Upside risk and return timing in Bitcoin

Daniel Batista*[1,2] and Marcelo Fernandes†[3]

[1] Université de Genève

[2] Swiss Finance Institute

[3] Sao Paulo School of Economics, FGV

March, 2026

Abstract

This paper studies the benefits of distinguishing between upside and downside volatility when timing returns in Bitcoin. Standard volatility management implicitly treats volatility spikes as signals of adverse states, thereby reducing exposure when realized volatility increases. In Bitcoin, however, volatility spikes are frequently driven by positive price movements and, as such, often indicate positive returns in the next period. We show that semivolatility timing rules that not only control for downside risk but also exploit upside risk yield substantially stronger risk-adjusted performance than both buy-and-hold and volatility-managed strategies. Neither unusually persistent volatility nor standard linear return predictability explain our findings. Instead, we document that Bitcoin exhibits an asymmetric volatility structure in which high upside-driven volatility states are disproportionately associated with positive subsequent returns.

*E-mail: daniel.batistadasilva@unige.ch

†E-mail: marcelo.fernandes@fgv.br

1 Introduction

Investors resort to risk-timing rules to improve on their risk-adjusted performance. Moreira and Muir (2017a) show that Sharpe ratios typically improve if investors decrease their exposure to risk factors when volatility is high. See also Barroso and Santa-Clara (2015), Liu et al. (2019), Cederburg et al. (2020b), Eisdorfer and Misirli (2020), and Barroso and Detzel (2021). In turn, Wang and Yan (2021) propose to time downside volatility given it is relatively better at predicting future returns than volatility. In Batista and Fernandes (2025), we entertain several forms of semivolatility management. In particular, we show that timing by the ratio of upside volatility to downside variance offers the most robust risk-adjusted performance gains for it reflects both skewness and downside risk concerns (Bianchi et al., 2022).

The above papers consider exclusively equity factors and anomaly portfolios, though. In this note, we examine Bitcoin (BTC) returns with a simple motivation in mind. While high-volatility episodes in equity markets are mainly driven by downside variation, the opposite applies to Bitcoin in that upside variation often drives large price movements. As such, Bitcoin is a fertile ground for timing both downside and upside risks. We indeed show that scaling by the upside-volatility-to-downside-variance ratio performs unusually well relative to standard volatility management given that the latter penalizes both downside and upside volatilities.

The economic intuition is straightforward. Volatility management scales exposure inversely with realized variance, implicitly treating volatility spikes as signals of unfavorable states. However, if a large fraction of volatility originates from positive price movements, such a timing rule would reduce exposure precisely when Bitcoin enters favorable states. By decomposing volatility into upside and downside components, semivolatility management not only avoids indiscriminate deleveraging after upside-driven volatility shocks, but also maintain or even increase exposure when total volatility reflects positive returns in the near future.

Using monthly Bitcoin returns and realized semivolatility measures, we show that exploiting upside volatility timing yields unusually large improvements in risk-adjusted performance relative to buy-and-hold, volatility management, and downside-volatility management. Neither unusually persistent volatility or standard linear return predictability explain these gains, though. They instead derive from the asymmetric volatility structure in Bitcoin returns, with high upside-driven volatility states disproportionately associated with subsequent positive returns. This explains why volatility management performs poorly in Bitcoin and why distinguishing between upside and downside volatility entails substantially gains. Altogether, we document an implication of the stylized fact that volatility spikes in Bitcoin often relate to positive return states.

2 Baseline Evidence

This section documents the main empirical fact we examine in this note. To do so, we assess the performance of risk timing rules in Bitcoin. In particular, we consider a simple buy-and-hold strategy, as well as volatility management as in Moreira and Muir (2017b), semivolatility management that controls only for downside risk as in Wang and Yan (2021), and semivolatility management that also exploits upside risk as in Batista and Fernandes (2025). The first two scale the Bitcoin investment respectively by the inverse of the total and downside variances, whereas the third considers the upside-volatility-to-downside-variance ratio. See Batista and Fernandes (2025) for methodological details on how we exactly run each strategy at the monthly frequency using realized measures based on daily returns.

