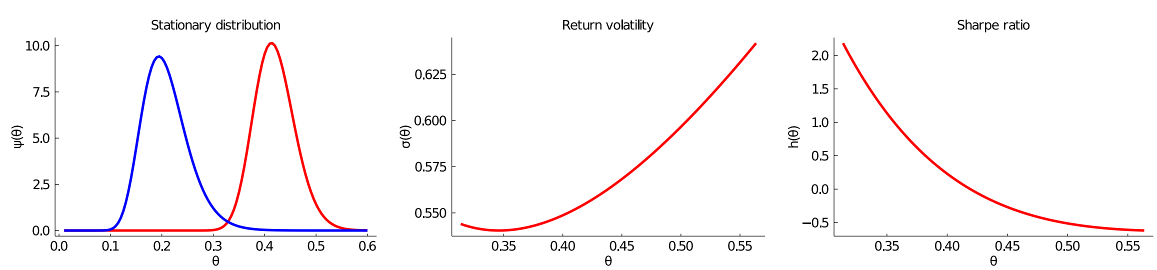

Figure 3: Stationary distribution of θ_t

The left panel depicts the stationary distribution of investors' estimate of Sharpe ratio θ_t. The blue and red line correspond to the case of $\gamma = 0$ and $\gamma = 0.2, \bar{h} = 0.5$. Other parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, $k = 0.05$. The middle and right panel show volatility and the Sharpe ratio restricted to values of θ_t in the interquantile range $[0.001, 0.999]$.

Proposition 2 reports the expected time for the process θ_t to first arrive at $\underline{\theta}$ or $\overline{\theta}$ from its initial value θ_0.

Proposition 2: *Suppose the process θ_t starts at $\theta_0 > 0$ at $t = 0$. The expected time it first hits $\underline{\theta} < \theta_0$ is given by:*

$$ET_{\underline{\theta}} = \frac{2}{k^2} \int_{\underline{\theta}}^{\theta_0} \frac{1 - \Psi(\theta)}{\psi(\theta)} d\theta \tag{19}$$

The expected time it first hits $\overline{\theta} > \theta_0$ is given by:

$$ET_{\overline{\theta}} = \frac{2}{k^2} \int_{\theta_0}^{\overline{\theta}} \frac{\Psi(\theta)}{\psi(\theta)} d\theta \tag{20}$$

where $\Psi(\theta) = \int_0^\theta \psi(y)dy$ is the CDF of the stationary distribution.

Proof. See the Appendix.

The left panel of Figure 4 plots the expected time for θ_t to reach a particular value θ when it starts from its median value. The expected time quickly increases and goes to infinity as θ moves away from the median value. For example, it would take on average 100 years for θ_t to reach the value of $\theta = 0.3$ that corresponds to the Sharpe ratio value of 2.7. The right panel of Figure 4 plots the expected time for θ_t to reach the median value when it starts from a particular initial value θ_0. We can see that it would take on average 2 years for θ_t to reach the median if it starts from $\theta_0 = 0.3$. Taken together, Panel (a) and (b) provide information about the possible duration of bubble cycles.

The knowledge of the stationary distribution (18) allows us to compute steady-state empirical moments of returns. Figure 5 shows the mean excess return, mean excess log-return, and mean Sharpe ratio. Provided that $\mu < \rho$ the mean excess log-return is negative. Therefore, the price is expected to decline over time. This is expected since the asset is not productive, and the expected growth of wealth is below the consumption rate.

In contrast to the mean excess log-return, the mean excess return and mean Sharpe ratio are positive. The sign difference between the mean log and simple returns might seem surprising at first glance. While θ_t spends a roughly equal amount of time in the positive and negative expected return area, the risk premium and Sharpe ratio are highly skewed and negatively correlated with investors' beliefs. When investors are pessimistic and invest a relatively low fraction of their wealth in the speculative asset, the expected returns are high and positive. But when investors are optimistic and invest a relatively large share of their wealth, the expected returns become negative.

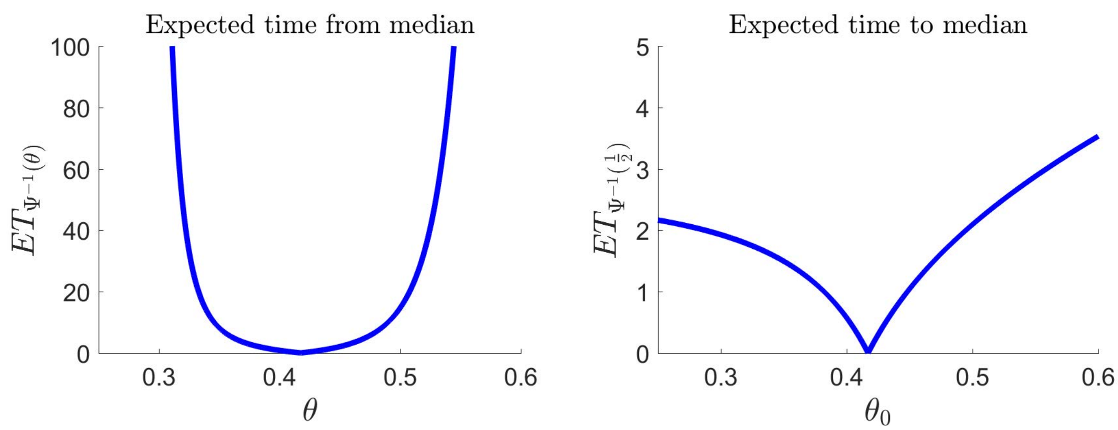

Figure 4: Expected first-hitting-time of θ_t

The left panel plots the expected time for θ_t to reach a value θ when it starts from its median value. The right panel plots the expected time for θ_t to reach the median value when it starts from an initial value θ_0. The parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, $k = 0.05$, $\gamma = 0.2$, and $\bar{h} = 0.5$.

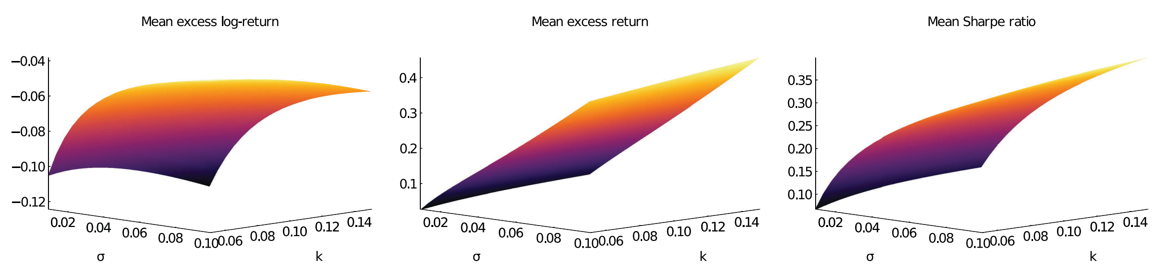

Figure 5: Empirical moments

The left, middle, and right panel depict the mean excess return, mean excess log-return, and mean Sharpe ratio under the stationary distribution (18). The parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\gamma = 0.2$, and $\bar{h} = 0.5$.

Since the logarithm is a concave function, it penalizes negative returns more than rewards positive returns. As a result, when returns are highly skewed it is possible to have negative expected log-returns but positive expected simple returns.

2.4. Arbitrage Opportunities

Because the model has only boundedly rational investors, the market can potentially be very inefficient with many arbitrage opportunities present. If this is indeed the case and it is easy to take advantage of irrational investors, one might question the plausibility of the model. In this section, we study return predictability and show that inefficiencies are generally small.

We start with computing the correlation of returns at different horizons. The closed-form solutions for correlations are not available. Therefore, we present the results based on simulations. Simulation of returns is straightforward and fast since both volatility and the Sharpe ratio are available in closed form. We simulate 1000 sample paths of 10,000 years of daily returns. In simulations, we sample the starting value of θ_t from the stationary distribution and set the same parameter values as before: $\rho = 0.02, \mu = 0, \sigma = 0.05, k = 0.05, \gamma = 0.2, \bar{h} = 0.5$. With these parameter values, we obtain that the autocorrelations of returns at daily, weekly, and monthly frequency are -0.1%, -0.6%, and -2.7%, correspondingly. The returns tend to exhibit mean-reversion but the magnitudes are small. Therefore, strategies that do not depend on the knowledge of the state of the economy, θ_t, generate negligible profit.

Next, we examine and compare the average long-run profit of investors who have some knowledge of the economy with that of feedback investors. We consider two cases. First, we look at the profit of a rational investor who knows θ_t and its true dynamics. Second, we look at the profit of an investor who does not know the Sharpe ratio and needs to estimate it using the Kalman filter (7), but who can choose the optimal parameter values.

