ECONOMETRIC ANALYSIS OF CRYPTOCURRENCY MARKETS: A TIME SERIES APPROACH

EMRE INAN

A DISSERTATION SUBMITTED TO
THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

GRADUATE PROGRAM IN ECONOMICS
YORK UNIVERSITY
TORONTO, ONTARIO

May 2026

Abstract

Cryptocurrency markets have grown at unprecedented rates in the early 2020s due to increased demand and become further integrated into the broader financial system through new practices, regulatory developments, and institutional adoption. To help develop risk assessment tools and optimize trading strategies, this dissertation studies the stability, efficiency, and predictability of cryptocurrency markets using time series methods.

The stability of cryptocurrency markets is analyzed in the context of stablecoins to better understand whether this class of cryptocurrencies lives up to its name and to mitigate the risks associated with the current practices. Daily stablecoin prices are modelled as a univariate double autoregressive process (DAR) with time-varying parameters and a set of simple plug-in measures of stability is introduced. The local analysis reveals significant temporal and cross-coin variation in the stability of large-cap stablecoins.

To evaluate the efficiency of cryptocurrency markets, a statistical test for the no-arbitrage hypothesis is developed based on a time series model of fundamental and idiosyncratic risk factors. The test is conducted both globally and locally for major cryptocurrency markets via a method of simulated moments. The empirical application shows that the fragmented nature of the spot markets presents significant arbitrage opportunities for cryptocurrency traders, especially when coupled with higher price volatility.

Motivated by their similarity with stablecoin prices, the out-of-sample predictability of funding rates of Bitcoin perpetual futures is evaluated using the one-step-ahead point forecasts generated from a set of DAR models. The results show that the discrepancy between

the spot and futures prices of Bitcoin leads to predictable patterns in historical funding rates, with the degree of predictability varying over time and across exchange platforms.

Acknowledgments

I am deeply grateful to my supervisor, Dr. Joann Jasiak, and my co-supervisor, Dr. Antoine Djogbenou, for their mentorship and generous support over the course of my doctoral studies. This dissertation benefited greatly from their constructive feedback and reached its final form through countless hours of insightful discussions and encouragement.

I would like to thank my parents, Halil and Sevim Inan, whose unconditional love and belief in me made this journey possible. I will forever be indebted to my father, my very first teacher, for introducing me to the beauty of mathematics and nurturing my curiosity for learning from an early age.

Special thanks to my partner and my best friend, Ashlea, for standing by me through every high and low, and for filling my PhD journey with faith, hope, and love.

Contents

List of Tables

List of Figures

Introduction

Propelled by the heightened interest during the early phases of the COVID-19 crisis, crypto markets turned into a trillion-dollar sector, making them economically more relevant than ever before. While the total market capitalization was slightly over $200B in the winter of 2020, it surpassed the $1.00T mark as of January 2021 and reached a peak of $2.86T by the end of the same year. The inflow of new funds not only benefited the likes of Bitcoin but also led stablecoins (a relatively newer type of cryptocurrency) to finally take off and become an integral part of the industry. Tether, the first stablecoin created, capitalized on the increasing demand from investors and established itself as the third-largest cryptocurrency after Bitcoin and Ethereum.

Today, the cryptocurrency ecosystem is quite diverse, comprising a myriad of native tokens, stablecoins, and derivative instruments, with over one thousand crypto exchanges offering services worldwide. The ecosystem is still evolving at a rapid rate and projected to grow even larger in the post-pandemic era. Considering its ever-changing dynamics, this dissertation conducts the analysis not only globally but also locally to disentangle the complexities of cryptocurrency markets and document their evolution over time. Methodologically, the common theme is the use of Double Autoregressive (DAR) processes to accommodate the persistence in the cryptocurrency time series data used throughout the analysis.

This dissertation is composed of three chapters and each chapter itself is a stand-alone essay addressing one of the following questions:

1. Can we develop an empirically robust econometric model to be used by investors

and policymakers to assess the stability of stablecoins in real time and make short-term predictions?

2. Are cryptocurrency spot markets efficient? Is there any difference between stablecoin markets and those of other major cryptocurrencies in terms of efficiency?

3. Does the price discrepancy between the spot and futures markets of cryptocurrencies lead to predictable patterns in the historical funding rates of perpetual contracts?

The first chapter examines the dynamics of Tether, the stablecoin with the largest market capitalization. The analysis is conducted locally for daily Tether/USD rates to detect local trends, time-varying volatility, and persistence, revealing that the distributional and dynamic properties had been evolving over the sampling period. To accommodate these dynamic patterns, the time-varying parameter Double Autoregressive tvDAR(1) process under the assumption of local stationarity is introduced. Based on the Lyapunov exponent of the model, a set of simple plug-in measures of stability is developed to assess and compare the stability of Tether and other stablecoins. The findings show that the tvDAR model provides a good fit to data with time-varying persistence and volatility, and produces reliable out-of-sample forecasts at short horizons. In the empirical application, the proposed stability measures provide reliable indicators of local explosive patterns for stablecoin monitoring.

The second chapter investigates the efficiency of cryptocurrency markets by testing the no-arbitrage hypothesis in a general framework for cryptocurrencies traded on multiple exchange platforms. The moment conditions implied by the no-arbitrage hypothesis are tested by applying a simulated method of moments and a general specification test. For each of the five cryptocurrencies considered, the test is applied globally using the full sample and locally via a rolling window approach. The results show that the null hypothesis of no-arbitrage is rejected globally for each cryptocurrency, whereas it is rejected more often for the native cryptocurrencies than for stablecoins. This implies that the fragmented nature of the markets presents significant arbitrage opportunities for cryptocurrency traders, especially when coupled with higher price volatility.

The third chapter studies the out-of-sample predictability of perpetual futures funding rates with a particular focus on Bitcoin contracts traded on Binance and Bybit. Throughout the analysis, one-step-ahead point forecasts are generated from a set of double autoregressive (DAR) models and evaluated against standard benchmarks. According to the results, the DAR model-based predictions outperform the no-change model both in terms of forecast error and directional accuracy, providing strong evidence for the predictability of the next period's funding rate levels. However, the local analysis shows that the stability of the funding rates is evolving over the evaluation period, suggesting a time-varying degree of their predictability.

Chapter 1

Time-Varying Coefficient DAR Model and Stability Measures for Stablecoin Prices: An Application to Tether

Coauthorship Statement

This chapter is based on joint work titled “Time-Varying Coefficient DAR Model and Stability Measures for Stablecoin Prices: An Application to Tether" with Dr. Antoine Djogbenou and Dr. Joann Jasiak. All authors contributed equally to the conceptualization, data preparation, empirical modeling, analysis, and writing of the manuscript.

A version of this chapter is published in `Journal of International Money and Finance`.

1.1 Introduction

The total market capitalization of cryptocurrencies is currently over 1 trillion US dollars, with the top three cryptocurrencies in terms of market capitalization being Bitcoin (BTC), Ethereum (ETH), and Tether (USDT). While Bitcoin and Ethereum display high volatility, Tether is a stablecoin, i.e. a cryptocurrency designed to maintain a stable price compared to other cryptocurrencies. It is the first and by far the largest stablecoin in the market with the highest daily volume of over $100 billion. For price stability, the value of Tether is pegged 1-to-1 with the US dollar. There also exist other stablecoins, e.g. USD Coin (USDC) and Dai (DAI), pegged to a currency or gold and managed by either a single authority (usually the service provider) or a network of participants (the whole protocol).

Stable cryptocurrencies are alternative financial instruments that can potentially increase public confidence and foster growth and innovation, provided that crypto markets are adequately regulated [Allen et al. (2022)]. The joint report of the US President's Working Group on Financial Markets (PWG), Federal Deposit Insurance Corporation (FDIC), and Office of the Comptroller of the Currency (OCC) from November 2021 highlights various risks that need to be addressed by the regulators. These are user protection and run risk, payment system risk, systemic risk, and concentration of economic power. The report makes various recommendations, including an insured depository institution's requirement on stablecoin issuers [see, the President's Working Group (2021) for more details]. This measure would address the reserve-related risks. Li and Mayer (2022) show that collateralized stablecoins, such as Tether, could create systemic risk if their issuers do not have sufficient reserves to maintain stability. Chen et al. (2022b) point out the important role of Tether in increasing the trading volume of Bitcoin compared to US dollars, and thus increasing the link between Tether and Bitcoin since 2017. Despite the increased attention of regulators and their calls for more scrutiny, the stablecoins have not yet achieved their full price stability. For example, Terra USD, an algorithmic stablecoin, collapsed in May 2022.

The stability is directly related to stablecoin predictability. A stablecoin price is theoretically an exchange rate intended to be perfectly predictable, in contrast with the traditional exchange rate, assumed to have nonstationary martingale dynamics and limited predictability in the sense current price is the best prediction of tomorrow's price. The issuance and redemption mechanisms of stablecoins are designed to achieve and maintain the stability and predictability objectives over time. Perfect predictability may be achievable "globally", i.e. in the long run. However, in the short run, i.e. "locally", the stablecoin prices deviate from the peg, and those deviations can be as large as to cause a crash, as evidenced by Terra USD's collapse.

To address the local instability issue, this paper makes the following contributions. First, we illustrate the instability by studying the local properties of Tether prices. Our empirical "local" analysis reveals the time-varying mean, volatility, and persistence of Tether prices that characterize its deviations from the peg. We show that locally, Tether prices exhibit strong persistence, associated with martingale behavior during local price trends and increased volatility episodes.

Next, we propose a model for stablecoin prices that accommodates the time-varying dynamic properties of Tether and is also locally consistent with the traditional approach to exchange rate modeling. The time-varying parameter Double Autoregressive (tvDAR) model is an extension of the Double Autoregressive (DAR) process of Ling (2004), which combines the autoregressive dynamics with conditional heteroscedasticity. The DAR model nests the ARCH(1) and autoregressive of order one (AR(1)) models, including the unit root model, i.e. the AR(1) model with an autoregressive coefficient equal to 1. When the autoregressive coefficient is one, and the ARCH effect is non-zero, the DAR becomes a stationary martingale process with infinite marginal variance [Gourieroux and Jasiak (2019)]. Then, the parameter estimators of DAR remain consistent and the associated test statistics preserve their asymptotic distributions, unlike in the traditional AR(1) model [Ling (2004), Ling and Li (2008)]. The advantage of the DAR model is that it ensures valid inference on the parameters when the autoregressive coefficient is one. By extending the DAR model to a time-varying parameter tvDAR model, we can capture the dynamics of local deviations of stablecoin from the peg and strong persistence. The tvDAR model also provides reliable short-term predictions.

Our third contribution is a set of three simple plug-in measures of stability for stablecoins. Two of them are based on the expression of the Lyapunov exponent of the tvDAR model.

The third one is a simple testable statistic based on the local stationarity condition.

Let us now discuss these contributions in detail. The proposed extension of the DAR to a deterministic time-varying parameter tvDAR model relies on the assumption of local strict stationarity of the process, following the approach of Dahlhaus (2000) and Dahlhaus et al. (2019). The time-varying coefficient modeling based on the assumption of local stationarity distinguishes our approach from other applied studies relying on the assumption of global strict stationarity of Tether prices. For instance, Baumöhl and Vyrost (2020) use high-frequency data to compute a spectral density-based quantile dependence measure under a global strict stationarity condition, which does not seem to be satisfied by Tether. Bianchi et al. (2022) estimate a Bayesian VAR with stochastic volatility and Student-t distributed shocks (BVAR-SV-t). However, their conditional volatility equation is constrained to unit root dynamics, which is inconsistent with the empirical evidence provided in this paper.

The quasi-maximum likelihood (QML) estimation method is used to estimate the tvDAR model locally. The QML estimators are consistent and asymptotically normally distributed even when the autoregressive parameter is close or equal to one, as documented in Tether prices. Therefore, asymptotically valid confidence intervals with critical values from the normal distribution can be constructed. The rectangular and Epanechnikov kernels are used to implement the estimator locally. One-step ahead forecasts of future Tether prices are easily obtained from the tvDAR model by updating the locally estimated parameters in the forecast formula. The model is also estimated on other stablecoin price processes, showing a good fit to the data.