We collect BTC daily returns from Yahoo Finance, which sources crypto data from CoinMarketCap. Given the usual concerns with crypto data (Schwenkler et al., 2026), it is worth stressing that our empirical results remain virtually the same quantitatively using Coin Codex data. Figure 1 plots the Bitcoin cumulative returns not only for the buy-and-hold benchmark, but also for each timing strategy. A striking pattern emerges. Accounting for both downside and upside risks yields substantially larger cumulative returns, whereas controlling only for total or downside variation fails to improve the buy-and-hold strategy significantly.

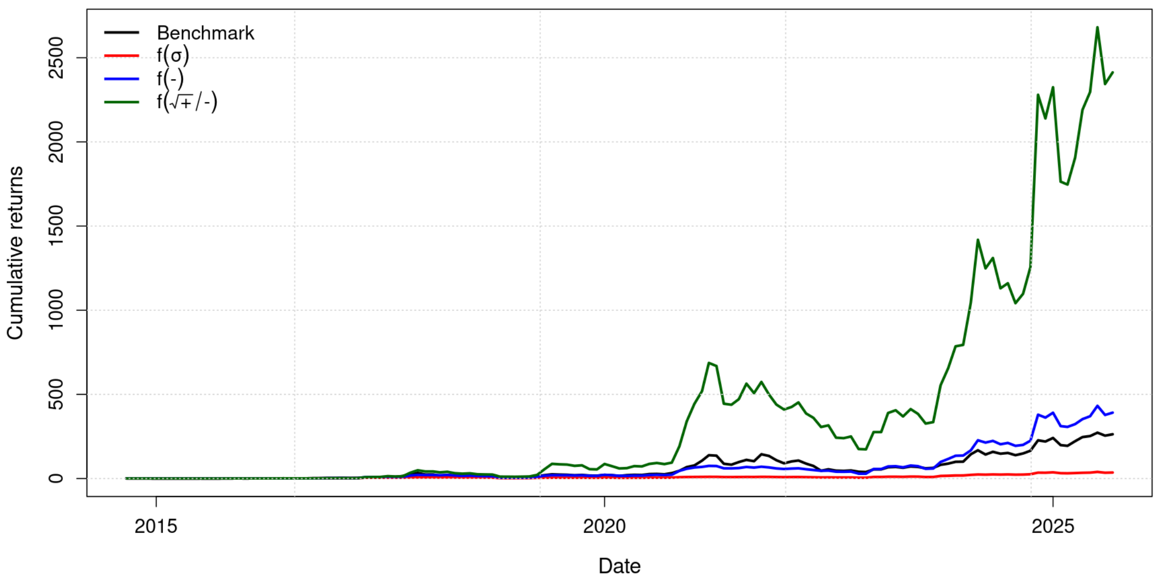

Figure 1: Bitcoin cumulative returns for both buy-and-hold and managed strategies

The black line tracks the cumulative returns on a buy-and-hold strategy for the Bitcoin, whereas the colored lines refer to the managed strategies. In particular, we depict the cumulative returns corresponding to volatility management in red, downside volatility in blue, and both upside and downside management in green.

Table 1 summarizes how they perform in terms of Sharpe and Sortino ratios. As it turns out, volatility management reduces risk-adjusted performance relative to buy-and-hold, with Sharpe and Sortino ratios declining from 0.74 and 1.26 to 0.55 and 0.90, respectively. Timing only downside volatility does not seem to help much as well, with Sharpe and Sortino ratios of 0.72 and 1.39, respectively. In contrast, accounting for both upside and downside volatilities pays off, delivering the best risk-adjusted performance with Sharpe and Sortino ratios of 0.81 and 1.70, respectively. Interestingly, these gains are not driven by volatility timing, but rather by distinguishing between upside and downside variations. In fact, using Wang and Yan's (2021) alpha decomposition, we find that the return-timing component is large and positive at 2.97 per year, whereas the volatility-timing component is negative at −0.45 per year. As such, the downside-and-upside volatility management does not generate alpha by systematically reducing risk; if anything, volatility timing slightly deteriorates the BTC performance.