We assume that investors do not have wealth shocks ($\sigma = \beta_i = 0$) and choose their portfolio policies to maximize the expected logarithmic return on their wealth. Let

$$logR^i = \lim_{T\to\infty} \frac{1}{T} E \log(W_T^i / W_0^i), \quad i \in \{r, f, o\}, \tag{21}$$

be the average logarithmic return of a rational (r), feedback (f), and feedback investor who chooses the optimal parameter values (o), and let π_t^i, $i \in \{r, f, o\}$, be their respective optimal portfolio strategies. Proposition (3) reports the results.

Proposition 3: *The optimal portfolio policies and expected logarithmic return on the wealth of a rational, feedback, feedback investor who chooses the optimal parameter values are given by the following expressions:*

$$\pi_t^r = \frac{h(\theta_t)}{\sigma(\theta_t)}, \qquad logR^r = r + \mu - \rho + \frac{1}{2} E[h^2(\theta_t)], \tag{22}$$

$$\pi_t^o = \frac{E[h(\theta_t)]}{\sigma(\theta_t)}, \qquad logR^o = r + \mu - \rho + \frac{1}{2} (E[h(\theta_t)])^2, \tag{23}$$

$$\pi_t^f = \frac{\theta_t}{\sigma(\theta_t)}, \qquad logR^f = r + \mu - \rho + \frac{k + 2\gamma}{2k} E[\theta_t^2] - \frac{\gamma}{k} \bar{h} E[\theta_t] - \frac{k}{2}. \tag{24}$$

where the expectation is taken with respect to the stationary measure (18).

Proof. See the Appendix.

For any investor, the optimal portfolio policy to maximize the expected log-return is to invest a share of wealth in the speculative asset equal to the ratio of her estimate of the Sharpe ratio and volatility. Since a rational investor knows the link between equilibrium quantities and θ_t and observes return volatility, which is a function of θ_t, she can deduce θ_t from the knowledge of $\sigma^P(\theta_t)$. As a result, she can also deduce the current value of the Sharpe ratio, $h(\theta_t)$.

A feedback investor uses her estimate of the Sharpe ratio, θ_t. In the proof, we show if a feedback investor could choose parameters optimally, she would be better off if she fixes her estimate of the Sharpe ratio at its mean value and does not update it. The reason for this is that, as we have shown in (2), investors' estimate of the Sharpe ratio and the true Sharpe ratio are inversely related. By setting her estimate of the Sharpe ratio to its mean, an investor can minimize her average mistake. Notice that to implement this strategy, the investor would need to know the mean of the Sharpe ratio, which might not be available. Fixing the estimate of the Sharpe ratio away from its mean results in a lower expected return. For sufficiently large deviations, the resulting expecting return can be lower than that of the feedback investors who use past returns to update their estimate.

For parameter values $\rho = 0.02, \mu = 0, \sigma = 0.05, k = 0.05, \gamma = 0.2$, and $\bar{h} = 0.5$, we obtain the following values for the expected excess log-return: $logR^r - r = 8.64\%$, $logR^o - r = -1.51\%$, and $logR^f - r = -6.43\%$. Consistent with our previous results, the excess log-return of feedback investors is negative. An investor who knows the mean value of the Sharpe ratio correctly can achieve a better but still negative excess return. Finally, a fully rational investor can achieve an excess log-return of about 8.64%. The average volatility of the portfolio that delivers this return is 45%. Overall, our results show that opportunities to make trading profits in this market are limited.

3. Application to Cryptocurrency Markets

In this section, we apply our model to cryptocurrency markets. Consistent with the assumption of the model, cryptocurrencies do not offer any payoff to their investors. Further in support of model assumptions, Makarov and Schoar (2020) show that in cryptocurrency markets, there is a strong positive relation between investors flows and prices. The net order flow explains 80% of return variation — a substantially greater proportion than in traditional capital markets. Liu and Tsyvinski (2020) show that cryptocurrency returns do not significantly load on macroeconomic factors or cryptocurrency production factors but correlate with measures of investor sentiment. Bhambhwani et al. (2019) show that cryptocurrency returns load positively on the U.S. market factor.

In our empirical analysis, we consider the three most liquid and largest cryptocurrency markets: bitcoin, ethereum, and ripple, and particularly focus on bitcoin, which is the oldest and most prominent cryptocurrency. Our model makes a number of specific predictions about volatility. In particular, it predicts that volatility is a U-shape function of investors' beliefs θ_t. Holding the aggregate wealth constant, the price is an increasing function of θ_t. Therefore, our model predicts that volatility should be high at the height and the bottom of cryptocurrency price cycles.

Since in practice prices are observed only at discrete dates, volatility is not directly observable and should be estimated. To estimate volatility we use model-free daily realized volatility based on 5-minute returns (5-minute RV). This choice of estimator as a measure of realized volatility is motivated by Liu et al. (2015) who consider over 400 different estimators and find little evidence that 5-minute RV is outperformed by any other measures. To construct 5-min RV, we use tick-level transaction prices from Kaiko and Bitcoincharts.com, which obtain the data by querying APIs provided by exchanges.[2]

While cryptocurrencies are traded on many exchanges, cryptocurrency exchanges are nonintegrated. As a result, as shown by Makarov and Schoar (2020), cryptocurrency markets exhibit periods of large, recurrent arbitrage opportunities across exchanges. To avoid capturing price variation across exchanges, we therefore use price originating from a single exchange. The choice of the exchange is motivated by liquidity and data availability. For bitcoin, we use transaction data from Bitstamp in the period from 1/1/2014 to 31/12/2020; for ethereum we use data from Coinbase in the period from 1/1/2017 to 28/05/2020; finally, for ripple we use data from Bitstamp in the period from 1/09/2017 to 28/05/2020.

Cryptocurrency markets trade 24 hours a day, seven days a week. We take midnight to midnight Coordinated Universal Time (UTC) as a trading day. For each day, whenever there are several trades within the same second, we take the median price for this second. Next, we create a one-minute grid whereby we find the latest available price before each minute on this day. We use these 24×60 prices to compute five different 5-minute RV based on returns starting at 00.00, 00.01, ..., and 00.04 UTC. Each 5-minute RV is the sum of squared 5-min returns. Finally, we average these five "5-minute sub-sampled RV" to obtain the final estimate of the realized variance.

Figure 6 plots daily log price and realized volatility for the three cryptocurrency markets. The right upper panel shows the bitcoin price. The bitcoin price dynamics exhibits the well-documented and publicized cycles. There were large run-ups in the bitcoin price prior to 2014 and from 2015 to early 2018, followed by steep declines. As

[2]See Makarov and Schoar (2020) for more details about data description.