The tvDAR model provides three simple plug-in measures of stability for stablecoins. Two of these measures, based on the Lyapunov exponent, assume the error distribution is either known or unknown. The Lyapunov exponent is commonly used to assess the stability of deterministic dynamical systems and to test for chaos [see, e.g., Sprott (2014)]. The expression of the Lyapunov exponent for the AR(1)-ARCH model was derived by Borkovec and Kluppenberg (2001) [see also Borkovec (2000)]. It was also considered by Nelson

(1990) in the context of the IGARCH model and by Cline and Pu (2004) in a non-parametric framework. The Lyapunov exponent was also used as a stability measure in application to the Vector Autoregressive VAR model by Dechert and Gencay (1992)[see also Lebaron (1994) for introduction to chaos]. The third stability measure is more conservative, and based on the local second-order stationarity condition of the process. It can be easily tested, by using its asymptotic normal distribution, derived in the paper. These stability measures are computed for Tether prices by plugging in the values of locally estimated parameters and compared with the stability measures of other stablecoin processes.

The proposed stability measures can be computed by stablecoin issuers and regulators and used for assessing the stability of a stablecoin over time, or for comparing the stability of different stablecoins. The episodes when the stability measures exceed a critical level are particularly important for risk management. Then, stablecoin issuers and regulators could take action to avoid explosive price patterns.

Investors expecting the stability of the Tether price at 1, may be interested in the predictability of Tether deviations from the one-dollar peg, given the past Tether prices. This "inefficiency" provides arbitrage opportunities for investors, i.e. possibilities of profit-making. The forecasting tools developed in this paper offer investors a convenient approach to predicting the deviations from the peg. For the crypto market analysis, we show evidence associating the periods of instability of stablecoin prices with important events in the cryptocurrency and financial sector, which are of concern to investors, regulators, and stablecoin issuers. We believe that active daily monitoring of Tether and other stablecoin prices, based on the proposed stability measures, model estimates, and forecasts can reduce the associated risks.

The paper is organized as follows. Section 1.2 makes an introduction to the stablecoin ecosystem. Section 1.3 describes recent risks in crypto markets and their effect on price fluctuations. Section 1.4 documents the local dynamic analysis of Tether prices. Section 1.5 discusses the modeling approach, estimation procedures, and stability measures. Section

1.6 presents the empirical results based on the estimation and inference on the DAR(1) and tvDAR(1) models, including the sample stability measures. Section 1.7 conducts some robustness analyses and compares the tvDAR model and stability measures applied to prices of various stablecoins other than Tether. Section 1.8 concludes. Appendix A contains the proofs and technical results. Appendix B provides simulation results illustrating the finite sample densities of the parameter and stability measure estimators. Additional empirical results on the application of the tvDAR model and stability measures to other stablecoins are given in Appendix C.

1.2 Stablecoin Ecosystem

Stablecoins are a type of cryptocurrency designed to maintain a stable price and reduced volatility, compared to other cryptocurrencies such as Bitcoin and Ethereum. Conventionally, stablecoin companies peg the value of their coins to that of a physical asset such as a fiat currency or gold with the assumption that the market price of their coins will eventually stabilize, establishing equivalency with the reference asset. The strategies used to achieve price stability of stablecoins are discussed below.

There is currently no standard in place that private enterprises should comply with to qualify as a legitimate stablecoin company. This leaves stablecoin enterprises with unlimited design options to choose from to differentiate their business models. Currently, business models of stablecoin companies differ in their economic design, the quality of backing they maintain, stability assumptions they rely on, and legal protection they provide for coin holders (Catalini and de Gortar , 2021). While the underlining business models may be diverse and complex, there is interest in the elements of such models to understand their economic implications. For example, one element of interest is the mechanism stabelcoins rely on to stabilize price and another is how the responsibilities are distributed over stablecoin protocols.

There exist two alternative mechanisms used by stablecoin companies to achieve price stability. They either hold collaterals in their reserves to back the value of their coins or they adjust the supply of coins through software codes to restore the peg with the reference asset. When the market value of a cryptocurrency is backed by collaterals, the cryptocurrency is referred to as a collateralized stablecoin. Conventionally, collateralized stablecoins are split into two sub-categories including off-chain collateralized stablecoins and on-chain collateralized stablecoins. Off-chain collateralized stablecoins are backed by a set of collaterals that have an economic value outside of the blockchain. The reserves of this type of stablecoins usually consist of a fiat currency such as the US dollar for Tether or a commodity such as gold for PAX Gold. Stablecoins are labelled as on-chain collateralized if the underlying collaterals are composed of crypto assets that are created in a digital form and recorded on a distributed ledger. For instance, Dai, the largest on-chain stablecoin project, supports 18 collateral assets including not only cryptocurriencies such as Ethereum and Chainlink but also stablecoins such as Tether, USD Coin, TrueUSD and PAX dollar.

Some projects opt for developing software codes to minimize price fluctuations instead of collateralizing their coins. This type of cryptocurrencies is called an algorithmic stablecoin as they try to stabilize their price around the peg by contracting or expanding the coin supply with the help of computer algorithms embedded in their design. Terra USD was until May 2022 the only example of an algortihmic stablecoin with a market capitalization over a billion US dollar.

In terms of distribution of responsibilities, stablecoins can be categorized as centralized or decentralized. Centralized stablecoins rely on a single legal entity to maintain the price stability, to manage and protect the collaterals, and to fulfill its obligations to users. For instance, Tether Limited is the legal entity that has the authority as well as the responsibility over every Tether in the circulation. Unlike centralized stablecoins, decentralized stablecoins distribute these responsibilities within their network through smart contracts. This allows network participants to take an active role in determining the rules of the stablecoin protocol such as the set of eligible collaterals and the minimum collateral requirements. The decentralized stablecoin DAI grants users who hold its governance token Maker the right to vote on the changes to its protocol. Next subsection provides further explanation on the issuance and redemption of stablecoins.

1.2.1 Issuance and Redemption of Stablecoins

Issuance and redemption are the two fundamental market activities that determine the equilibrium quantity of a stablecoin in the market. The equilibrium quantity goes up when new coins are issued, and it goes down when the existing coins are redeemed. While the equilibrium quantity changes with issuance and redemption, the price of a stablecoin is held constant. This constant price policy is the result of the pegging strategy explained in the previous section.

Issuance and redemption of stablecoins are presumably initiated by users.[1] How these transactions are executed depends on whether stablecoin has a centralized or decentralized structure. Issuance takes place following the transfer of funds by a user to the stablecoin enterprise. These funds are deposited into banking accounts of a custodian for centralized stablecoins or into a cryptographical vault for decentralized ones. For example, an individual or a business who wants to buy Tether should transfer the funds, specifically the US dollar, to Tether Limited's accounts at Cathay Bank and Hwatai Bank in Taiwan. The collection of these funds constitutes the reserves of the stablecoin enterprise, and they are meant to be kept as a collateral to back the value of every coin in the circulation. Once the funds are successfully deposited, stablecoins are issued through smart contracts and credited into the user's wallet. For centralized stablecoins, it is the issuer or the agent that authorizes the issuance of the coins whereas it is done automatically by the blockchain technology for decentralized stablecoins.

Redemption of stablecoins is also initiated by a user but the difference is that the trans-

[1] See Griffin and Shams (2020) for further discussion on whether Tether issuances are supply-driven or demand-driven.

actions take place in the reverse order. To redeem stablecoins, users place an order on the blockchain to exchange their stablecoins for the collateral. Upon the order, the stablecoin enterprise becomes obliged to withdraw the stablecoins from circulation and give the user the corresponding amount of collateral in return. The stablecoins that the user redeems are destroyed subsequently from the protocol. Stablecoin projects pledge in their whitepapers that their coins are 100% redeemable and redemptions can be performed any time users want at the predetermined price. Hence, one can argue that redeemability becomes the liability of stablecoin enterprises and plays a key role in the sustainability of their projects. It is the issuer that is liable to users for redemptions in centralized stablecoins. Decentralized stablecoins, on the other hand, have no single legal entity that shoulders the responsibility. It is the whole network that is responsible for undertaking redemptions.

1.2.2 Market Price of Stablecoins

Stablecoin prices are usually fluctuating around the target value. While their value in terms of the reference asset is fixed during issuance and redemptions, they are often traded at a premium or a discount on the exchange platforms.[2] This splits the market for stablecoins between the primary market and the secondary market, which can be considered analogous to the market for traditional securities. The primary market for stablecoins is where stablecoins are created (issuance) or destroyed (redemption) at the fixed exchange rate predetermined by the stablecoin initiative. The secondary market is where the users trade stablecoins within and across cryptocurrency exchanges such as Binance, Coinbase Exchange, Kraken and Bitfinex. The price of stablecoins in the secondary market is determined as a result of the market activities. According to Griffin and Shams (2020), the secondary market activities account for most of the aggregate Tether flow from 2014 to 2018. Hence, the price of Tether and other stablecoins in the secondary market can be considered as the effective rate at which individuals or businesses can buy and sell stablecoins on a day-to-day basis.

[2]See Lyons and Viswanath-Natraj (2020) for the detailed analysis of premium and discount on stablecoin prices.

1.3 Risks in Crypto Markets and Price Fluctuations

The extreme price volatility of some cryptocurrencies and increased adoption rates have raised concerns about their potential risk to financial stability. The financial stability review of the International Monetary Fund (2022) showed evidence of co-movement between the crypto and equity markets and of transmission of shocks to the financial system. The IMF observed that spillovers from Bitcoin's price volatility to the S&P 500 and MSCI emerging markets indices have risen by about 12 to 16 percentage points since late 2019. During the same period, the spillovers from Tether to these indices have increased by about 4 to 6 percentage points. The ECB (2022) financial stability review noted that the development of the crypto-asset ecosystem and the increased involvement of institutional investors posit the need to understand potential systemic risks that could be transmitted to the traditional financial sector.

Similar concerns were conveyed by the financial policy committee of the Bank of England (2022) in their report on the role of crypto assets and decentralized finance on financial stability. It pointed out that there are various types of risks to financial stability. More specifically, there are risks to systemic financial institutions, which can result from financial losses and operational disruption due to their involvement in facilitating client trading in crypto-asset derivatives. There also exist risks to core financial markets due to the increased portion of crypto assets in institutional investors' portfolios. Other risks arise when public confidence in the payment system is undermined due to the failure of an important stablecoin issuer to meet its obligations. The report also points out the risk to the real economy resulting from households and businesses holding significant ownership in crypto assets financed with debt.

On January 3, 2023, the US Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency issued a joint statement on crypto-asset risks to banking organizations. The Fed drew their attention to various risks,[3] including the risk of fraud and scams among market participants, the risk of misleading representations and disclosures by crypto-asset issuers, and the risk associated with the lack of maturity and robustness in the management and governance practices of the crypto-asset industry.

Li and Mayer (2022) noted that the introduction of stablecoins was comparable to "the unregulated creation of safe assets to meet agents' transactional demands" known as shadow banking. Unlike stablecoin issuers, shadow banks must provide credit guarantees to protect them from insolvency. The issuers of stablecoins designed to be on par with a fiat currency use reserves to sell or buy the currency to achieve its price stability. This mechanism helps stablecoin companies, like Tether Limited, to maintain the market value of their coins and protects them against the highly volatile nature of cryptocurrency markets. It resembles a fixed exchange rate regime currently implemented in Panama, Qatar, and Saudi Arabia. In the fixed exchange rate regime however, the central bank can also use its foreign reserves to buy or sell the domestic currency to maintain the fixed parity with the currency peg. When the reserve system fails, the domestic currency can be devaluated. Lyons and Viswanath-Natraj (2020) commented that contrary to central banks, stablecoin issuers do not have any macroeconomic policy instruments. For example, stablecoin companies cannot use the interest rate or devaluation policy to control the exchange rate.