Table 1: Performance of Bitcoin strategies

management strategy	mean	volatility	Sharpe	Sortino
buy-and-hold benchmark	0.503	0.679	0.740	1.256
total volatility	0.323	0.584	0.554	0.900
downside volatility	0.539	0.745	0.723	1.386
downside-and-upside volatility	0.703	0.866	0.811	1.704

A natural question is whether these performance gains also arise for other cryptocurrencies, or it is unique to Bitcoin. To address this, we collect daily data for 87 crypto coins in Yahoo Finance.[1] For each cryptocurrency i at time t, we first estimate in the first half of the sample the following asymmetry measure:

$$A_i = \Pr(r_{i,t} > 0 \mid UV_{i,t-1} \text{ in top decile}) - \Pr(r_{i,t} > 0 \mid RV_{i,t-1} \text{ in top decile}), \quad (1)$$

where UV and RV respectively denote realized upside and total variances. The idea is to capture whether positive returns at time t are more likely given a high upside variance at time $t-1$ than given a high total variance. Higher values of A_i indicate that upside-driven volatility states contain more favorable information about future returns than total

[1] In particular, apart from BTC, Yahoo Finance publishes daily prices for 1INCH, AAVE, ADA, ALGO, AMP, AR, ATOM, AVAX, AXS, BAT, BCH, BNB, BSV, BTC, CAKE, CFX, CHZ, CKB, CRO, CRV, CVX, DAI, DASH, DCR, DOGE, DOT, EGLD, ENS, ETC, ETH, FET, FIL, FLOW, FTT, GLM, GNO, GT, HBAR, HNT, ICP, INJ, IOTX, JST, KAVA, KCS, KSM, LDO, LEO, LINK, LPT, LTC, MANA, MX, NEAR, NEO, NEXO, OKB, PAXG, QNT, QTUM, RAY, RSR, RUNE, SAND, SHIB, SNX, SOL, SUN, TFUEL, THETA, TRX, TUSD, TWT, USDC, USDT, VET, WEMIX, XAUt, XCN, XDC, XEC, XLM, XMR, XRP, XTZ, ZEC, and ZIL.

volatility states. The asymmetry in BTC is very pronounced, but not extremely so in that it corresponds to the 83rd percentile of the cross-sectional empirical distribution. After sorting every cryptocurrency into quartiles based on the A_i estimates, we then assess the Sharpe ratio differentials relative to timing exclusively volatility or downside volatility in the second half of the sample.

Figure 2 reports the average out-of-sample performance across quartiles. Both volatility and downside-volatility management dominate downside-and-upside variance management in every quartile. However, differences are much smaller for the last quartile that congregates the crypto assets with stronger asymmetry in the first half of the sample. The performance gains of downside-and-upside volatility management clearly stand out for Bitcoin, even if they are not exactly unique. In many assets, upside-driven volatility coincides more frequently with favorable future returns, but with typically weaker and less stable effects.

We next ask whether the stronger performance of the downside-and-upside volatility management for Bitcoin is due to some unusually high persistence in BTC volatility. Figure 3 examines this possibility by looking at where BTC lies in the distribution of half-life estimates for the realized variances and semivariances across the 87 cryptocurrencies. It turns out that Bitcoin does not stand out, with a half-life right within the cross-sectional distribution of other crypto coins. Accordingly, we cannot attribute the performance gains of the downside-and-upside semivolatility management exclusively to persistence in the BTC volatility. In the next section, we further investigate why downside-and-upside volatility management works particularly well for Bitcoin, relative to other crypto coins.