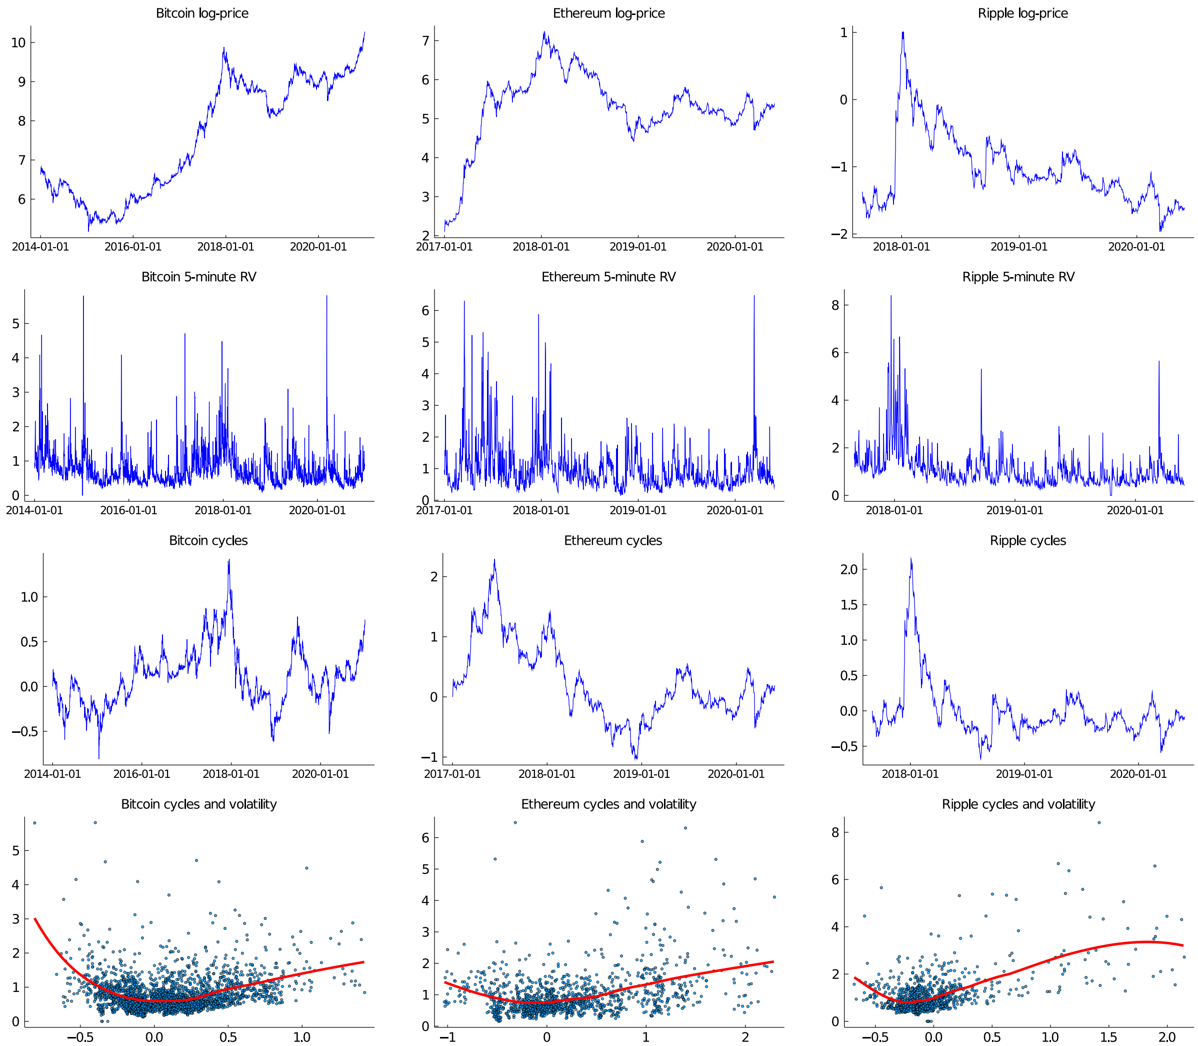

Figure 6: Bitcoin price and returns

The panels depict daily log price, 5-minute RV estimate, our estimate of cryptocurrency cycles, and scatter plots of daily realized volatility against our measure of crypocurrency cycles for bitcoin from 1/1/2014 to 31/12/2020, ethereum from 1/1/2017 to 28/05/2020, and ripple from 1/09/2017 to 28/05/2020.

of December 2020, the bitcoin price is again at all-time high. Ethereum and ripple prices follow similar dynamics.

The second panel of Figure 6 shows the 5-minute RV estimate. Eyeballing the behavior of realized volatility and prices in Figure 6 we can notice that consistent with the model's predictions, volatility is indeed generally higher both at the height and the bottom of cryptocurrency price cycles. To verify this behavior, we construct a simple measure of cryptocurrency cycles by taking the deviation of cryptocurrency price from its exponentially-weighted moving average with the weight parameter equal to 0.99. Notice that the dynamics of beliefs (7) imply that investors' beliefs θ_t can be written as an exponentially moving average of past returns. Therefore, one should expect our simple measure of cryptocurrency cycles to be correlated with θ_t. The results are shown in the third panel of Figure 6. The deviations are highest during bitcoin run-ups in 2017, 2019, and the end of 2020, and lowest during bitcoin crashes in 2015 and late 2018.

The last panel of Figure 6 shows a scatter plot of daily realized volatility against our measure of cryptocurrency cycles along with the LOESS (locally estimated scatterplot smoothing) curve. We can see that the LOESS curve is a U-shaped function, which confirms our observation that volatility is higher at the height and the bottom of cryptocurrency price cycles. Thus, volatility in cryptocurrency markets displays a different pattern from that in traditional markets, where rising asset prices are accompanied by declining volatility, and vice versa (the "leverage effect"). We conjecture that our model's prediction about the U-shaped behavior of volatility is not limited to cryptocurrency markets but holds more generally for all assets affected by strong speculative demands.

Next, we estimate and fit the model to data. In this exercise, we restrict our attention to the bitcoin market. Our goal here is to see if the model can match the dynamics of prices and moments of returns. Denote the time-t log bitcoin price by p_t. The joint process (p_t, θ_t) is Markov. Its evolution is described by the following system of stochastic differential equations

$$
\begin{aligned}
dp_t &= \left(r + \sigma(\theta_t)h_t(\theta_t) - \frac{1}{2}\sigma(\theta_t)^2\right)dt + \sigma(\theta_t)dB_t, \\
d\theta_t &= \big(\gamma(\bar{h} - \theta_t) + k(h_t(\theta_t) - \theta_t)\big)dt + kdB_t,
\end{aligned} \tag{25}
$$

where functions $h_t(\theta_t)$ and $\sigma(\theta_t)$ are given by equations (15) and (17). Of the two processes, p_t and θ_t, only the price is observed.

Since the closed-form solution for the transition density is not available one cannot apply the maximum likelihood method directly. The literature proposed several methods to approximate the transition density. These include simulated maximum likelihood estimation, Pedersen (1995), Brandt and Santa-Clara (2002); Markov chain Monte Carlo, Eraker (2001), Elerian et al. (2001); and using Hermite function with coefficients obtained using higher order Ito-Taylor expansions, Ait-Sahalia (2002). In this paper, we take advantage of high frequency data and use the Euler scheme to estimate the unknown parameters in (25) and infer the unobserved states θ_t and s_t.

The Euler scheme uses a first-order approximation given by the discrete-time process

$$p_{(i+1)\Delta} = p_{i\Delta} + \left(r + \sigma(\theta_{i\Delta})h_t(\theta_{i\Delta}) - \frac{1}{2}\sigma(\theta_{i\Delta})^2 \right)\Delta + \sigma(\theta_{i\Delta})\sqrt{\Delta}\varepsilon_i,$$
$$\theta_{(i+1)\Delta} = \theta_{i\Delta} + \left(\gamma\bar{h} - (\gamma + k)\theta_{i\Delta}\right)\Delta + k\left(p_{(i+1)\Delta} - p_{i\Delta} - r\Delta\right)/\sigma(\theta_{i\Delta}), \qquad (26)$$

where ε_i, $i = 1, \ldots, n$ are independent, identically $N(0,1)$-distributed. We expect this approximation to work sufficiently well at the daily frequency (Florens-Zmirou (1989)).

The parameters ρ and μ cannot be identified separately in the data. Therefore, set $\rho = 2\%$ and $r = 1\%$, and use the maximum likelihood estimator for the discretized system (26) to estimate other parameter values. We obtain the following estimates: $\mu = 5.4\%$, $\sigma = 0.012$, $k = 0.08$, $\gamma = 0.19$, $\bar{h} = 0.55$, and $\theta_0 = 0.73$. The high estimate of μ is the consequence of large bitcoin appreciation over the considered period, and is unlikely to hold in the long-run.[3]

The imputed unobserved investors' beliefs θ_t are plotted in Figure 7. The system starts at the state where the perceived Sharpe ratio θ_t is high. This predicts lower expected returns and leads to a sharp price decline. The early 2018 and late 2020 peaks of bitcoin prices realize again when θ_t reaches its highest values.

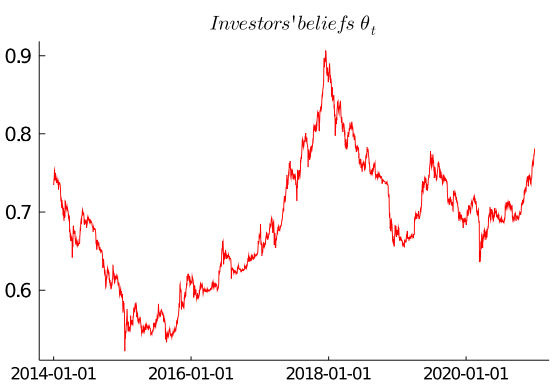

Figure 7: Investors' beliefs
The figure depicts the evolution of investors' beliefs θ_t obtained by fitting the model to the bitcoin data using the maximum likelihood estimator for the Euler scheme (26).

In the model, the expected return and volatility are functions of investors' beliefs, θ_t. Having obtained the parameter estimates and investors' beliefs, we can check these predictions by regressing bitcoin realized returns and volatility on their model predicted values. Regressions (27) and (28) report the results. The t-stats are given in

[3]See our discussion in Section 4.

parentheses.

$$r_t = \underset{(-1.0)}{-0.002} + \underset{(2.07)}{2.08}\hat{r}(\theta_{t-1}) + \varepsilon \tag{27}$$

$$\sigma_{rv,t} = \underset{(-6.5)}{-0.86} + \underset{(12.3)}{2.18}\hat{\sigma}(\theta_{t-1}) + \varepsilon. \tag{28}$$

If the model is correct, one should expect the coefficients at the predicted values from the model to be equal to one. Looking at the regression results, we can see that the coefficients are indeed reasonably close to one and significant at the 5% level. Overall, regressions (27) and (28) confirm that the model is able to fit the data.