The stabilization strategies for stablecoins are not always successful. As shown in Section 3.1, Tether price tends to be affected by various events in the crypto world, leading to deviations from the peg. Section 3 below presents further empirical evidence based on the dynamic analysis of the data.

1.4 Tether Dynamics

Our empirical analysis of Tether is applied to publicly available data on daily closing rates of Tether against US Dollar between November 9, 2017, and July 31, 2021, recorded by

[3] https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20230103a1.pdf

coinmarketcap.com and supplied by Yahoo Finance. The sample size consists of $T = 1361$ daily closing prices.

On July 2021, Tether was the largest stablecoin in terms of market capitalization with the market value around 64.3 billion USD. The second largest stablecoin was USD Coin, with a market value lower than 25 billion USD. Tether has been widely used by crypto traders to transfer investments between different cryptocurrencies or move investments into or out of fiat currencies.

Figure 1.1 displays the evolution of daily closing rates of Tether against the US Dollar over the sampling period. We observe episodes of explosive dynamics mixed with more stable periods as well as the convergence of the process to a smooth path with values close to 1 at the end of the sampling period. The convergence of Tether towards the peg and its reduced range after year 2021 occur in conjunction with increased volume levels displayed in Figure 1.2.

Figure 1.1: Tether/USD daily closing rates

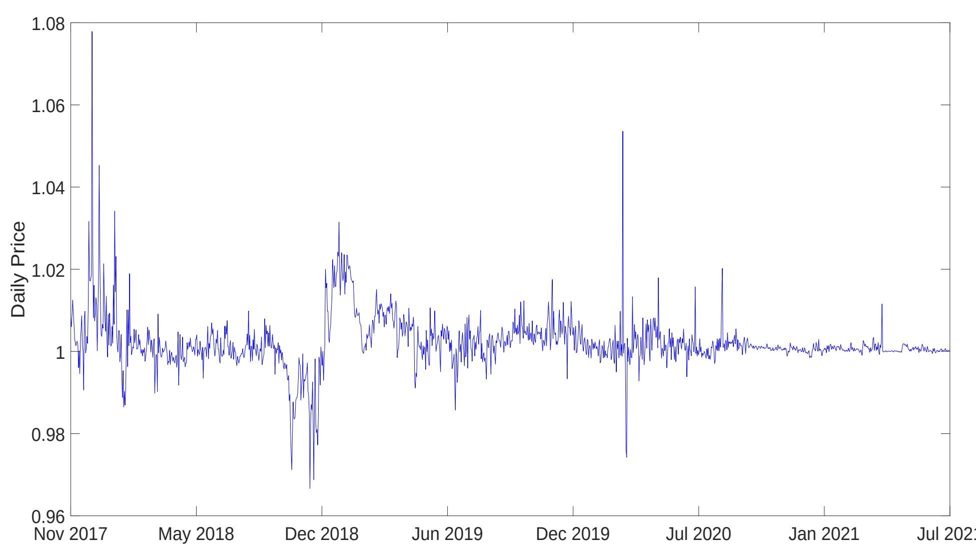

Note: This figure presents the dynamics of the daily closing price of Tether in US dollars between 11/09/2017 and 31/07/2021.

During the sampling period, the lowest and highest Tether/USD rates were 0.9666 and 1.0779, respectively. Although the associated deviations of 0.0334 and 0.0779 from one US dollar parity seem small, they can result in significant arbitrage opportunities for investors holding large positions in Tether. While the average price over this period is 1.0022 and is close to one as expected, the volatility around the mean is 0.066, yielding a coefficient of variation of 25%. The observed deviations from the peg can result from changes in demand, delayed intervention by Tether Limited to maintain price stability, and the effects of the financial environment.

The fluctuation of the Tether/USD rate around the one-dollar peg can be compared with the European snake in the tunnel currency system created in April 1972. To increase the convergence among the different currencies in the European Economic Community (EEC), a single currency band was created, within which all the EEC currencies could fluctuate without deviating too much from the peg. The peg was defined first with respect to gold and, later on, the US dollar. More details on this system can be found in Day (1976). To achieve stability around the peg, European central banks had to use their reserves to intervene by buying or selling individual currencies. The system was difficult to sustain, and several countries ultimately left the agreement. The difficulties in maintaining the snake in the tunnel currency system speak to the challenge of maintaining stablecoin prices around its peg using the reserves only. To understand the price movements of Tether, we first discuss the factors that affect its demand during the sampling period.

Figure 1.2 displays the daily volume of Tether measured as the value of daily transactions in US Dollars. Tether's daily volumes increased drastically from less than $10 billion in 2018 to as high as $279 billion in 2021. While the daily volume has been increasing almost steadily in years 2018 and 2019, that growth exhibited a rather volatile pattern in years 2020 and 2021. In the first half of 2021, the daily volume series was showing a positive trend. In early 2021, the daily volume of Tether more than doubled in just a few months and reached a peak of 99.3 billion USD on March 13, a day after the infamous "Crypto Black Thursday". The

Figure 1.2: Time varying volume of Tether in US $

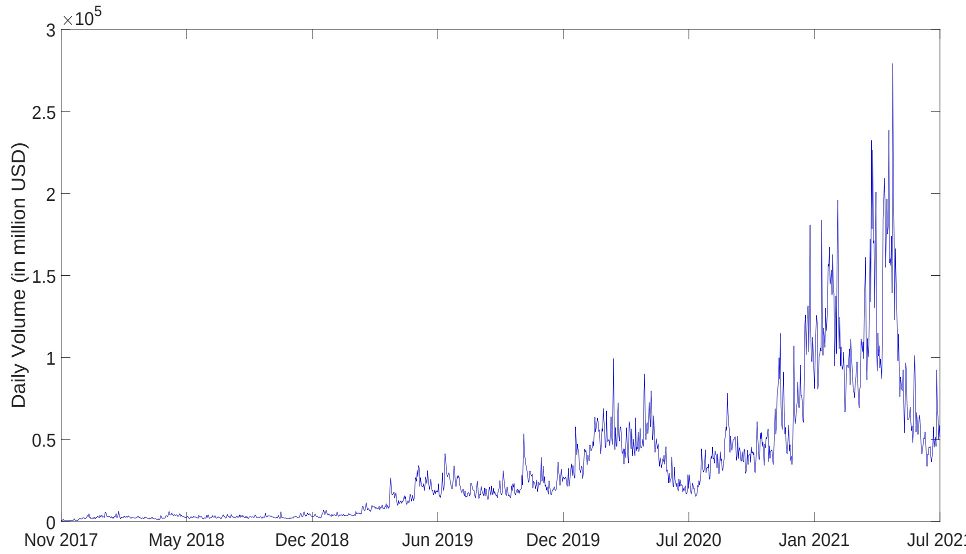

Note: This figure shows the total value of daily transactions of Tether in US dollars between 11/09/2017 and 31/07/2021.

highest daily volume of $279 billion in transactions was recorded on May 19, 2021, when the Chinese government cracked down on its domestic cryptocurrency market. That ended the upward trend and the daily volume plunged to as low as $33.7 billion in July 2021. After that decrease between May and July 2021, the volume of Tether returned to its pre-pandemic level. In the year 2021 alone, it varied roughly between $15.4 billion and $279 billion US Dollars.

The high levels of daily volume of Tether in the years 2020 and 2021 could be explained by an increasing demand from investors as the market for cryptocurrencies grew substantially during the initial stages of the COVID-19 pandemic.

To further explore Tether price dynamics and determine the modeling strategy, the next Section examines the evolution of its local mean and variance. We study to what extent the effects of idiosyncratic market events influence the local mean and variance of Tether prices, and how Tether's past and current prices are related.

1.4.1 Local Analysis of Tether Price

This subsection analyzes the local dynamics of Tether price series. We examine its local means, variances, and autocorrelations.

1.4.1.1 Local mean and variance

We can identify local patterns in the Tether/USD rate series by considering time-varying descriptive statistics computed by rolling over a window of 50 days. The locally estimated marginal mean μ_t and variance σ_t^2 are displayed in panels a) and b) of Figure 1.3 along with the 95% confidence interval (CI) for the local mean[4].

Figure 1.3: Local mean and variance of Tether prices

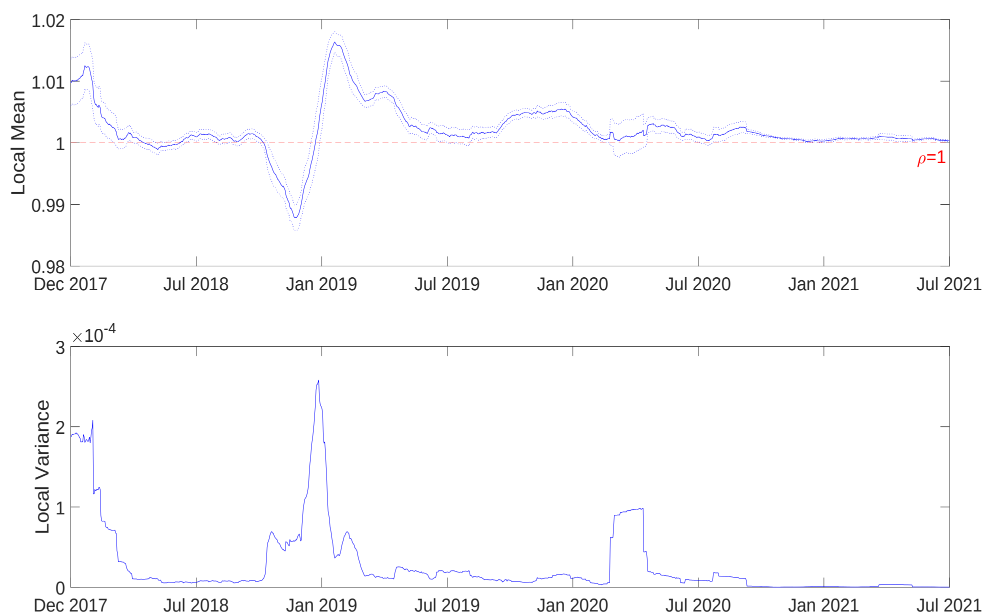

Note: The top panel shows the local mean of Tether price series (solid line) and CI at 95% (dashed line). The bottom panel shows the local variance. Estimation by rolling with window=50 obs.

The top panel of Figure 1.3 shows the local mean of price series and illustrates its evolution

[4]The lower and upper bounds of the confidence interval at level 95% are calculated under the identical and independent distribution (i.i.d.) assumption as $\hat{\mu}_t \mp 1.96\frac{\hat{\sigma}_t}{\sqrt{n}}$ for each window of n days.

over the sampling period. We observe that the local mean of Tether varies over time and Tether displays local trends. Especially, in the first half of the sampling period, we observe several positive and negative local trends. The local trends at the end of 2018/beginning of 2019 reflect the behavior of Tether price series documented in Figure 1.1 and akin to the pattern of financial bubbles observed in stock prices (rational stochastic bubbles as in Blanchard and Watson (1982) or intrinsic bubbles as in Froot and Obstfeld (1991)). The local trends become weaker towards the end of the sampling period when the local mean becomes more stable and close to the target value of 1. Overall, the periods when the target value of 1 falls within the 95% confidence intervals of local mean are: April 19, 2018 to May 5, 2018, May 9, 2018 to June 16, 2018, October 10, 2018 to October 17, 2018, January 2, 2019 to January 3, 2019, and April 2, 2020 to May 1, 2020.

The local variance of price series is plotted in the bottom panel of Figure 1.3. It shows variation over time, which is higher in the first half of the sampling period. For example, in 2018, Tether had a period of high volatility, which lasted from January to March, followed by a period of low volatility from April to November. Tether is less volatile during the second half of the sampling period as the local variance takes smaller values except for a short period of increased volatility between mid-March and early May in 2020.