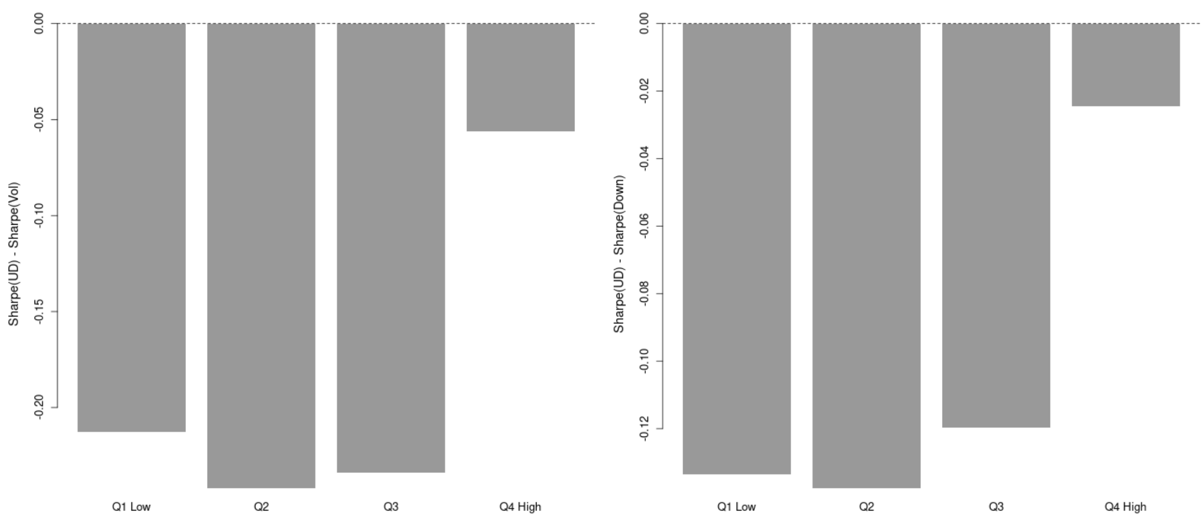

Figure 2: Relative performance of downside-and-upside volatility management across cryptocurrencies
After sorting cryptocurrencies into quartiles based the asymmetry measure estimates we obtain in the first half of the sample, we compute the average Sharpe ratio differences relative to volatility management (left panel) and to downside volatility management (right panel) in the second half of the sample.

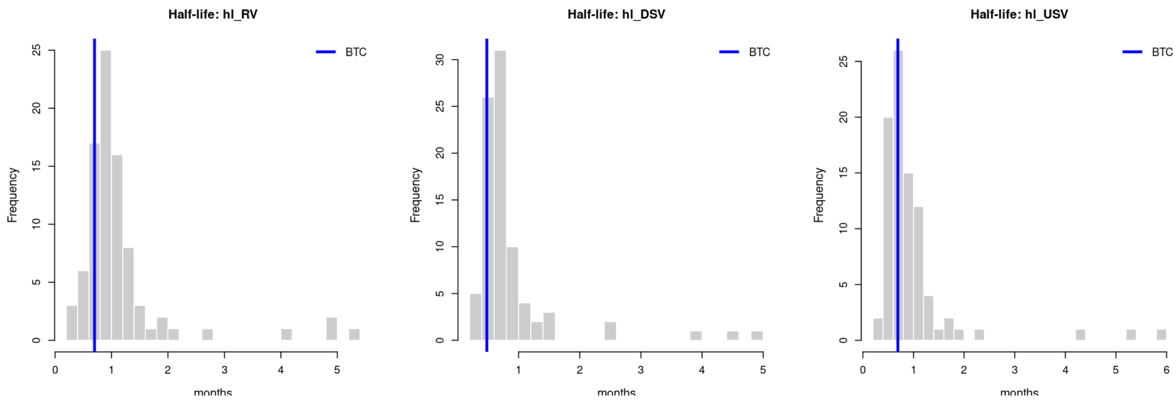

Figure 3: Empirical distribution of the half-lives of the realized (semi)variances across crypto coins
The first panel displays the empirical distribution of the half-life estimates of the realized variance across the 87 crypto currencies. The second and third panel exhibits similar histograms for the downside and upside volatility, respectively. Both panels highlight the position of the BTC half-life estimate in the empirical distribution.

3 Mechanism

We typically associate high volatility with adverse states of nature, so that high-volatility months should precede negative returns. We start by examining how much downside and upside variances help forecast future returns in a linear predictive regression. Figure 4 reports not only the coefficient estimates from the predictive regressions of future returns on downside and upside semivariance measures across cryptocurrencies, but also the corresponding adjusted R^2 values. Most coefficient estimates are noisy and statistically insignificant, including those for Bitcoin. In addition, the adjusted R^2 of the BTC predictive regression fits well in the bulk of the distribution, with a slightly negative value. Accordingly, the performance gains of timing both downside and upside variances do not seem to come from standard linear predictability.