Next, we show that the model can also generate time-series momentum. Liu and Tsyvinski (2020) document that there is time-series momentum in cryptocurrency markets at horizons from one week to four weeks. In the model, expected returns are positively autocorrelated because they are a function of persistent state-variables, θ_t, and s_t. The realized returns and expected returns tend to be negatively correlated because the true Sharpe ratio and investors' estimate of the Sharpe ratio are inversely related. The momentum profit is determined by the interplay of these two effects.

When the risk premium is high or low, the realized returns also tend to be high or low and continue to be high or low for some time in the future, leading to momentum in returns. When the risk-premium is around zero, the market moves from slight positive expected returns to slight negative expected returns following positive realized returns and vice versa. This creates reversals in returns. In steady state, the second effect dominates. As a result, returns are slightly negatively correlated. But on the way to the steady-state returns exhibit momentum.

To demonstrate this, we simulate 1,000,000 samples of weekly returns over the six-year horizon using the dynamics (26) and parameter values we obtained from fitting the model to the bitcoin market. As initial values for θ_t we set $\theta_0 = 0.3$. Similar to Liu and Tsyvinski (2020), we regress cumulative $i-$week return $R_{t,t+i}$ on weekly return R_t, $i = 1, 2, 3,$ and 4. The average coefficients along with the coefficients observed in the data are reported in the table below.

weekly	$R_{t,t+1}$	$R_{t,t+2}$	$R_{t,t+3}$	$R_{t,t+4}$
model: R_t	0.03	0.06	0.9	0.12
data: R_t	0.04 (0.70)	0.13 (1.71)	0.20 (1.96)	0.10 (0.77)

Finally, we note that the model predictions are also consistent with the large price impact observed in the bitcoin market. Makarov and Schoar (2020) estimate that a buy order of 1,000 bitcoin leads to a 30bs increase in the bitcoin price. The total amount of active bitcoins during the estimation period was about 16.5M, which translates into a price impact of 50. In the model, price impact is equal to the ratio σ_t^P/σ, and is about 50 for the fitted parameter values.

4. Discussion and extensions

In the current model, we have made a number of simplifying assumptions. In particular, we have assumed (1) that investors have the same expectations of the Sharpe ratio, and (2) that all investors participate in the market from date 0. As a consequence of these assumptions, there are two dimensions where our model is at odds with the data. First, as in any setting with homogeneous agents, trading volume is zero in our model. Second, the bubble growth in the model realizes through the dynamics of investors' beliefs and the mean of wealth shocks.

In practice, a large portion of bubble growth occurs when investors, previously unaware of a speculative asset, jump on the bandwagon and invest in it. As a result, the total aggregate wealth of investors who invest in a speculative asset grows at a high rate in the early stages of the bubble. Over time, as more and more investors become aware of the speculative asset, the growth of "new" money slows down, and the growth rate of total aggregate wealth converges to its "natural" rate determined by inflation and the growth rate of the economy. Thus, our estimate of μ in Section 3 is likely biased upwards.

One relatively straightforward way to make our model more realistic would be to assume that aggregate wealth shocks have not constant but stochastic drift μ_t, where innovations to the drift process are positively correlated with the wealth shocks. It can be shown that in this extended setup, the expressions (15) for the portfolio policy and volatility remain unchanged, and the Sharpe ratio (16) becomes

$$h(\theta_t, \mu_t) = \frac{\mu_t - \rho}{\sigma} + \left(1 + \frac{\gamma}{k\sigma}(\bar{h} - \theta_t)\right) k \frac{\pi'(\theta_t)}{\pi(\theta_t)} + \frac{1}{2\sigma} k^2 \frac{\pi''(\theta_t)}{\pi(\theta_t)}. \tag{29}$$

As in the main setup, the drift process only changes the asset's expected return but not its volatility.

Another way to address the current shortcomings of the model is to allow investors to have heterogeneous beliefs and model the entry of investors using the dynamics of their beliefs. For example, one can imagine retail and institutional investors have different priors and different speeds with which they update their beliefs. Initially, retail investors are more optimistic, so they drive bubble growth in the early periods. Over time, institutional investors seeing large run-up in the price slowly update their preferences and enter the market, fueling bubble growth at the later stage. Developing a fully-fledged realistic model with heterogeneous investors is an interesting direction for future research but goes beyond the scope of this paper.

In the rest of this section, we consider a simple extension of our model to the case of heterogeneous beliefs. Despite its simplicity, we show that this model can successfully explain the joint qualitative behavior of prices and volume in the bitcoin market.

Suppose that in addition to the group of investors who update their beliefs according to equation (7), there is another group of investors who hold constant beliefs about the Sharpe ratio. Denote their beliefs by $\hat{\theta}$. Similar to the first group, investors in the second group have logarithmic utility, face the same budget constraint (1), and neglect the correlation of investors' wealth shocks and asset returns when making their portfolio decisions. Thus, the optimal portfolio policies of the feedback and constant-belief investors are

$$\pi_t = \frac{\theta_t}{\sigma_t^P}, \tag{30}$$

$$\widehat{\pi}_t = \frac{\hat{\theta}}{\sigma_t^P}, \tag{31}$$

respectively.

Denote the aggregate wealth of the feedback investors by W_t and that of the constant-belief investors by $\widehat{W}_t$. The market clearing condition (12) becomes

$$P_t = \pi_t W_t + \widehat{\pi}_t \widehat{W}_t. \tag{32}$$

With the two groups of investors present, their wealth ratio, $\widehat{W}_t/W_t$, becomes an additional state variable. Denote this ratio by X_t. Applying Ito's lemma to ratios, we obtain the dynamics of X_t:

$$\frac{dX_t}{X_t} = (\hat{\theta} - \theta_t)\left(h_t - \sigma - \theta_t\right) dt + (\hat{\theta} - \theta_t) dB_t, \tag{33}$$

where h_t is the Sharpe ratio, which is now a function of both θ_t and X_t. Proposition 4 shows that the extended model preserves the tractability of our base model.

Proposition 4: *There is a unique equilibrium, where the price follows an Ito process and the feedback investors' beliefs evolve according to (7). In equilibrium, all quantities are functions of the feedback investors' estimate of the Sharpe ratio θ_t and the wealth* *ratio X_t. In particular,*

$$\sigma^P(\theta_t, X_t) = \frac{\hat{\theta} X_t + \theta_t}{Z(\theta_t, X_t)}, \tag{34}$$

$$h_t(\theta_t, X_t) = \frac{\mu - \rho}{\sigma} + \left(k + \frac{\gamma}{\sigma}(\bar{h} - \theta_t)\right)\frac{Z_\theta}{Z} \tag{35}$$

$$+ \frac{1}{\sigma Z}\Big[\frac{1}{2}k^2 Z_{\theta\theta} + k(\hat{\theta} - \theta_t)X_t Z_{\theta X} + \frac{1}{2}(\hat{\theta} - \theta_t)^2 X_t^2 Z_{XX}\Big]. \tag{36}$$

where $Z(\theta_t, X_t)$ is given by

$$Z(\theta_t, X_t) = \frac{1}{k} e^{-\frac{(\theta+\sigma)^2}{2k}} \int_{\underline{\theta}(\theta_t, X_t)}^{\theta} (\hat{\theta} X_t + y) e^{\frac{(y+\sigma)^2}{2k}} dy, \tag{37}$$

and $\underline{\theta}(\theta_t, X_t)$ solves

$$e^{\frac{(\hat{\theta}-\underline{\theta})^2}{2k}} \underline{\theta} = -\hat{\theta} X_t e^{\frac{(\hat{\theta}-\theta_t)^2}{2k}}. \tag{38}$$

Proof. See the Appendix.

The main economic mechanism in the heterogeneous beliefs model is similar to that in the baseline model. However, the heterogeneous beliefs model displays considerably more complicated dynamics. The asset pricing moments are no longer just functions of investors' beliefs but depend on the wealth ratio as well. When X_t goes to zero, feedback investors dominate the economy, and the model solution converges to the solution in the main setup. When X_t goes to infinity, the economy is dominated by the constant-beliefs investors. In this economy, volatility and the Sharpe ratio are constant and are equal to $\sigma_t^P \equiv \hat{\theta} + \sigma$ and $(\mu - \rho)/\sigma$, respectively.