Although the rolling window estimators help detect local trends, the results need to be interpreted with caution. For a window size of n days, the first n-1 observations in the dataset are eliminated due to rolling. In addition, using a longer rolling window (e.g., 100 days) may over-smooth the changes in the mean and variance compared to a shorter window (e.g., 50 days). Therefore, we use the window of 50 days for further computations.[5]

The distributional changes in Tether also concern the range and quantiles of the series. Overall, we observe that:

1. The local mean is changing over time and is close to 1 between April and June 2018,

[5] Also, note that n=50 is large enough to estimate parameters within each window consistently. Furthermore, n=50 divided by the sample size of T=1361 is the bandwidth $b_T = n/T = 0.0367$ in our local analysis, which will be discussed later. In the literature, an optimal choice should satisfy $Tb_T^3 = o(1)$. In our case, we have $Tb_T^3 = 0.0675$, which is relatively small. See page 1039 of Dahlhaus et al. (2019) for more details.

and after January 2021.

2. The variance is time varying and diminishes over time.

In Section 3.2, we identify a series of events that are closely related to Tether, and help explain the changes observed in the local statistics.

1.4.2 Event Analysis

The dynamics of Tether are influenced by various events, which can be associated with the episodes of distinct patterns in its local mean and variance.

Figure 1.4 shows the evolution of local mean and variance of Tether prices along with the series of events. A total of 12 events are marked on the graph including 8 events affecting the local mean and 4 events influencing the local variance.

In 2018, the local mean of Tether prices decreased for most of the year, except for a brief recovery period between early May and late September. Indeed, the year 2018 was a very tumultuous year for Tether Limited and its business partners. Tether Limited was scrutinized by the media and scholars for its reserves quality and close ties to the cryptocurrency exchange Bitfinex. Moreover, Tether was accused of holding insufficient reserves to back all of its coins in circulation and manipulating the price of Bitcoin by pumping unbacked supply of Tether into the market through Bitfinex to buy Bitcoins. Amid these controversies, the decrease of the local mean of Tether for most of year 2018 could be linked directly or indirectly to a shift in investor sentiment towards Tether.

There was a brief recovery period between early May and late September 2018 when the owners of Tether Limited showed their willingness to address the investors' concerns about the accountability of their business and that of Bitfinex. In particular, on May 7, 2018, they hired Peter Warrack, an anti-money laundering specialist, previously employed at RBC Royal Bank, as the Chief Compliance Officer of the cryptocurrency exchange Bitfinex. Upon this news, Tether enjoyed a recovery period when its local mean increased above the one-dollar peg. Nevertheless, the local mean of Tether started to decrease again in September and reached its lowest level in November 2018. That decrease of Tether prices could be associated with the introduction of USD Coin (USDC) on September 26, 2018. USD Coin, is also designed to be on par with the US dollars, and is similar to Tether despite it claims to offer a better transparency of its business activities.

In 2019 Tether made up quickly for the losses. The local mean increased from its lowest level in November 2018 to its highest level in January 2019 in less than two months. This increase in the local mean of Tether prices could reflect a positive reaction of investors to Bloomberg's announcement that contrary to the allegations, Tether Limited was holding sufficient reserves to back its coins in the circulation.[6] However, this positive attitude towards Tether did not last long, as the local mean started to decline again after January 2019. The downward pressure on Tether's price continued until August 2019 despite Tether Limited's efforts to be more transparent about the composition of its reserves. According to coindesk.com,[7] in mid-February 2019 Tether announced on its website that while every Tether is always 100% backed by their reserves, these reserves are not necessarily composed of fiat currency only, as they claimed before, but also include cash equivalents, other assets, and receivables from loans. This update, however, did not stop the New York State Attorney General (NYSAG) on April 24, 2019, from suing Bitfinex and Tether Limited based on an ongoing fraud.[8] After the decrease in local mean, an upward local trend becomes noticeable around September 24, 2019, when iFinex Inc, the parent company of Bitfinex and Tether Limited, won a motion in the court against NYSAG giving them a non-disclosure right to documents related to its business activities.[9] Following this court decision, the local mean of Tether diverged from the peg and kept increasing until the end of 2019, although at a slower rate.

[6] https://www.bloomberg.com/news/articles/2018-12-18/crypto-mystery-clues-suggest-tether-has-the-billions-it-promised

[7] https://www.coindesk.com/markets/2019/03/14/tether-says-its-usdt-stablecoin-may-not-be-backed-by-fiat-alone/

[8] https://ag.ny.gov/press-release/2019/attorney-general-james-announces-court-order-against-crypto-currency-company

[9] https://www.forbes.com/sites/michaeldelcastillo/2019/09/24/bitfinex-and-tether-win-appeal-from-new-york-supreme-court-in-900-million-case/?sh=7c8bf41132bc

Figure 1.4: Important events and the local mean and variance of Tether prices

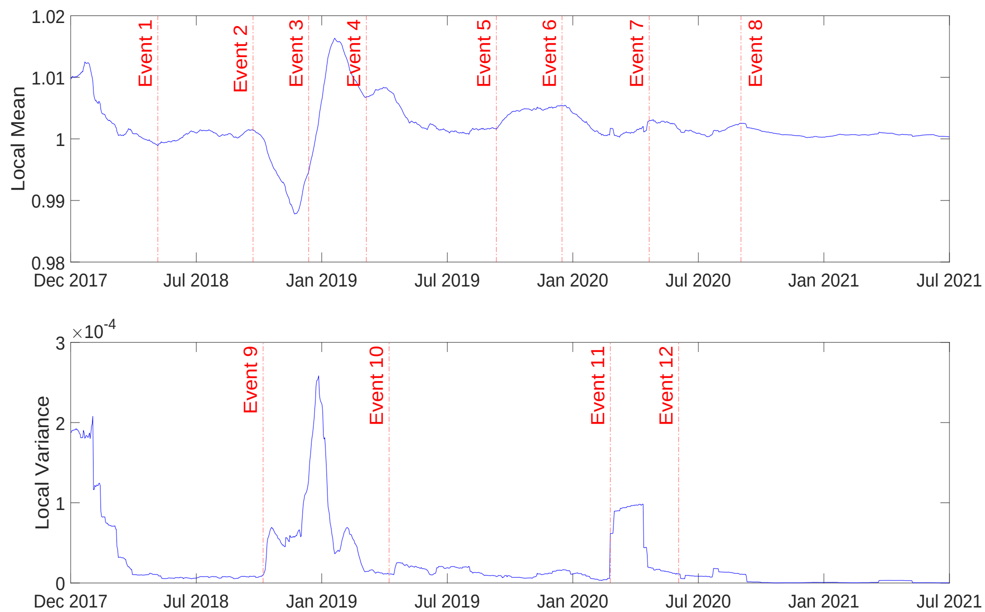

Note: This note provides a description of the events.
Event 1: Peter Warrack was hired by Bitfinex as the Chief Compliance Officer on May 7, 2018.
Event 2: USD Coin was launched on September 26, 2018.
Event 3: Bloomberg suggested on December 12, 2018 that Tether could be fully backed.
Event 4: On March 14, 2019, Tether made changes to its backing policy on its official website.
Event 5: Bitfinex won a motion in the New York Supreme Court to delay submission of its business documents.
Event 6: WHO made an announcement on Twitter on December 31, 2019 to acknowledge the cases of pneumonia in Wuhan, China.
Event 7: Bitcoin shed over 10% of its value in a matter of minutes on May 9, 2020, which was followed by an outage in the cryptocurrency exchange Coinbase.
Event 8: S&P Cryptocurrency Broad Digital Market (BDM) Index increased more than fivefold between late September 2020 and May 2021.
Event 9: Bitfinex "temporarily paused" EUR, USD, JPY, and GBP wire deposits on October 11, 2018.
Event 10: Tether went live on the Tron network on April 17, 2019.
Event 11: Black Thursday; Bitcoin's price reduced by around 50% in less than a day on March 12, 2020.
Event 12: On June 22, 2020, Tether announced on Twitter that they would implement a chain swap for a sizable amount of USDT from Tron TRC20 to ERC20 protocol on June 29th.

On December 31, 2019, the World Health Organization (WHO) announced on its official Twitter account that they were informed of cases of pneumonia of unknown cause in Wuhan City, China.[10] With this announcement, WHO acknowledged the problem of an epidemic in China, which would soon turn into a global health and economic crisis of COVID-19 pandemic. Subsequently, the local mean of Tether started to decline cancelling out the gains from the last quarter of 2019.

Tether price was relatively stable over the period 2020-2021. Improved stability is also observed in other frequently traded stablecoins such as USD Coin, Binance USD, True USD, Pax Dollar starting from July 2020, and in Dai starting from December 2020.

This period of relative stability overlaps with the period of bullish run in the cryptocurrency market and increased Bitcoin and Ethereum prices in US Dollars. The cryptocurrency market indices such as the S&P Cryptocurrency Broad Digital Market (BDM) Index recorded a more than fivefold increase between September 2020 and May 2021.[11] The strong demand for cryptocurrencies also benefited the stablecoin companies as the total market capitalization of the top 10 stablecoins went up from approximately $20 billion in September 2020 to slightly over $100 billion in July 2021.[12]

The stability of all stablecoins in the period 2020-2021 is depicted as relative because of some impactful interruptions. The price of Bitcoin exceeded 60,000 dollars for the first time on March 13, 2021, and fluctuated around this peak level until the cryptocurrency market crash on April 17, 2021. Investors took advantage of higher Bitcoin prices in March and traded Bitcoin for stablecoins. Consequently, this could have led to larger fluctuations in stablecoin prices, we observe in all stablecoins from mid-March onwards as well as the spike on April 17, 2021. On April 17, 2021, the cryptocurrency market shrunk by over $300 billion in less than 24 hours.[13] While the waves of sell-offs caused the price of Bitcoin to plummet

[10] see https://twitter.com/who/status/1213795226072109058?lang=en for the original tweet from WHO

[11] https://www.spglobal.com/spdji/en/indices/digital-assets/sp-cryptocurrency-broad-digital-market-index/#overview

[12] https://www.statista.com/statistics/1255835/stablecoin-market-capitalization/

[13] https://www.forbes.com/sites/jonathanponciano/2021/04/18/crypto-flash-crash-wiped-out-300-billion-in-less-than-24-hours-spurring-massive-bitcoin-liquidations/?sh=7d60735b2c89

by 10.5%, the price of all stablecoins increased simultaneously. This can be considered as evidence that stablecoins provide hedging opportunities for cryptocurrency investors against Bitcoin's volatility [see, e.g., Wang et al. (2020)]. However, one could also argue that the risk-mitigating properties of stablecoins are closely linked to the comovements between the prices of stablecoins and Bitcoin. For example, stablecoins showed resilience against another market crash in mid-May 2021. On May 19, 2021, the Chinese government announced that the banks in China are banned from providing cryptocurrency services to their clients.[14] The market reacted to this news immediately and Bitcoin lost 30% of its value over the course of the day. During that crash, all stablecoins except for Terra USD managed to keep their prices stable around the one-dollar peg.

1.4.3 Persistence in Tether Series

Let us now examine the serial correlation in Tether prices to reveal the impact of past prices on the current one. The autocorrelation function (ACF) of Tether is computed from the entire sampling period 2017-2021 as well as from each calendar year separately. Figure 1.5 presents the computed ACF functions: the ACF over the entire period (panel (a)), and in years 2017-2021 in panels (b) to (f), consecutively.

The ACF calculated from the entire sample exhibits long-range persistence. In contrast, the subperiod analysis reveals that persistence in Tether prices is strong in the years up to and including 2019 [15], whereas the series has short memory in 2020 and 2021.

This finding, combined with the results from the previous section suggests that the long-range persistence in Tether coincides with the periods characterized by strong local trends and high volatility as documented in Figure 1.3. When the variation is small and the local mean is stabilized around the one-dollar peg as in 2020 and 2021, Tether tends to display shorter memory.

The empirical results show that an analysis of Tether based on the global statistics

[14] https://www.theguardian.com/technology/2021/may/19/bitcoin-falls-30-after-china-crackdown

[15] In 2017, there are only 53 observations, which could be the reason why persistence appears weak.