There is no reason to restrict attention to linear relationships, though. Figure 5 plots the cross-sectional empirical distribution of the relative frequency of positive returns given the past variance or upside variance is in the top decile. Bitcoin lies in the right tail of both cross-sectional distributions, so that the likelihood of a positive return in the next month is unusually high if the volatility is high enough this month, especially if driven by upside variation. As such, by scaling exposure inversely with realized variance, standard volatility management reduces exposure precisely in states that subsequently deliver positive returns. In contrast, timing rules that distinguish between upside and downside variation avoid indiscriminate deleveraging after upside-driven volatility spikes.

To further assess whether the performance gains come indeed from this nonlinear relationship between returns and past upside variance, we check how downside-and-upside

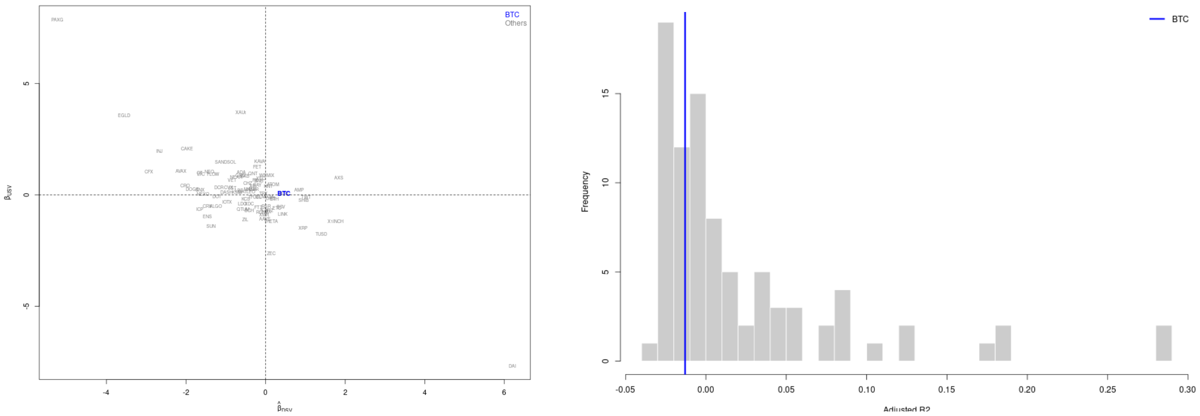

Figure 4: Predictive regressions of returns on past downside and upside semivariances

The left panel depicts the coefficient estimates of the past downside and upside variances, whereas the right plot displays the cross-sectional empirical distribution of the adjusted R^2 values of these predictive regressions. We highlight Bitcoin in blue, whereas we omit the coefficient estimates for two stable coins tethered to the US dollar (namely, USDT and USDC) to avoid distorting the scatter plot in the left panel due to their extreme values: $(46.73, -41.22)$ and $(18.36, -15.26)$, respectively.

volatility management compares with the buy-and-hold strategy for the Bitcoin across bull and bear regimes. Figure 6 shows that accounting for both downside and upside variances yields a similar Sharpe ratio to the buy-and-hold strategy in bull markets, whilst improving on the Sortino ratio. In contrast, both Sharpe and Sortino ratios improve significantly in bear markets. Altogether, it seems that downside-and-upside volatility management reduces exposure in bear regimes, thereby curbing negative performance, because they mainly relate to downside-driven volatility. In turn, it increases exposure in bull regimes only if upside volatility is high enough. This explains why it improves the Sortino ratio, but not the Sharpe measure. As the former scales the average positive returns only by downside variance, it does not penalize the months of positive returns due to upside-driven volatility. Instead, the Sharpe ratio scales by total variance, so that it weights down the average positive return even if the total volatility is high because of the positive semivolatility.