To illustrate the model dynamics consider the case where $\gamma = 0$, $\mu = 0$, and $\hat{\theta}$ is equal to the mean of the Sharpe ratio in the economy populated only by the feedback investors. We showed in Section 2.4 that in the economy populated only by the feedback investors, the portfolio strategy of the constant-beliefs investors would deliver a higher logarithmic return than any other feedback strategy. Thus, in partial equilibrium, the constant-beliefs investors eventually would dominate the feedback investors. In general equilibrium, however, as the relative wealth of the constant-beliefs investors grows, volatility and the Sharpe ratio converge to constant values. In this case, the portfolio strategy of the feedback investors, who update their beliefs, starts outperforming the portfolio strategy of the constant-beliefs investors. As a result, the economy converges to a steady state where no type of investors completely dominates the other in the long run.

Figure 8 plots volatility and the Sharpe ratio in the heterogeneous beliefs economy given by equations (34). For any fixed wealth ratio X, volatility exhibits the familiar U-shape pattern: It is largest for small and large values of θ_t. Volatility generally declines with X since the price needs to adjust more to induce the feedback investors to rebalance their portfolio. As in the main setup, the Sharpe ratio goes to infinity as θ_t and the price decline. Different from the main setup, where feedback investors never take a short position, now investors sometimes can short the asset. The price, however, stays always positive.

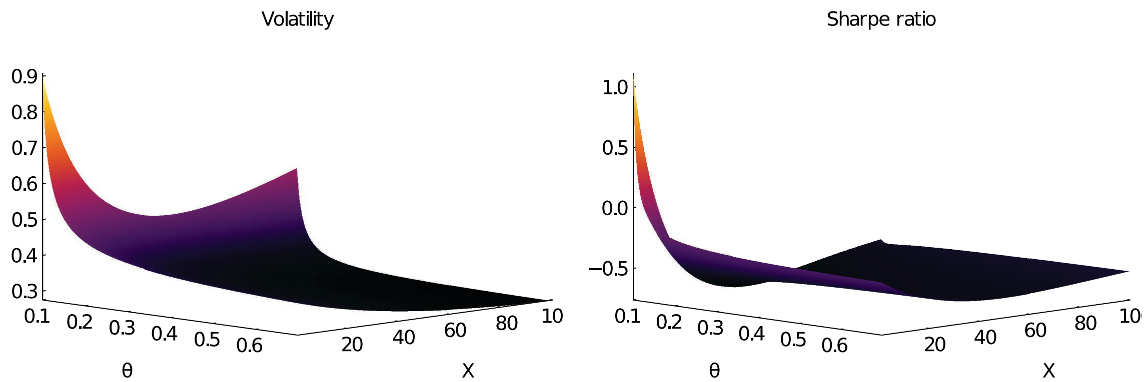

Figure 8: Volatility and Sharpe ratio

The left and right panel depict volatility and the Sharpe ratio, respectively, as functions of feedback investors' estimate of Sharpe ratio θ_t and wealth ratio X_t. The parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, $k = 0.05$, $\gamma = 0.0$, $\bar{\theta} = 0.21$.

Given portfolio policies of the investors (30), the fraction of the speculative asset held by the constant belief investors is equal to

$$\frac{\widehat{\pi}_t \widehat{W}_t}{\widehat{\pi}_t \widehat{W}_t + \pi_t W_t} = \frac{\hat{\theta} X_t}{\hat{\theta} X_t + \theta_t}. \tag{39}$$

Applying Ito's lemma to ratios, and using (7) and (33), we find that its dynamics is given by

$$d\left(\frac{\hat{\theta} X_t}{\hat{\theta} X_t + \theta_t}\right) = (\ldots)dt - \frac{\hat{\theta} X_t}{(\hat{\theta} X_t + \theta_t)^2} \cdot \Big(k + \theta_t(\theta_t - \hat{\theta})\Big)\, dB_t. \tag{40}$$

In the model, the feedback and constant-beliefs investors constantly trade with each other. Because we use Brownian motion to model uncertainty, trading volume in fact is infinite over any period. Therefore, to study volume, we follow Lo and Wang (2003) and use the coefficient at the dB_t in the above equation as a proxy for volume. Its plot is shown in the left panel of Figure 9.

Similar to volatility, the volume goes to infinity when $\hat{\theta} X_t + \theta_t$ and the price tend to zero. Also, volume tends to be high when θ_t is high. In general, volume tends to be high when volatility is high and vice versa. This can be proved using explicit expressions for volatility and volume derived above. To visualize this relationship, we simulate the model using the dynamics equations (7) and (33) for θ_t and X_t. In simulations, we set as before, $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, and $k = 0.05$, and simulate the model over 10,000 years to obtain the stationary distribution of θ_t and X_t. For each θ_t and X_t we compute volatility and volume. The scatter plot of log volatility versus log volume is shown in the right panel of Figure 9. The figure confirms a strong positive relationship between volume and volatility in the model.

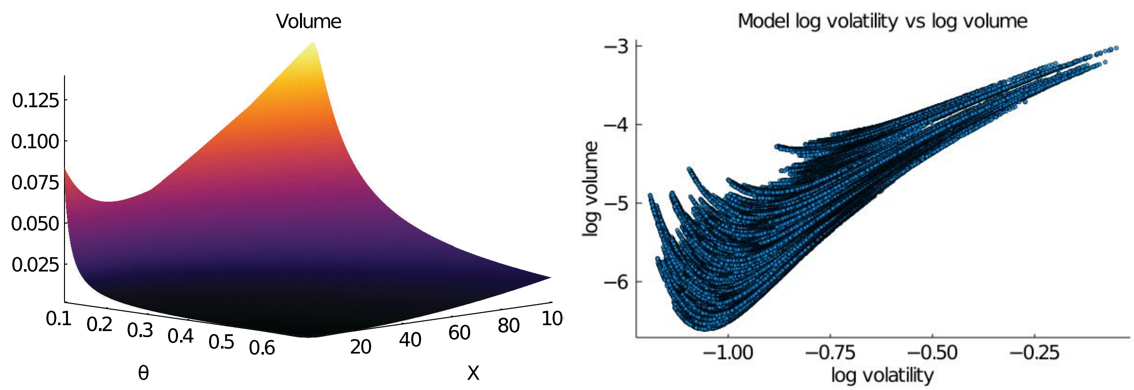

Figure 9: Volume

The left panel depicts volume as a function of feedback investors' estimate of the Sharpe ratio θ_t and wealth ratio X_t. The right panel shows the scatter plot of log volatility versus log volume from model simulations. The parameter values are set as follows: $\rho = 0.02$, $\mu = 0$, $\sigma = 0.05$, $k = 0.05$, $\gamma = 0.0$, $\bar{\theta} = 0.21$.

A strong positive relationship between volume and volatility has been documented in many markets; see Karpoff (1987) for a survey of empirical evidence. Next, we show that it also exists in the bitcoin market. To calculate bitcoin volume, we use bitcoin volume data from 17 largest and most liquid crypto exchanges provided by Kaiko. Our data cover the period from 1/1/2014 to 28/05/2020. Following Makarov and Schoar (2020), each day on each exchange we compute bitcoin volume to major fiat currencies (USD, EUR, JPY, and KRW). We then sum all these volumes and take the log. The resulting time series are shown in the left panel of Figure 10.

Similar to volatility, volume tends to be high at the bottom and the height of bitcoin cycles. To verify this, the middle panel plots volume fit against our measure of bitcoin cycles we used in Section 3. The LOESS fit is shown in red. As in the case of volatility, it is a U-shaped function. This is not a random coincidence since in the bitcoin market volume and volatility are strongly related. The latter is confirmed in the right panel of Figure 10 which plots volume against log volatility. Comparing Figures 9 and 10 we

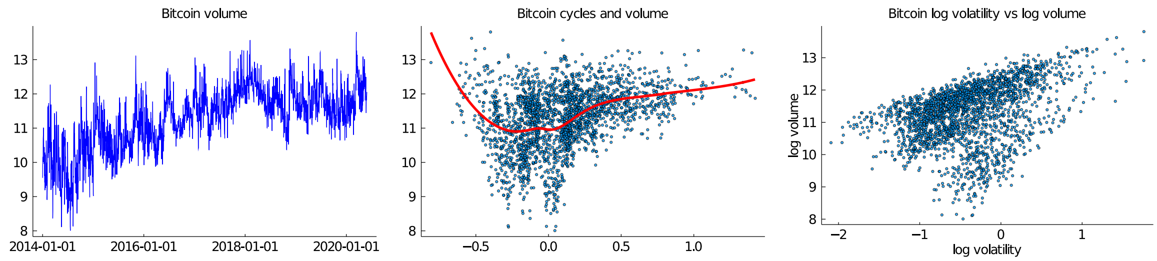

Figure 10: Bitcoin market volume

The left panel depicts the log bitcoin volume for the period from 1/1/2014 to 28/05/2020. The middle panel plots volume fit against our measure of bitcoin cycles. The right panel shows the scatter plot of log volatility versus log volume from data. Bitcoin volume data includes those from the 17 largest and most liquid cryptocurrency exchanges and is provided by Kaiko.

can see that the joint behavior of volume and volatility in the model is similar to that in the bitcoin market.