Figure 1.5: ACF: Autocorrelation functions of Tether prices over different periods

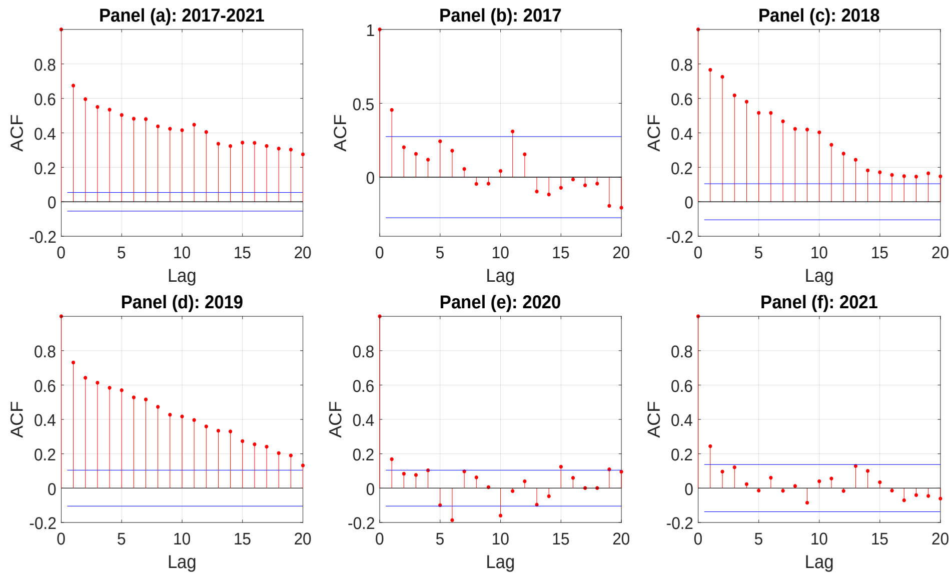

Note: This figure presents the autocorrelation functions (ACF) of Tehter prices computed from the entire period (panel (a)) and in years 2017-2021 in panels (b) to (f), consecutively. The dashed lines represent the 95% CI of statistical significance.

computed from the whole sample would provide unreliable results, especially concerning the serial correlation of prices. For example, serial correlation estimated from Tether prices recorded in 2021 have different values and range compared to 2018-2019.

Let us now focus on the autocorrelation values. In particular, the autocorrelation at lag one of Tether prices can also be estimated from the autoregressive coefficient of an autoregressive of order 1 (AR(1)) model. We first consider the autoregressive coefficient estimated from the AR(1) model fitted to the whole sample of demeaned Tether prices $x_t = y_t - \mu$ obtained by subtracting the global mean of 1.0022:

$$x_t = \rho x_{t-1} + \sigma_e e_t, \tag{1.4.1}$$

where e_t is assumed to be a white noise with mean 0 and variance 1. The AR(1) process is stationary when $|\rho| < 1$ and nonstationary and explosive when $\rho = 1$. Model (1.4.1) is estimated globally by the OLS, or equivalently by maximizing the Normality-based Maximum

Likelihood (ML) as follows:

$$\hat{\theta}_T = Argmax_\theta \sum_{t=1}^{T} l(x_t|x_{t-1};\theta) = Argmax_\theta \sum_{t=2}^{T} -\frac{1}{2}\left(\log\left(2\pi\sigma_e^2\right) + \frac{(x_t - \rho x_{t-1})^2}{\sigma_e^2}\right)$$

where $\theta = (\rho, \sigma_e^2)$. The estimated parameter values are $\hat{\rho}_T = 0.674$ and $\hat{\sigma}_{e,T}^2 = 0.545$. Next, model (1.4.1) is fitted locally and estimated by rolling with a window of length 50 and displayed in Figure 1.6.

Figure 1.6: AR(1): rolling parameter estimators $\hat{\rho}(t)$ and $\hat{\sigma}_e(t)$

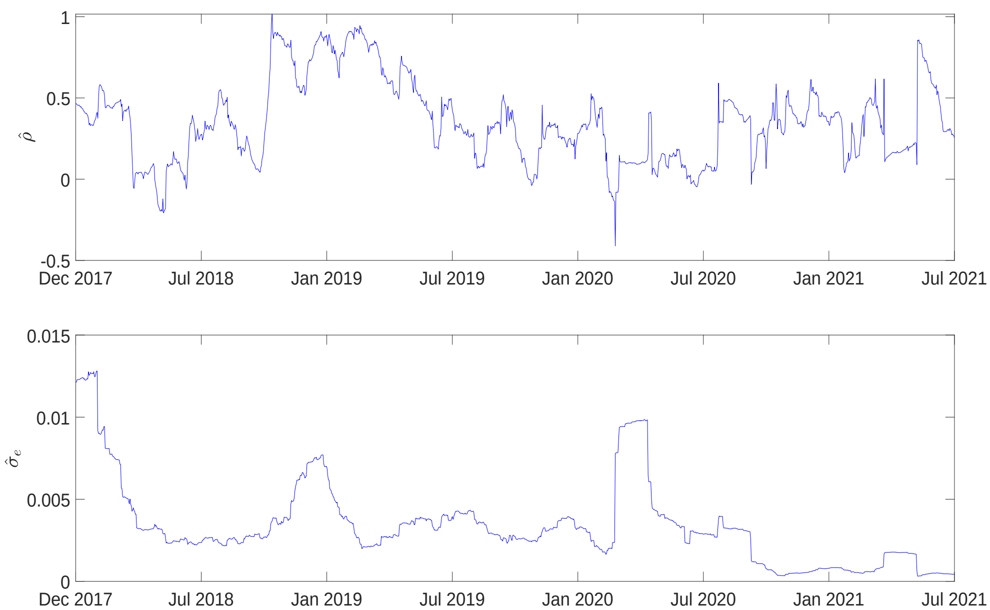

Note: The top panel presents the estimated autoregressive coefficient $\hat{\rho}(t)$ of AR(1), or autocorrelation at lag 1. The bottom panel shows the estimated conditional (error) standard deviation $\hat{\sigma}_e(t)$. Estimation by rolling, window=50 observations.

The top panel of Figure 1.6 shows the estimated autoregressive coefficient, or autocorrelation at lag 1 estimated by rolling with a window of 50 observations. We observe that the autoregressive coefficient is either close, or equal to 1 during the explosive episodes in 2018 and 2019, which violates the weak stationarity condition of the AR(1) process. This suggests that locally, the Tether prices have unit root dynamics, like the stock prices and exchange rates of traditional currency. The autoregressive coefficient tends to be lower at the end of the sampling period.

In comparison with the results presented in Section 1.4.1, estimating the autocorrelation at lag one from the autoregressive coefficient of AR(1) model gives us greater flexibility to assess the changing persistence of Tether prices as we can examine easily its evolution over short periods of 50 days. For example, the estimated values of the AR(1) coefficient ρ suggest that in 2018 the autocorrelation at lag one ranges between -0.20 and 1.02, whereas in Section 1.4.1, the ACF at lag 1 was estimated to be slightly over 0.76 for the entire year.

The bottom panel of Figure 1.6 shows the square root of conditional (error) variance $\hat{\sigma}_e(t) = \sqrt{\frac{1}{50}\sum_{\tau=t-49}^{t}(x_\tau - \hat{\rho}(\tau)\, x_{\tau-1})^2}$ of the AR(1) model, estimated by rolling with a window of 50 observations. This plot exhibits periods of low and high volatility. Especially after October 2020, $\hat{\sigma}_e(t)$ decreases and becomes close to 0. This result is consistent with the period of stability in Tether prices during the cryptocurrency bull market, as described in Section 3. If the volatility was computed over the full sample we would get a constant estimate, concealing the changes in volatility over time. To accommodate this feature, we consider a time-varying volatility model, which also captures the time-varying persistence parameter while providing valid inference when it takes or exceeds the value 1.

1.4.4 Properties of Rolling Estimators

Let us now examine the properties of rolling estimators for time varying parameters written as deterministic functions of time. First, we estimated by rolling the time varying marginal mean and variance functions, hoping to approximate $m(t)$ and $\sigma^2(t)$ in a simple model

$$y_t = m(t) + \sigma(t) u_t,$$

under the simplifying assumption of Normally distributed i.i.d. process u_t with mean 0 and variance 1. Then, in finite sample, we have

$$\hat{m}(t) = \frac{1}{50}\sum_{\tau=t-49}^{t} y_\tau \sim N\left(\frac{m(t)+\cdots+m(t-49)}{50}, \frac{\sigma^2(t)+\cdots+\sigma^2(t-49)}{250}\right)$$

We see that $\hat{m}(t)$ is biased of m(t) towards an integrated mean. By the same argument it can be shown that $\hat{\sigma}^2(t)$ is biased of both $\sigma^2(t)$ and integrated variance.

Let us now consider the time varying parameters $\theta(t) = (\rho(t), \sigma_e^2(t))$ of a time varying parameter AR(1) model:

$$x_t = \rho(t)x_{t-1} + \sigma_e(t)e_t,$$

where $x_t = y_t - m(t)$. Then, the normality-based ML estimator $\hat{\theta}^*(t)$ obtained by rolling can be written as the following kernel ML estimator (see Fan et al. (1998) for the method of local kernel-weighted likelihood estimation using local polynomial fitting).

$$\begin{aligned}
\hat{\theta}^*_T(t) &= Argmax_\theta \frac{1}{50}\sum_{\tau=t-49}^{t} l(x_\tau|x_{\tau-1};\theta) \\
&= Argmax_\theta \frac{1}{50}\sum_{\tau=1}^{T} 1_{t-49\leq\tau\leq t}\, l(x_\tau|x_{\tau-1};\theta) \\
&= Argmax_\theta \sum_{\tau=1}^{T}\left[\frac{1}{50}1_{-49\leq\tau-t\leq 0}\, l(x_\tau|x_{\tau-1};\theta)\right] \\
&= Argmax_\theta \sum_{\tau=1}^{T}\frac{1}{50}\, 1_{-49/50\leq(\tau-t)/50\leq 0}\, l(x_\tau|x_{\tau-1};\theta) \\
&= Argmax_\theta \sum_{\tau=1}^{T}\frac{1}{50}\, K\left(\frac{\tau-t}{50}\right) l(x_\tau|x_{\tau-1};\theta),\ t=1,...,T
\end{aligned}$$

with the kernel $K(u) = 1_{[-1,0]}(u)$.

It is easy to see that the dimension of the parameter of interest $[\theta(1),...,\theta(T)]$ depends on the number of observations T. To circumvent this difficulty, we can replace the functional parameter $[\theta^*(1),...,\theta^*(T)]$ with $t \in N$ by an alternative functional parameter $\theta(c), c \in (0,1)$

on [0,1], such that $\theta^*(t) = \theta(t/T)$. The rolling ML estimator is such that

$$\hat{\theta}^*_T(t) = \hat{\theta}_T(t/T) = Argmax_\theta \sum_{\tau=1}^{T} \frac{T}{50} K\left(\frac{\tau/T - t/T}{50/T}\right) l(x_\tau|x_{\tau-1};\theta),\ t = 1, ..., T.$$

This formula can be extended to any value of argument $c \in [0, 1]$:

$$\hat{\theta}_T(c) = Argmax_\theta \sum_{\tau=1}^{T} \frac{T}{50} K\left(\frac{\tau/T - c}{50/T}\right) l(x_\tau|x_{\tau-1};\theta),\ c \in [0, 1].$$

Dahlhaus (2000) and Dahlhaus et al. (2019) show that under regularity conditions the functional parameters $\theta(c)$ of a locally stationary process can be consistently estimated. Instead of considering the time varying parameters in calendar time, we have defined the functional parameters in a deformed time $t \to t/T$. The functional parameter is now independent of the number of observations T, while the effect of T is introduced by considering a triangular array approach. This leads to a sequence of models indexed by T:

$$x_{t,T} = \rho(t/T)x_{t-1,T} + \sigma_e(t/T)e_{t,T},$$

where $x_{t,T} = y_{t,T} - m(t/T)$.

This approach motivates our modelling approach presented in the next section.