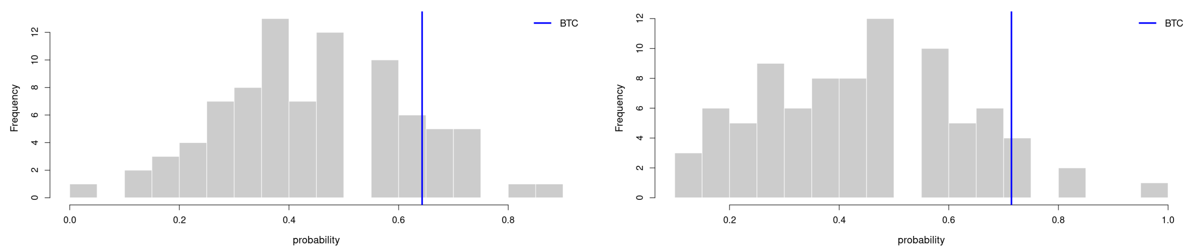

Figure 5: Empirical distribution of the relative frequency of positive returns given past (upside) variance

The left panel plots the cross-sectional empirical distribution of the relative frequency of positive returns given the past variance is in the top decile, whereas the right panel displays the corresponding distribution given the past upside variance is in the top decile. We highlight in blue where Bitcoin lies in each cross-sectional distribution.

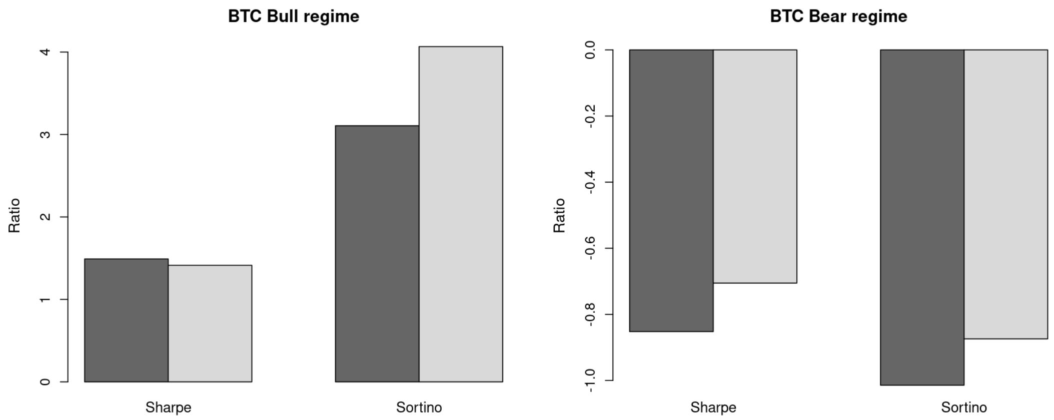

Figure 6: Risk-adjusted performances in bull and bear regimes

We display the Sharpe and Sortino ratios for Bitcoin investments using either the buy-and-hold strategy or the downside-and-upside variance management. The right panel considers only periods in which BTC returns are positive (i.e., bull regime), whereas the left plot refers to the periods in which BTC returns are negative (i.e., bear regime).

Finally, we also check whether our findings are robust to different leverage caps. Cederburg et al. (2020a) argues it is very hard to carry out volatility management without imposing leverage restrictions. This happens because, as long as the realized (semi)volatility is low enough, one would have to increase BTC investments to unrealistic levels if timing only for total (or downside) volatility, without giving proper consideration to upside variation. To avoid unrealistic positions, we impose a maximum leverage of 100% as in Batista and Fernandes (2025). Figure 7 shows how the cumulative returns change if we increase maximum leverage to 200% or decrease to 50%. It is apparent that performance is very sensitive to the leverage cap, indicating that the performance gains depend more on whether we allow exposure to rise in favorable states, than on pure risk reduction.