5. Conclusion

The paper develops a novel model of bubbles. In the model, boundedly rational investors have CRRA preferences. They observe the price process, but do not fully understand how its volatility and expected returns are determined in equilibrium. Investors learn about the expected returns by fitting a parsimonious model to the history of past prices. The model has an explicit solution, which allows us to compute many objects of interest in closed form and to efficiently estimate the model. We fit the model to cryptocurrency markets and show that it is able to explain the joint behavior of expected returns, trading volume, volatility, momentum, and large price impact in these markets.

References

Ait-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approach. *Econometrica*, 70(1):223–263.

Amromin, G. and Sharpe, S. A. (2008). Expectation of risk and return among household investors: Are their sharpe ratios countercyclical? Working paper.

Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24.

Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, 129(2):203–227.

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.

Bhambhwani, S., Delikouras, S., and Korniotis, G. M. (2019). Do fundamentals drive cryptocurrency returns? CEPR Discussion Paper 13724, Center for Economic and Policy Research.

Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.

Brandt, M. W. and Santa-Clara, P. (2002). Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets. *Journal of Financial Economics*, 63(2):161–210.

Brunnermeier, M. K. and Oehmke, M. (2013). Bubbles, financial crises, and systemic risk. In *Handbook of the Economics of Finance*, volume 2, pages 1221–1288. Elsevier.

Choi, J. and Mertens, T. (2013). Extrapolative expectations and the equity premium. Working paper.

Cong, L. W., Li, Y., and Wang, N. (2020). Tokenomics: Dynamic adoption and valuation. NBER Working Paper 27222, National Bureau of Economic Research.

Cutler, D. M., Poterba, J. M., and Summers, L. H. (1990). Speculative dynamics and the role of feedback traders. NBER Working Paper 3243, National Bureau of Economic Research.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance*, 45(2):379–395.

Elerian, O., Chib, S., and Shephard, N. (2001). Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, 69(4):959–993.

Eraker, B. (2001). Mcmc analysis of diffusion models with application to finance. *Journal of Business & Economic Statistics*, 19(2):177–191.

Florens-Zmirou, D. (1989). Approximate discrete-time schemes for statistics of diffusion processes. *Statistics*, 20(4):547–557.

Fuster, A., Hebert, B., and Laibson, D. (2012). Natural expectations, macroeconomic dynamics, and asset pricing. *NBER Macroeconomics Annual*, 26:1–48.

Greenwood, R. and Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, 27(3):714–746.

Harris, M. and Raviv, A. (1993). Differences of opinion make a horse race. *The Review of Financial Studies*, 6(3):473–506.

Haruvy, E., Lahav, Y., and Noussair, C. N. (2007). Traders' expectations in asset markets: experimental evidence. *American Economic Review*, 97(5):1901–1920.

Hirshleifer, D., Li, J., and Yu, J. (2015). Asset pricing in production economies with extrapolative expectations. *Journal of Monetary Economics*, 76:87–106.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6):2143–2184.

Hong, H. and Stein, J. C. (2003). Differences of opinion, short-sales constraints, and market crashes. *Review of Financial Studies*, 16(2):487–525.

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. Academic Press.

Jin, L. J. and Sui, P. (2019). Asset pricing with return extrapolation. Working paper.

Kandel, E. and Pearson, N. D. (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy*, 103(4):831–872.

Karlin, S. and Taylor, H. E. (1981). *A second course in stochastic processes*. Elsevier, 1st edition.

Karpoff, J. M. (1987). The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis*, 22(1):109–126.

Khasminskii, R. (1980). *Stochastic stability of differential equations*. Springer Science & Business Media, 2nd edition.

Li, K. and Liu, J. (2019). Extrapolative asset pricing. Working paper.

Liao, J., Peng, C., and Zhu, N. (2021). Price and volume dynamics in bubbles. Working paper.

Liu, L. Y., Patton, A. J., and Sheppard, K. (2015). Does anything beat 5-minute rv? A comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311.

Liu, Y. and Tsyvinski, A. (2020). Risks and returns of cryptocurrency. *Review of Financial Studies*.

Lo, A. W. and Wang, J. (2003). Trading volume. In Dewatripont, M., Hansen, L. P., and Turnovsky, S. J., editors, *Advances in economics and econometrics: theory and applications, Eighth World Congress*, volume 2, chapter 6, pages 206–277. Cambridge University Press.

Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887.

Pagnotta, E. (2020). Decentralizing money: Bitcoin prices and blockchain security. *Review of Financial Studies*.

Pagnotta, E. and Buraschi, A. (2018). An equilibrium valuation of bitcoin and decentralized network assets. Working paper, Imperial College.

Pedersen, A. R. (1995). A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scandinavian Journal of Statistics*, 22(1):55–71.

Rabin, M. and Vayanos, D. (2010). The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies*, 77(2):730–778.

Scheinkman, J. A. and Xiong, W. (2003). Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6):1183–1220.

Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151.

Sockin, M. and Xiong, W. (2020). A model of cryptocurrencies. NBER Working Paper 26816, National Bureau of Economic Research.

Vissing-Jorgensen, A. (2003). Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions. *NBER Macroeconomics Annual*, 18:139–194.

Xiong, W. (2013). Bubbles, crises, and heterogeneous beliefs. NBER Working Paper 18905, National Bureau of Economic Research.

Appendix. Proofs

Proof of Lemma 1

The observation equation (6) is

$$dz_t = h_t dt + \sigma dB_t. \tag{A1}$$

Denote the correlation between dB_t and $d\hat{B}_t$ by ρ. Then, investors' estimate of the state variable evolves (Jazwinski (1970), p.239) according to

$$d\theta_t = \gamma(h_t - \theta_t)dt + k(dz_t - \theta_t dt), \tag{A2}$$

where

$$k = \sqrt{\gamma^2 + \hat{\sigma}^2 + 2\gamma\rho\sigma\hat{\sigma}} - \gamma. \tag{A3}$$

Q.E.D.

Proof of Proposition 1

The general solution of the linear ODE (14) is

$$\pi(\theta_t) = \frac{1}{k} e^{-\frac{(\theta_t+\sigma)^2}{2k}} \Big(\int_0^{\theta_t} x e^{\frac{(x+\sigma)^2}{2k}} dx + C \Big). \tag{A4}$$

In equilibrium, the portfolio choice $\pi(\theta) = \theta/\sigma(\theta)$ is equal to the ratio of asset price and aggregate wealth P_t/W_t. Applying Ito's Lemma to $\pi(\theta)$ and substituting in (7), (3) and (11),

$$\begin{aligned} d\pi(\theta_t) &= \pi'(\theta_t)d\theta_t + \frac{1}{2}\pi''(\theta_t)(d\theta_t)^2 \\ &= \left[\big(\gamma\bar{h} + kh(\theta_t) - (\gamma+k)\theta_t\big)\pi'(\theta_t) + \frac{k^2}{2}\pi''(\theta_t) \right] dt + k\pi'(\theta_t)dB_t, \end{aligned} \tag{A5}$$

$$\begin{aligned} d\pi(\theta_t) &= d\left(\frac{P_t}{W_t}\right) \\ &= \frac{1}{W_t}dP_t - \frac{P_t}{W_t^2}dW_t - \frac{1}{W_t^2}dP_t dW_t + \frac{1}{2}\frac{2P_t}{W_t^3}(dW_t)^2 \\ &= \frac{P_t}{W_t}\Big[\big(\rho - \mu + (\sigma_t^P - \theta_t)h_t - (\theta_t+\sigma)(\sigma_t^P - \theta_t - \sigma)\big)dt + \big(\sigma_t^P - \theta_t - \sigma\big)dB_t \Big]. \end{aligned} \tag{A6}$$

Comparing the coefficients of dt terms in two expressions of $d\pi(\theta_t)$, substituting in (14) and simplifying, we arrive at (16).

Consider the limit of π_t when θ_t tends to zero. If this limit is nonzero, the Sharpe ratio h_t does not tend to infinity when $\theta_t \to 0$. Therefore, equation (7) that describes the evolution of investors' beliefs implies that both positive and negative realizations of θ_t are possible. Since investors perceive the expected excess return to be negative (positive) for negative (positive) θ their portfolio choice $\pi(\theta)$ should change sign at $\theta_t = 0$, contradiction.

Thus, $\pi(0) = 0$, and therefore, there is only one solution in (A4), which is consistent with equilibrium. Namely,

$$\pi(\theta_t) = \frac{1}{k} e^{-\frac{(\theta_t+\sigma)^2}{2k}} \int_0^{\theta_t} x e^{\frac{(x+\sigma)^2}{2k}} dx$$

which is equation (15).