Overall, the results in Section 1.4.1 highlight the existence of a time-varying pattern in the mean and volatility of Tether prices. Section 1.4.2 documents that these changes tend to be linked to exogenous events in the crypto market and transmitted by past Tether prices. Furthermore, we find in Section 1.4.3 that the autocorrelation of Tether prices changes over time, and modeling this feature is crucial. Figure 1.6 suggests that the autoregressive coefficient of the AR(1) model estimated by rolling has the potential to replicate the time-varying autocorrelation at lag 1. However, the estimated autoregressive coefficient was shown to take values equal and close to one, which makes the inference based on the ML estimator of AR(1) model unreliable. Moreover, the AR(1) model does not incorporate the effect of past prices on volatility and cannot explain its variation over time. Hence, Section 1.5 extends this preliminary model to a Double Autoregressive (DAR) model with conditional heteroskedasticity. To ensure unbiased estimation, we follow the approach of Dahlhaus (2000) presented in Section (1.4.4) and consider functional parameters estimated from windows with more weight assigned to observations around each estimation date rather than with constant weight on all observations in a window, as it was done so far.

1.5 tvDAR(1) Model for Tether Prices and Stability Measures

To account for time-varying and past-dependent conditional mean and volatility, we introduce the Double Autoregressive (tvDAR) process of order 1 with time-varying parameters. The first part of this section recalls the constant parameter DAR model before describing the tvDAR. Next, the estimation of the model is discussed. The last part of this section presents the stability measures and their estimators.

1.5.1 Model with Constant Parameters

We consider the DAR process of order 1 for the demeaned Tether price series:

$$x_t = \phi x_{t-1} + \eta_t \sqrt{\omega + \alpha x_{t-1}^2}, \tag{1.5.1}$$

where $\omega > 0, \alpha > 0$, and $\eta_t, t = 1, ..., T$ is an independent and identically distributed (i.i.d.) sequence with mean 0 and variance 1. The parameter ϕ captures the conditional mean dependence. Parameter α represents the past dependence in the conditional variance. The model is semi-parametric and conditionally heteroskedastic.[16] Borkovec and Kluppenberg

[16]More on the models with conditional heteroscedasticity and their applications in finance can be found in Gourieroux (1997). The ML and least squares estimators for conditionally heteroscedastic DAR models

(2001) and Ling (2004) show that there exists a unique strictly stationary and ergodic solution to model (1.5.1) when the following assumptions hold:

Assumption A.1: η_t has a symmetric and continuous density with mean 0 and variance 1.
Assumption A.2: The parameter space is $\Theta = \{\theta = (\phi, \omega, \alpha) : E(ln|\phi + \eta_t\sqrt{\alpha}|) < 0$ with $|\phi| \leq \tilde{\phi}, \underline{\omega} \leq \omega \leq \tilde{\omega}, \underline{\alpha} \leq \alpha \leq \tilde{\alpha}$ where $\tilde{\phi}, \underline{\omega}, \tilde{\omega}, \underline{\alpha}, \tilde{\alpha}$ are positive constants.

Assumption A.1 is not a stringent assumption. It is satisfied in particular if $\eta_t, t = 1, ..., T$ are normally distributed. Assumption A.2 is the condition of existence and negativity of the Lyapunov exponent ensuring the existence and uniqueness of a stationary solution [Borkovec and Kluppenberg (2001)]. The region of ϕ, α that satisfy the negativity condition is displayed in Figure 1, p. 64 Ling (2004) and Figure 1, p. 191 Chen et al. (2014). It includes the cases when $\phi \geq 1$ as well as $E(x_t^2) = \infty$. The processes x_t that satisfy Assumption 2 are strictly stationary. Some of these processes are also weakly (second-order) stationary and satisfy additionally the condition $\phi^2 + \alpha < 1$ ensuring that $E(x_t^2) < \infty$. Thus, the marginal variance of those processes is finite.

When $\phi = 1$, and $E(\ln|1 + \eta_t\sqrt{\alpha}|) < 0$, the process x_t is a strictly stationary martingale process with volatility induced "mean-reversion" [Gourieroux and Jasiak (2019)]. Model (1.5.1) is non-stationary when the Lyapunov exponent is non-negative. In particular, it is nonstationary at the boundary points $(\phi, \alpha) = (\pm 1, 0)$ and nests the standard unit root models at these two points. When $\phi = 0$, the process is an ARCH(1) model. Moreover, process (x_t) is strictly stationary when x_0 is drawn from a stationary distribution.

Under assumptions A.1 and A.2, the parameter space Θ is compact and there exists a unique strictly stationary solution of the model for any $\theta \in \Theta$. In addition, we assume that:
Assumption A.3: The model is well-specified, i.e. the process satisfies equation (4.2) for the true value of parameter $\theta_0 = (\phi_0, w_0, \alpha_0)$ and the true density ψ_0 of η. The true parameter value θ_0 is an interior point in Θ.
Assumption A.4: The observed process is the unique, strictly stationary solution associ-

with thresholds are given in Zakoian (1994).

ated with (θ_0, ψ_0).

These two conditions are introduced for the identification of the model and parameter estimation.

1.5.2 Model with Time-Varying Parameters

The DAR(1) model can be extended to a time-varying parameter model by using the triangular array approach for locally stationary processes [Dahlhaus (2000) and Dahlhaus et al. (2019)]. The time varying tvDAR(1) model is written for locally demeaned observations $x_{t,T}, t = 1, ..., T$ indexed by t and T (triangular array) and defined by:

$$x_{t,T} = \phi(t/T)x_{t-1,T} + \eta_{t,T}\sqrt{\omega(t/T) + \alpha(t/T)x_{t-1,T}^2}, \tag{1.5.2}$$

where for each time T, $(\eta_{t,T})$ is a strong (i.i.d) white noise with mean zero, unit variance and a symmetric distribution invariant in T. $\phi(c), \omega(c) > 0, \alpha(c) > 0$, $c \in [0, 1]$ are deterministic functions. We assume that these functions are smooth.

Assumption a.1: The functions $\phi(.), \omega(.), \alpha(.)$, are positive, deterministic and twice differentiable on $[0, 1]$.

Moreover, the trajectories of the process have to be little responsive to small changes of the parameters, which is ensured by a Lipschitz condition. More precisely, let us consider process $x_t(c)$ defined by:

$$x_t(c) = \phi(c)x_{t-1}(c) + \eta_t(c)\sqrt{\omega(c) + \alpha(c)x_{t-1}(c)^2}, \tag{1.5.3}$$

We assume that the following condition holds:

Assumption a.2:

Let the L_q norm for $q > 0$ be denoted by $||.||_q$. Then,

i) For each $c \in (0, 1)$, process $\{x_t(c)\}$ is stationary and ergodic.

ii) $c \rightarrow x_t(c)$ is continuous for any t and $||sup_c x_t(c)||_q < \infty$.

iii) There exists $\alpha, 1 \geq \alpha > 0$ and $C_B > 0$, such that

$||x_t(c) - x_t(c')||_q < C_B|c - c'|^\alpha$ uniformly in t and $c, c' \in (0,1)$.

Under Assumptions a.1 and a.2, if T is large and t/T in a small interval $(c - \epsilon, c)$, then the parameters are almost constant over that interval and locally model (1.5.3) is close to model (1.5.2) with $\phi = \phi(c), \omega = \omega(c), \alpha = \alpha(c)$. This explains the local stationarity.

When all the observations $x_{t,T}, t = 1, ..., T$ are considered, the variation of the parameters prevents the DAR process from being globally stationary. However, it is locally stationary, if Assumption A.2 is locally satisfied, i.e. $E[ln|\phi(c) + \eta_t(c)\alpha(c)|] < 0$ for any c. This is the condition on the negativity of the local Lyapunov exponent.

1.5.3 Estimation

1.5.3.1 Estimation of the Model with Constant Parameters

The parameter estimates of model (1.5.1) are obtained by maximizing the quasi-likelihood objective function written for normally distributed η_t:

$$L_T(\theta) = -\frac{1}{2}\sum_{t=2}^{T} \ln\left(\omega + \alpha x_{t-1}^2\right) - \frac{1}{2}\sum_{t=2}^{T} \frac{\left(x_t - \phi x_{t-1}\right)^2}{\left(\omega + \alpha x_{t-1}^2\right)}, \tag{1.5.4}$$

where $\theta = [\phi,\ \omega,\ \alpha]'$. The quasi-maximum likelihood (QML) estimators of model (1.5.1):

$$\hat{\theta}_T = Argmax_{\theta \in \Theta} L_T(\theta)$$

are consistent under Assumptions A.1 to A.4, and the vector of QML estimators $\hat{\theta}_T = [\hat{\phi}_T\ \hat{\omega}_T\ \hat{\alpha}_T]' \rightarrow \theta_0$ in probability, where $\theta_0 = [\phi_0\ \omega_0\ \alpha_0]'$ [Ling (2004, 2007)]. Moreover, if the following assumption:

Assumption A.5: $E(\eta_t^4) < \infty$ is satisfied, Ling and Li (2008) and Chen et al. (2014)

show that the QML estimators of θ are also asymptotically normal when $|\phi| \geq 1$[17] as well as $E(x_t^2) = \infty$. Then, we have

$$\sqrt{T}(\hat{\theta}_T - \theta_0) \to N(0, diag(\Sigma^{-1}, \kappa\Omega^{-1}))$$

where this convergence is in distribution, $\Sigma = E_0[x_{t-1}^2/(\omega_0 + \alpha_0 x_{t-1}^2)]$ and

$$\Omega = E_0\left(\frac{1}{(\omega_0 + \alpha_0 x_{t-1}^2)^2}\begin{bmatrix} 1 & x_{t-1}^2 \\ x_{t-1}^2 & x_{t-1}^4 \end{bmatrix}\right),$$

with $diag(\Sigma^{-1}, \kappa\Omega^{-1})$ denoting the block-diagonal matrix with Σ^{-1} as the upper left block and $\kappa\Omega^{-1}$ as the bottom right block and κ is the kurtosis less 1 of the distribution of η. In particular, $\kappa = 2$ when η_t is normal. These asymptotic results are valid for any true distribution of η, not necessarily a Normal distribution. The consistent estimators of Σ and Ω are

$$\hat{\Sigma}_T = \frac{1}{T-1}\sum_{t=2}^{T}[x_{t-1}^2/(\hat{\omega}_T + \hat{\alpha}_T x_{t-1}^2)],$$

$$\hat{\Omega}_T = \frac{1}{T-1}\sum_{t=2}^{T}\frac{1}{(\hat{\omega}_T + \hat{\alpha}_T x_{t-1}^2)^2}\begin{bmatrix} 1 & x_{t-1}^2 \\ x_{t-1}^2 & x_{t-1}^4 \end{bmatrix}.$$

The model residuals are defined as:

$$\hat{\eta}_{t,T} = (x_t - \hat{\phi}_T x_{t-1})/\sqrt{\hat{\omega}_T + \hat{\alpha}_T x_{t-1}^2}.$$

The model residuals $\hat{\eta}_{t,T}, t = 1, ..., T$ allow us to estimate non-parametrically the error density to verify ex-post the symmetry assumption. The parameter κ is estimated by $\hat{\kappa}_T = \frac{1}{T-1}\sum_{t=2}^{T}\hat{\eta}_{t,T}^2 - 1 = \frac{1}{T-1}\sum_{t=2}^{T}\frac{\left(x_t - \hat{\phi}x_{t-1}\right)^4}{\left(\hat{\omega}+\hat{\alpha}x_{t-1}^2\right)^2} - 1$ allowing us to accommodate the heavy tailed distribution of the stablecoin prices.

[17] The ML/OLS estimators of ϕ from a linear autoregressive AR(1) model with constant parameters are not asymptotically normally distributed when $\phi = 1$.

The distributional properties of DAR estimators in finite sample for T=50 and T=100 are examined by simulations and displayed in Figure B.3, Appendix B.