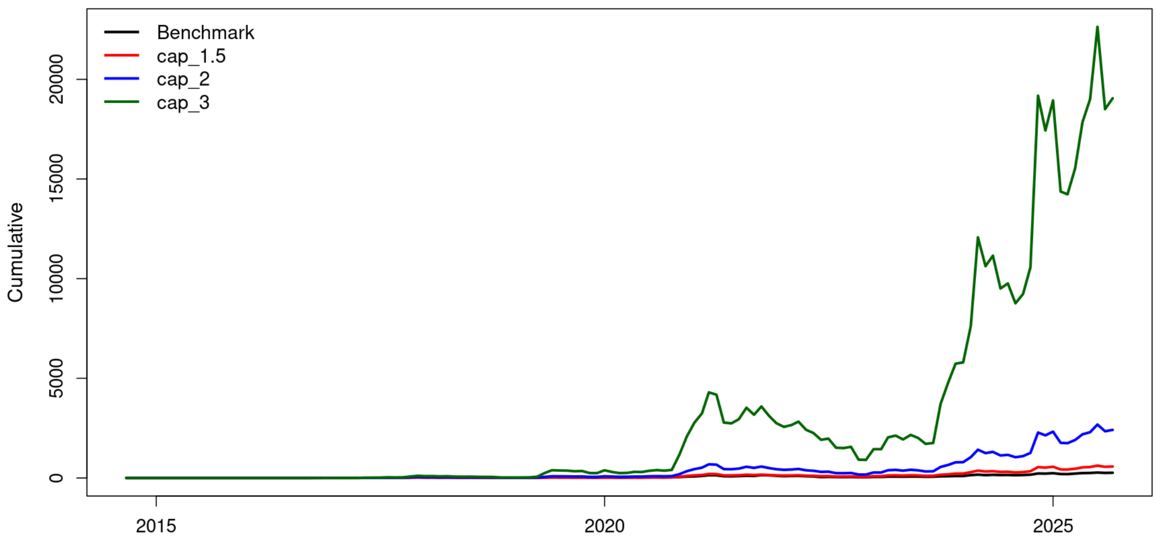

Figure 7: Cumulative returns for the Bitcoin under different leverage caps

We compute the cumulative returns under downside-and-upside volatility management for different leverage caps.

4 Conclusion

We find that downside-and-upside volatility management performs extremely well for Bitcoin. This results neither from unusually high persistence in volatility nor from linear return predictability. Instead, the performance gains we observe are consistent with an asymmetric volatility mechanism: namely, high upside volatility in Bitcoin is disproportionately associated with positive return realizations. This feature helps explain why rules based on total variance can underperform in Bitcoin and why conditioning exposure on upside relative to downside variation yields unusually strong gains. More broadly, the results highlight that volatility spikes do not necessarily correspond to adverse states in all assets, particularly in speculative markets such as cryptocurrencies.

References

Barroso, P., Detzel, A., 2021. Do limits to arbitrage explain the benefits of volatility-managed portfolios? Journal of Financial Economics 140, 744–767.

Barroso, P., Santa-Clara, P., 2015. Momentum has its moments. Journal of Financial Economics 116, 111–120.

Batista, D., Fernandes, M., 2025. Semivolatility-managed portfolios. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4891824. working paper.

Bianchi, D., De Polis, A., Petrella, I., 2022. Taming momentum crashes. Available at SSRN, doi:http://dx.doi.org/10.2139/ssrn.4182040.

Cederburg, S., O'Doherty, M.S., Wang, F., Yan, X., 2020a. On the performance of volatility-managed portfolios. Journal of Financial Economics 138, 95–117.

Cederburg, S., O'Doherty, M.S., Wang, F., Yan, X.S., 2020b. On the performance of volatility-managed portfolios. Journal of Financial Economics 138, 95–117.

Eisdorfer, A., Misirli, E., 2020. Distressed stocks in distressed times. Management Science 66, 2452–2473.

Liu, F., Tang, X., Zhou, G., 2019. Volatility-managed portfolio: Does it really work? Journal of Portfolio Management 46, 38–51.

Moreira, A., Muir, T., 2017a. Volatility-managed portfolios. Journal of Finance 72, 1611–1644.

Moreira, A., Muir, T., 2017b. Volatility-managed portfolios. Journal of Finance 72, 1611–1644.

Schwenkler, G., Shah, A., Yang, D., 2026. Aggregate confusion in crypto market data. conditionally accepted at Management Science .

Wang, F., Yan, X.S., 2021. Downside risk and the performance of volatility-managed portfolios. Journal of Banking and Finance 131. Article 106198.