Next, we prove that $\pi(\theta_t)$ is an increasing function of θ_t.

$$\begin{aligned}
\pi(\theta) &= \frac{1}{k} e^{-\frac{(\theta+\sigma)^2}{2k}} \int_0^{\theta} x e^{\frac{(x+\sigma)^2}{2k}} dx \\
&< \frac{1}{k} e^{-\frac{(\theta+\sigma)^2}{2k}} \int_0^{\theta} \frac{\theta}{\theta+\sigma} (x+\sigma) e^{\frac{(x+\sigma)^2}{2k}} dx \\
&= \frac{\theta}{\theta+\sigma} e^{-\frac{(\theta+\sigma)^2}{2k}} \int_0^{\theta} e^{\frac{(x+\sigma)^2}{2k}} d\left(\frac{(x+\sigma)^2}{2k}\right) \\
&= \frac{\theta}{\theta+\sigma} \left(1 - e^{-\frac{\sigma^2+2\sigma\theta}{2k}}\right) < \frac{\theta}{\theta+\sigma}
\end{aligned} \tag{A7}$$

The first inequality follows from the fact that $x/(x+\sigma)$ is increasing for $x \in [0, \theta]$. Combining (14) and (A7), we arrive at $\pi'(\theta) > 0$.

Finally, consider the dynamics of θ_t. Equation (7) implies that

$$d\theta_t = \mu_\theta(\theta_t)dx + \sigma_\theta(\theta_t)dB_t, \tag{A8}$$

where

$$\mu_\theta(\theta) = \gamma \bar{h} + kh(\theta) - (\gamma + k)\theta, \tag{A9}$$

$$\sigma_\theta(\theta) = k. \tag{A10}$$

Define

$$s'(\theta) = \exp\left\{-2\int_{\theta_0}^{\theta} \frac{\mu_\theta(y)}{\sigma_\theta^2(y)}dy\right\} = \exp\left\{-\frac{2}{k^2}\int_{\theta_0}^{\theta}(\gamma\bar{h} + kh(y) - (\gamma+k)y)dy\right\},$$
$$s(\theta) = \int_{\theta_0}^{\theta} s'(y)dy. \tag{A11}$$

The function $s(\theta)$ is known as the scale function of the diffusion θ_t. It can be verified that

$$\lim_{\theta\to\infty} s(\theta) = \infty, \tag{A12}$$
$$\lim_{\theta\to 0} s(\theta) = -\infty. \tag{A13}$$

Therefore, θ_t is recurrent on $(0,\infty)$ (Khasminskii (1980), p.95). A stationary density of θ_t satisfies the Kolmogorov forward equation

$$0 = \frac{\partial^2}{\partial\theta^2}\left[\sigma_\theta(\theta)^2\psi(\theta)\right] - \frac{\partial}{\partial\theta}\left[\mu_\theta(\theta)\psi(\theta)\right]. \tag{A14}$$

The solution can be written (Karlin and Taylor (1981), p. 221) as

$$\psi(\theta) = \frac{C}{\sigma_\theta^2(\theta)s'(\theta)} = C\exp\left\{\frac{2}{k^2}\int_{\theta_0}^{\theta}(\gamma\bar{h} + kh(y) - (\gamma+k)y)dy\right\}, \quad x \geq 0. \tag{A15}$$

Q.E.D.

Proof of Proposition 2

For $\theta_t \in [\underline{\theta}, \overline{\theta}]$, define

$$g(\theta_t) = \frac{2}{k^2}\left[\frac{\int_{\underline{\theta}}^{\overline{\theta}}\left(\psi(\theta)^{-1}\int_{\underline{\theta}}^{\theta}\psi(y)dy\right)d\theta}{\int_{\underline{\theta}}^{\overline{\theta}}\psi(\theta)^{-1}d\theta}\int_{\underline{\theta}}^{\theta_t}\psi(\theta)^{-1}d\theta - \int_{\underline{\theta}}^{\theta_t}\left(\psi(\theta)^{-1}\int_{\underline{\theta}}^{\theta}\psi(y)dy\right)d\theta\right] \tag{A16}$$

It can be verified that $g(\underline{\theta}) = g(\overline{\theta}) = 0$, and that $g(\theta_t)$ satisfies

$$\mu(\theta_t)g'(\theta_t) + \frac{1}{2}\sigma^2(\theta_t)g''(\theta_t) = -1 \tag{A17}$$

Let $T_{\underline{\theta}\overline{\theta}} = \{\inf t | \theta_t = \underline{\theta} \text{ or } \theta_t = \overline{\theta}\}$. $T_{\underline{\theta}\overline{\theta}}$ is a stopping time with $E^{\theta}[T_{\underline{\theta}\overline{\theta}}] < \infty$. Then

Dynkin's formula holds:

$$E^{\theta}[g(\theta_{T_{\underline{\theta}\overline{\theta}}})] = g(\theta_0) + E^{\theta}\left[\int_0^{T_{\underline{\theta}\overline{\theta}}}\left(g'(\theta_t) + \frac{1}{2}\sigma(\theta_t)^2 g''(\theta_t)\right)dt\right]$$
$$= g(\theta_0) + E^{\theta}\left[\int_0^{T_{\underline{\theta}\overline{\theta}}}(-1)dt\right] \quad \text{(A18)}$$
$$ET_{\underline{\theta}\overline{\theta}} = g(\theta_0) - E^{\theta}[g(\theta_{T_{\underline{\theta}\overline{\theta}}})] = g(\theta_0) \quad \text{(A19)}$$

Taking the limit $\overline{\theta} \to \infty$ in (A16), we arrive at (19). Taking the limit $\underline{\theta} \to 0$, we arrive at (20) . *Q.E.D.*

Proof of Proposition 3

Investor i's wealth evolve according to

$$d\log(W_t^i) = \left(r + \mu - \rho + \pi_t^i\sigma_t^P h_t - \frac{1}{2}(\pi_t^i\sigma_t^P)^2\right)dt + \pi_t^i\sigma_t^P dB_t, \quad \text{(A20)}$$

Therefore,

$$\log R^i = \lim_{T\to\infty}\frac{1}{T}E\log(W_T^i/W_0^i)$$
$$= r + \mu - \rho + \lim_{T\to\infty}\frac{1}{T}\int_0^T E\left[\pi_t^i\sigma_t^P h_t - \frac{1}{2}(\pi_t^i\sigma_t^P)^2\right]dt$$
$$= r + \mu - \rho + E\left[\pi_t^i\sigma_t^P h_t - \frac{1}{2}(\pi_t^i\sigma_t^P)^2\right]. \quad \text{(A21)}$$

Substituting rational investors' portfolio choice $\pi_t^r = h_t/\sigma_t^P$ into (A21), we obtain (22).

Consider an investor whose belief $\hat{\theta}_t$ about the Sharpe ratio obey the following dynamics

$$d\widehat{\theta}_t = \widehat{\gamma}(\bar{h} - \widehat{\theta}_t) + \widehat{k}(h(\theta_t) - \widehat{\theta}_t)dt + \widehat{k}dB_t. \quad \text{(A22)}$$

Applying Ito's lemma to $\widehat{\theta}_t^2$,

$$d\widehat{\theta}_t^2 = 2\widehat{\theta}_t d\widehat{\theta}_t + (d\widehat{\theta}_t)^2 = (\widehat{k}^2 + 2\widehat{\gamma}\widehat{\theta}_t(\bar{h} - \widehat{\theta}_t) + 2\widehat{k}\widehat{\theta}_t(h(\theta_t) - \widehat{\theta}_t))dt + 2\widehat{k}\widehat{\theta}_t dB_t.$$

Her portfolio choice is $\widehat{\pi}_t = \widehat{\theta}_t/\sigma_t^P$ and her wealth evolve according to

$$d\log(\widehat{W}_t) = \left(r - \rho + \widehat{\theta}_t h_t - \frac{1}{2}\widehat{\theta}_t^2\right)dt + \widehat{\theta}_t dB_t$$
$$= \left(r + \mu - \rho - \frac{\widehat{k}}{2} + \frac{1}{2}\widehat{\theta}_t^2 + \frac{\widehat{\gamma}}{\widehat{k}}\widehat{\theta}_t(\widehat{\theta}_t - \bar{h})\right)dt + \frac{1}{2\widehat{k}}d\widehat{\theta}_t^2. \quad \text{(A23)}$$

Therefore her expected logarithmic return is equal to

$$\begin{aligned}\log \widehat{R} &= \lim_{T\to\infty} \frac{1}{T} E \log(\widehat{W}_T/\widehat{W}_0) \\ &= r + \mu - \rho + \frac{\widehat{k} + 2\widehat{\gamma}}{2\widehat{k}} E[\widehat{\theta}_t^2] - \frac{\widehat{\gamma}\bar{h}}{\widehat{k}} E[\widehat{\theta}_t] - \frac{\widehat{k}}{2}. \end{aligned} \tag{A24}$$

Setting $\widehat{k} = k$ and $\widehat{\gamma} = \gamma$, we have $\widehat{\theta}_t = \theta_t$ and from the above equation we obtain (24) for the feedback investor.