1.5.3.2 Estimation of the Model with Time-Varying Parameters

Let us consider the locally stationary tvDAR model. The dynamic model (1.5.2) of triangular arrays $x_{t,T}, t = 1, ..., T$ depends on the functional parameters $\phi(c), \omega(c) > 0, \alpha(c) > 0, c \in [0, 1]$ and the density function of the noise $\eta_{t,T}$. The estimation of $\phi(.), \omega(.), \alpha(.)$ can be done by the local-in-time QML estimators. We consider a kernel K defined on $[-1/2, 1/2]$ and bandwith b_T, $b_T > 0$ (following the notation used in Dahlhaus et al. (2019), page 1035). The local negative log-conditional quasi likelihood function is

$$L_{T,b}(c, \phi, \omega, \alpha) = \frac{1}{Tb_T} \sum_{t=2}^{T} K\left(\frac{t/T - c}{b_T}\right) \left[-\frac{1}{2} \ln\left(\omega + \alpha x_{t-1,T}^2\right) - \frac{1}{2} \frac{\left(x_{t,T} - \phi x_{t-1,T}\right)^2}{\left(\omega + \alpha x_{t-1,T}^2\right)}\right]. \tag{1.5.5}$$

Then, the local QML estimator of $\theta(c) = [\phi(c), \omega(c), \alpha(c)]$ is:

$$\hat{\theta}_{T,b}(c) = Argmax_{\theta} L_{T,b}(c, \phi, \omega, \alpha). \tag{1.5.6}$$

Under suitable regularity conditions given in [Dahlhaus et al. (2019)], this functional QML estimator $\hat{\theta}_{T,b_T}(c),\ c \in [0, 1]$ is consistent of $\theta(c), c \in [0, 1]$ and its limiting distribution is normal:

$$\sqrt{Tb_T}(\hat{\theta}_b(c) - \theta_0(c)) \to N(0, \int K(y)^2 dy\, J(c)^{-1} I(c) J(c)^{-1}),$$

where $\to$ denotes the weak convergence of processes indexed by c, $J(c)$ is the Hessian matrix and $I(c)$ is the outer product of scores, both evaluated at c. Under the symmetry assumption on the marginal distribution of η_t and for a strictly stationary and ergodic x_t, the information matrix $I(c)$ simplifies to a block diagonal matrix [Ling (2004), Remark 1].

The regularity conditions concern the functional parameter, the distribution of noise η_t, the kernel K and bandwidth b_T. They are given in Dahlhaus et al. (2019), since the DAR model is an example of a nonlinear autoregressive model discussed in Dahlhaus et al. (2019), Example 5.5, p. 1039. In particular, the bandwidth has to satisfy the conditions $b_T \to 0, Tb_T \to \infty, Tb_T^3 \to 0$ when T tends to infinity.

1.5.4 Stability Measures

The demand for stability in the cryptocurrency market became more pronounced in recent years as the total market cap of the top 10 stablecoins increased from \$5.6 billion in January 2020 to almost \$150 billion by the end of 2021.[18] This made it economically more relevant for regulators, stablecoin issuers, and investors to understand the consequences of stablecoin instability and the associated risks. Below, we discuss the stablecoins from this perspective and explain how the stability measures can be used to make informed decisions by these different interest groups.

Price stability is a key factor in the success of stablecoins in terms of gaining new functions as a currency. While some argue that stablecoins have the potential to be used as a form of mainstream payment (Kronick and Zelmer , 2023), the evidence shows that the use of stablecoins has been so far limited to the cryptocurrency market.[19] Initially, stablecoins, particularly Tether, were found to be used primarily as a store of value by crypto investors (Force, ECB Crypto-Assets Task , 2020). This is a straightforward result of stablecoin stability because investors seeking a relatively stable option to store their funds in the blockchain ecosystem have had no option but stablecoins to do so. More recently, it has been documented that in addition to being used as a store of value, stablecoins have also become the main source of liquidity for Decentralized Finance (DeFi) applications (Adachi

[18] https://www.statista.com/statistics/1255835/stablecoin-market-capitalization/

[19] The exception to this is that there have been some attempts from Circle, the issuer of USD Coin, to expand the scope of its stablecoin business beyond the cryptocurrency market by forging partnerships with American payment companies including Visa and Mastercard, and MoneyGram International Inc., a leading cross-border payment company. However, the economic size of these partnerships is yet to be documented. For more details, see Ho et al. (2022).

et al. , 2022). To understand the reason for this, let us now discuss how liquidity is supplied in DeFi. In its current state, liquidity in DeFi comes from the so-called liquidity pools which use the funds they collect from investors and pays them interest in return. When investors deposit highly volatile cryptocurrencies into liquidity pools, they are exposed to large price fluctuations, which can sometimes be so high that the fluctuation in the cryptocurrency price offsets the interest earnings. Stablecoin deposits can eliminate this risk due to improved stability. Therefore, investors, especially those wary of the high market volatility, would prefer stablecoins to alleviate risks. Because investors who use stablecoins can also have high price sensitivity, they can use a stability measure to do their due diligence while choosing the most stable alternative.

One way to assess the stability is to use a Lyapunov exponent. It measures the average logarithmic rate of separation or convergence of initially close trajectories in chaotic systems, and the sensitivity to initial conditions, in general. A negative value of the Lyapunov exponent indicates the stability of the dynamical system, while a positive value indicates chaos. The more negative the Lyapunov exponent, the more stable the system. Therefore, it is used for testing for chaos (Sprott , 2003). In this section, the Lyapunov exponent is proposed as a stability measure for Tether and other stablecoins.

In the framework of the DAR model, the Lyapunov exponent is:

$$\lambda = E(\ln|\phi + \eta\sqrt{\alpha}|).$$

The more negative the Lyapunov exponent, the less explosive the process[20] and more likely its marginal variance is finite.

The behavior of $E(\ln|\phi + \eta_t\sqrt{\alpha}|)$ as a function of ϕ, α can be examined analytically for selected densities of η, and/or simulated and illustrated graphically [see, Liu et al. (2018) for graphical illustration]. Appendix A.1 shows the derivation of the analytical formula of λ for a uniformly distributed sequence $\{\eta_t\}$ given below:

[20] A strictly stationary process can have infinite moments.

Proposition 1: If $\eta \sim U_{[-1,1]}$, the Lyapunov exponent is given by:

$$\lambda(\phi, \alpha) = \frac{1}{2\sqrt{\alpha}} \left[(|\phi| + \sqrt{\alpha}) \ln(|\phi| + \sqrt{\alpha}) - (|\phi| + \sqrt{\alpha}) - ||\phi| - \sqrt{\alpha}| \ln ||\phi| - \sqrt{\alpha}| + ||\phi| - \sqrt{\alpha}|\right]$$

Proof: See Appendix A.1.

This example clarifies that the Lyapunov exponent is a continuous function of ϕ, α, although with points of non-differentiability. It is illustrated graphically in Figure B.1, Appendix B. Moreover, it is easy to show that $E(\ln |\phi + \eta_t\sqrt{\alpha}|)$ is always an even function of ϕ, i.e. it takes the same value for ϕ and $-\phi$. To see that, consider a symmetric density function $\psi(\eta)$. Then,

$$E(\ln | - \phi + \eta_t\sqrt{\alpha}|) = \int (\ln | - \phi + \eta_t\sqrt{\alpha}|)\psi(\eta)d\eta.$$

Because $\psi(\eta)$ is symmetric, we can change the variable $\eta \to -\eta$:

$$E(ln|-\phi+\eta_t\sqrt{\alpha}|) = \int (\ln |-\phi-\eta_t\sqrt{\alpha}|)\psi(-\eta)d\eta = \int (\ln |\phi+\eta_t\sqrt{\alpha}|)\psi(\eta)d\eta = E(\ln |\phi+\eta_t\sqrt{\alpha}|).$$

This proves that the Lyapunov exponent is even in ϕ. The Lyapunov exponent can be computed by plugging-in the parameter estimates. Let us first consider the constant parameter DAR model. The following estimators can be considered:

a) Suppose that the true density function $\psi = \psi_0$ is known. Then, the estimator $\hat{\lambda}_{1,T}$ of the Lyapunov exponent is:

$$\hat{\lambda}_{1,T} = \int \ln |\hat{\phi}_T + \eta\sqrt{\hat{\alpha}_T}|\, \psi_0(\eta)\, d\eta.$$

Proposition 2: When the density $\psi(\eta) = \psi_0(\eta)$ is known, then under assumptions A.1 to A.4 and the following condition:

(A.6) $\exists\delta > 0, \text{such that} \int \sup_{\substack{\phi_0-\delta<\phi<\phi_0+\delta \\ \alpha_0-\delta<\alpha<\alpha_0+\delta}} |\ln|\phi+\eta\sqrt{\alpha}||\psi_0(\eta)d\eta < \infty$

the estimator $\hat{\lambda}_{1,T}$ converges in probability to the true value $\lambda_0 = \int \ln|\phi_0 + \eta\sqrt{\alpha}_0|\psi_0(\eta)d\eta$ of the Lyapunov exponent when $T \to \infty$.

Proof: See Appendix A.2.

b) When the density $\psi(\eta)$ is unknown, the model is semi-parametric. The Lyapunov exponent can be estimated by the estimator $\hat{\lambda}_{2,T}$ such that:

$$\hat{\lambda}_{2,T} = \frac{1}{T}\sum_{t=1}^{T} \ln|\hat{\phi}_T + \hat{\eta}_{t,T}\sqrt{\hat{\alpha}_T}|.$$

Therefore this estimator is equal to :

$$\hat{\lambda}_{2,T} = \frac{1}{T}\sum_{t=1}^{T} \ln\left|\hat{\phi}_T + \frac{x_t - \hat{\phi}_T x_{t-1}}{\sqrt{\hat{\omega}_T + \hat{\alpha}_T x_{t-1}^2}}\sqrt{\hat{\alpha}_T}\right| = \frac{1}{T}\sum_{t=1}^{T} g(x_t, x_{t-1}; \hat{\theta}_T)$$

where $g(x_t, x_{t-1}; \theta) = ln\left|\phi + \frac{x_t - \phi x_{t-1}}{\sqrt{\omega + \alpha x_{t-1}^2}}\sqrt{\alpha}\right|$. We also introduce the notation $G_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} g(x_t, x_{t-1}; \theta)$.

Proposition 3: Let us introduce the additional conditions:

(A.7) $E_{\theta_0} g(x_t, x_{t-1}; \theta) < \infty, \ \forall\theta \in \Theta$.

(A.8) Sufficient Lipschitz condition for stochastic equicontinuity: There exists a stochastic sequence B_T with $B_T = O_p(1)$ and an increasing function h from $[0, \infty)$ to $[0, \infty)$, continuous at 0 with $h(0) = 0$, such that for all $\tilde{\theta}$, $\theta \in \Theta$, $|G_T(\tilde{\theta}) - G_T(\theta)| \leq B_T h(d(\tilde{\theta}, \theta))$.

Then, under assumptions A.1-A.4 and conditions A.7 and A.8, the estimator $\hat{\lambda}_{2,T} = G_T(\hat{\theta}_T) \to G(\theta_0) = \lambda_0$ in probability.

Proof: See Appendix A.3.

The asymptotic distributions of estimators $\hat{\lambda}_{1,T}$ and $\hat{\lambda}_{2,T}$ cannot be obtained asymptotically from the Taylor series expansion because function $\phi, \alpha \to \int \ln|\phi + \eta\sqrt{\alpha}|\psi(\eta)d\eta$ does not satisfy the necessary differentiability assumption, as pointed out in Proposition 1. However, the distribution of $\hat{\lambda}_{1,T}$ and $\hat{\lambda}_{2,T}$ can be determined by simulations and used for hypothesis testing.