Solving (A22), we obtain the expression of $\widehat{\theta}_t$

$$\widehat{\theta}_t = \frac{\widehat{\gamma}}{\widehat{\gamma} + \widehat{k}} \bar{h} + \int_{-\infty}^{t} \widehat{k} e^{-(\widehat{\gamma}+\widehat{k})(t-s)} dz_s \tag{A25}$$

$$= \frac{\widehat{\gamma}}{\widehat{\gamma} + \widehat{k}} \bar{h} + \int_{-\infty}^{t} \widehat{k} e^{-(\widehat{\gamma}+\widehat{k})(t-s)} (h_s ds + dB_s). \tag{A26}$$

The unconditional expectation of $\widehat{\theta}_t$ is equal to

$$E[\widehat{\theta}_t] = \frac{\widehat{\gamma}}{\widehat{\gamma} + \widehat{k}} \bar{h} + \int_{-\infty}^{t} \widehat{k} e^{-(\widehat{\gamma}+\widehat{k})(t-s)} E[h_s] ds = \frac{\widehat{\gamma}}{\widehat{\gamma} + \widehat{k}} \bar{h} + \frac{\widehat{k}}{\widehat{\gamma} + \widehat{k}} E[h_t]. \tag{A27}$$

Now we look at the limit $\widehat{k} \to 0$ and $\widehat{\gamma} \to 0$. In this case the variance of $\widehat{\theta}_t$ tend to zero and the stationary distribution of $\widehat{\theta}_t$ tend to a degenerate distribution $\widehat{\theta}_t = E[h_t]$. We argue that this choice is optimal and the expected logarithmic return is equal to

$$\log \widehat{R} = r + \mu - \rho + \frac{1}{2}(E[h_t])^2. \tag{A28}$$

For the proof of optimality, we need to calculate the unconditional variance of $\widehat{\theta}_t$ for any $\widehat{k}$. The innovation in the return process dz_t is negatively autocorrelated. For time $s < t$, positive realized return at time s increases θ_t and therefore decreases $h(\theta_t)$

$$\begin{aligned}\mathrm{Cov}(dz_s, dz_t) &= \mathrm{Cov}(dz_s, h_t dt + dB_t) \\ &= \mathrm{Cov}(dz_s, h(\theta_t)) dt < 0. \end{aligned} \tag{A29}$$

$\widehat{\theta}_t$ is a weighted average of dz_t. The variance of $\widehat{\theta}_t$ is thus smaller than the sum of variance of its components

$$\begin{aligned}\operatorname{Var}[\widehat{\theta}_t] &< \int_{-\infty}^{t} \operatorname{Var}\left[\widehat{k} e^{-(\widehat{\gamma}+\widehat{k})(t-s)} dz_s\right] \\ &= \int_{-\infty}^{t} \widehat{k}^2 e^{-2(\widehat{\gamma}+\widehat{k})(t-s)} \operatorname{Var}[dz_s] = \frac{\widehat{k}^2}{2(\widehat{\gamma}+\widehat{k})} \end{aligned} \tag{A30}$$

For any positive $\widehat{k}$ and $\widehat{\gamma}$,

$$\begin{aligned}\log \widehat{R} &= r+\mu-\rho+\frac{\widehat{k}+2\widehat{\gamma}}{2\widehat{k}} E[\widehat{\theta}_t^2] - \frac{\widehat{\gamma}\bar{h}}{\widehat{k}} E[\widehat{\theta}_t] - \frac{\widehat{k}}{2} \\ &= r+\mu-\rho+\frac{\widehat{k}+2\widehat{\gamma}}{2\widehat{k}} (E[\widehat{\theta}_t])^2 - \frac{\widehat{\gamma}\bar{h}}{\widehat{k}} E[\widehat{\theta}_t] + \frac{\widehat{k}+2\widehat{\gamma}}{2\widehat{k}} \operatorname{Var}[\widehat{\theta}_t] - \frac{\widehat{k}}{2} \\ &= r+\mu-\rho+\frac{1}{2}(E[h_t])^2 - \frac{\widehat{\gamma}^2}{2(\widehat{\gamma}+\widehat{k})^2}(E[h_t]-\bar{h})^2 + \frac{\widehat{k}+2\widehat{\gamma}}{2\widehat{k}} \operatorname{Var}[\widehat{\theta}_t] - \frac{\widehat{k}}{2} \\ &< r+\mu-\rho+\frac{1}{2}(E[h_t])^2 + \frac{\widehat{k}+2\widehat{\gamma}}{2\widehat{k}} \cdot \frac{\widehat{k}^2}{2(\widehat{\gamma}+\widehat{k})} - \frac{\widehat{k}}{2} \\ &< \lim_{\substack{\widehat{k}\to 0 \\ \widehat{\gamma}\to 0}} \log \widehat{R} \end{aligned} \tag{A31}$$

In other words, choosing $\widehat{k} = \widehat{\gamma} = 0$ and fixing the estimate of Sharpe ratio at the unconditional estimate of Sharpe ratio $E[h_t]$, on average, yield a higher average logarithmic return than that for any positive $\widehat{k}$. *Q.E.D.*

Proof of Proposition 4

Define $Z_t = P_t/W_t$. Using the market clear condition (32) and substituting in the optimal portfolio policies of the feedback and constant-belief investors (30)-(31),

$$\begin{aligned} Z_t = \frac{P_t}{W_t} &= \frac{\pi_t W_t + \widehat{\pi}_t \widehat{W}_t}{W_t}, \\ &= \frac{\hat{\theta} X_t + \theta_t}{\sigma_t^P} \end{aligned} \tag{A32}$$

Applying Ito's Lemma to Z_t,

$$\begin{aligned} dZ_t =& \Big[(\gamma\bar{h} + kh_t - (\gamma + k)\theta_t)Z_\theta + (\widehat{\theta} - \theta_t)\left(h_t - \sigma - \theta_t\right) X_t Z_X \\ &+ \frac{1}{2}k^2 Z_{\theta\theta} + k(\widehat{\theta} - \theta_t)X_t Z_{\theta X} + \frac{1}{2}(\widehat{\theta} - \theta_t)^2 X_t^2 Z_{XX}\Big] dt \\ &+ \Big[kZ_\theta + (\widehat{\theta} - \theta_t)X_t Z_X\Big] dB_t, \qquad \text{(A33)} \\ dZ_t =& d\left(\frac{P_t}{W_t}\right) \\ =& \frac{P_t}{W_t}\Big[\big(\rho - \mu + (\sigma_t^P - \theta_t)h_t - (\theta_t + \sigma)(\sigma_t^P - \theta_t - \sigma)\big)dt + \big(\sigma_t^P - \theta_t - \sigma\big)dB_t\Big]. \qquad \text{(A34)} \end{aligned}$$

Comparing the coefficients of dB_t,

$$\begin{aligned} kZ_\theta + (\widehat{\theta} - \theta_t)X_t Z_X &= Z_t\big(\sigma_t^P - \theta_t - \sigma\big) \\ &= \hat{\theta}X_t + \theta_t - (\theta_t + \sigma)Z_t \qquad \text{(A35)} \end{aligned}$$

Along the characteristic curve $X(\theta) = C\exp(-\frac{1}{2k}(\theta - \widehat{\theta})^2)$,

$$\begin{aligned} k\frac{dZ}{d\theta} + (\theta + \sigma)Z =& \widehat{\theta}X(\theta) + \theta, \qquad \text{(A36)} \\ Z =& \frac{1}{k}e^{-\frac{(\theta+\sigma)^2}{2k}} \int_{\underline{\theta}(X,\theta)}^{\theta} (\widehat{\theta}X(y) + y)e^{\frac{(y+\sigma)^2}{2k}} dy. \qquad \text{(A37)} \end{aligned}$$

where $\underline{\theta}(X, \theta) < 0$ is the solution to

$$\widehat{\theta}X(\underline{\theta}) + \underline{\theta} = 0 \qquad \text{(A38)}$$

Note that $C = X\exp(\frac{1}{2k}(\theta - \widehat{\theta})^2)$. Therefore, $\underline{\theta}(X, \theta)$ solves (38). Comparing the coefficients of dt in (A33) and (A34), we obtain (36). The proof of uniqueness follows similar steps as in Proporsition 1. *Q.E.D.*