For illustration, we use a simulation experiment to compute and approximate the asymptotic distribution of $\hat{\lambda}_{2,T}(c)$ for a given set of fixed parameter values in the top panel of Figure B.2, Appendix B.

c) An alternative stability measure is $\xi = \phi^2 + \alpha$, which depicts the region of parameter space $\xi < 1$ where the marginal variance of x_t remains finite, so that the process is both strictly and weakly stationary. In that sense ξ is a more conservative measure of stability than the Lyapunov exponent because there is a region of parameter values ϕ, α where the condition $\xi < 1$ no longer holds, while the condition $\lambda < 0$ remains satisfied. The estimator $\hat{\xi}_T$ of ξ is:

$$\hat{\xi}_T = \hat{\phi}_T^2 + \hat{\alpha}_T.$$

Proposition 4: Under Assumptions A.1-A.5, the estimator $\hat{\xi}_T$ converges in probability to ξ_0 when $T \to \infty$ and it is asymptotically normally distributed:

$$\sqrt{T}(\hat{\xi}_T - \xi_0) \overset{\text{A}}{\sim} N(0, V_\xi),$$

where the formula of the asymptotic variance V_ξ is given in Appendix A.4.

The asymptotic variance provides the asymptotically valid standard errors that can be used to test the null hypothesis $\xi \geq 1$ using a Wald test statistic. The interpretation of measure ξ is similar to that of λ: the smaller ξ, the more stable x_t.

We use simulation experiments to illustrate the asymptotic distribution of $\xi(c)$ computed for a given set of fixed parameter values and shown in the lower panel of Figure B.2, Appendix B.

1.4.4.1 Model with time varying parameters

The Lyapunov exponent $\lambda_2(c)$ can be estimated locally by computing $\hat{\lambda}_{2,T}(c)$ from the plugged in parameter estimates and residual values of the tvDAR model. For example, the Lyapunov exponent can be estimated from a kernel-weighted formula

$$\hat{\lambda}_{2,T}(c) = \frac{1}{Tb_T}\sum_{t=1}^{T} K\left(\frac{t/T - c}{b_T}\right)(ln|\hat{\phi}(t/T) + \hat{\eta}_{t,T}\sqrt{\hat{\alpha}(t/T)}|),$$

with the Epanechnikov kernel $K(c) = \frac{3}{2}(1-(2c)^2)$ for $c \in [-1/2, 1/2]$ and $K(c) = 0$ otherwise, which satisfies the regularity condition for localizing kernel [see Assumption 2.6 in Dahlhaus et al. (2019)].

A similar approach can be used to estimate locally the measure $\xi(c) = \phi^2(c) + \alpha(c)$. The local plug-in estimator $\hat{\lambda}_{2,T}(c)$ is computed in the empirical analysis described in the next section.

1.6 Empirical Results

This section presents the estimation of the constant parameter DAR model and time-varying parameter tvDAR model. The time varying parameters tvDAR model is estimated first by rolling, which is equivalent to the use of an asymmetric rectangular kernel, as shown in the previous section. Next, the tvDAR model is estimated by using the Epanechnikov kernel which assigns higher weights to the observations close to the estimation date. The estimators of stability measures are also computed and illustrated graphically.

The estimation of the DAR(1) process with constant parameters proceeds as follows. First, we demean the price series of Tether by subtracting the global mean of 1.0022 and then fit the model (1.5.1) to the entire series of demeaned prices. We obtain the following results:

The standard deviations of the parameters are computed from the formula of asymptotic variance given in Section 1.5.3.

In the following Sections 1.6.1 and 1.6.2, the tvDAR model (1.5.2) is estimated by using

Table 1.1: DAR(1) model: Estimation of constant parameters

	DAR(1) parameters		
	ϕ	ω	α
Estimates	0.699	$9.102e^{-06}$	0.484
Standard error	(0.034)	$(2.174e^{-06})$	(0.205)

Note: This table presents the estimates of the constant parameter DAR model and their standard errors, based on the entire sample.

two types of kernels. The first one is the asymmetric rectangular kernel, equivalent to the rolling estimation over the window of 50. This window ensures good properties of the estimators, while preventing over-smoothing (Section 1.6.1). The second approach relies on the Epanechnikov kernel and produces consistent and asymptotically normally distributed parameter estimates used for hypotheses testing (Section 1.6.2).

1.6.1 Rectangular Kernel

The tvDAR(1) is estimated from the demeaned Tether price by rolling, which is a common practice in applied literature, with a window of length 50 days. This is equivalent to using an asymmetric rectangular kernel, $K(u) = 1_{(-1,0)}(u),\ b_T = 50/T$ which is computationally simple, but does not satisfy the smoothness conditions required for ensuring locally the validity of asymptotic distribution of the QML estimators. The rolling estimate and the confidence interval of functional parameter $\phi(t/T)$ are displayed in the top panel of Figure 1.7.

The middle panel shows the estimated square root of conditional variance $\sqrt{\hat{\omega}(t/T) + \hat{\alpha}(t/T)x_t^2}$ of Tether prices. The bottom panel in Figure 1.7 presents the local estimates $\hat{\lambda}_{2,T}(t/T)$ of the Lyapunov exponent computed by plugging in the local parameter estimators.

We observe that there are periods when the autoregressive coefficient is not significantly different from 1. For instance, this is the case between October and December 2018 and

Figure 1.7: tvDAR with rectangular kernel (rolling): Estimation Results

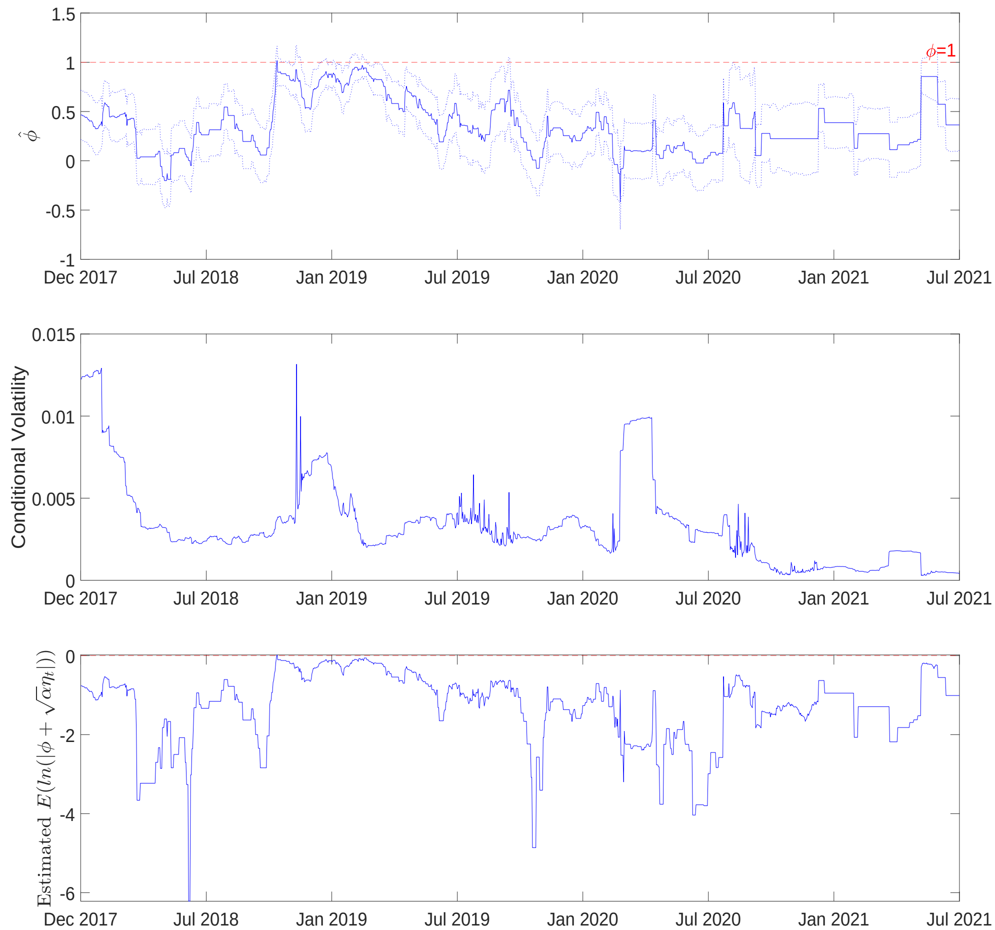

Note: The top panel presents the estimated autoregressive coefficient of tvDAR (solid line) and its 95% confidence interval (dashed line). The middle panel shows the estimated square root of conditional variance based on $\hat{\alpha}(t/T)$ and $\hat{w}(t/T)$. The bottom panel presents the Lyapunov exponent estimated from $\hat{\lambda}_{2,T}(t/T)$. Estimation by local-in-time QML with rectangular kernel, equivalent to rolling with window=50 obs.

between February and May 2019. These results from the tvDAR (1) model confirm our initial observation in Section 1.4 that the price series of Tether shows strong persistence in 2018 and 2019 whilst allowing us to pinpoint its timing more accurately. The estimated square root of conditional volatility has a pattern consistent with the local marginal variance estimator in Figure 1.3. The curve plotted in the bottom panel suggests that Assumption 2 holds and the Lyapunov exponent remains negative even at the dates of the highest persistence. The sample Lyapunov exponent is time-varying and becomes on average more negative at the end of the sampling period. This indicates that Tether achieves higher stability at the end of the sampling period.

Figure 1.8: Tether prices and one-step ahead out-of-sample forecasts

Full Sample

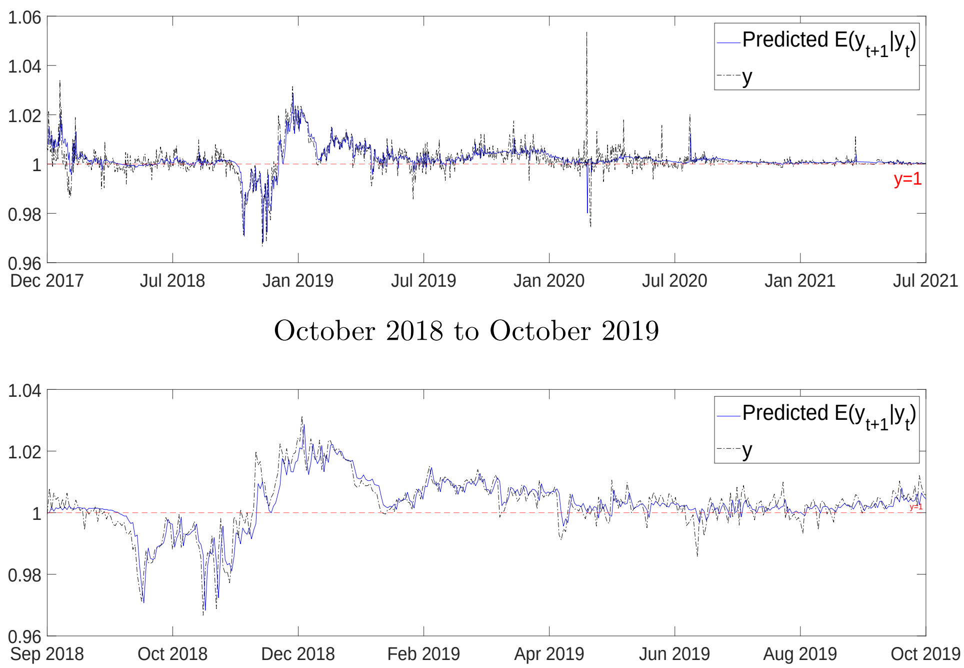

Note: This figure presents the one-step ahead out-of-sample forecasts of Tether prices (solid line). The top panel presents the forecasts over the entire period. The bottom panel shows the forecasts for the period September 2018 to October 2019. The dotted line depicts the true values of Tether prices.

The tvDAR(1) model estimated with the asymmetric rectangular kernel can be used for forecasting at short horizons, under the assumption that the parameter functions remain constant and equal to the last estimated value. Figure 1.8 presents the observed Tether price and the estimate $\hat{y}_{t+1}$ of the one-day-ahead forecast $E(y_{t+1}|y_t, \ldots)$ of the price of Tether from the rectangular kernel (rolling window=50 obs). To get $\hat{y}_{t+1}$, we compute $\hat{\phi}$ times $\hat{x}_{t+1}$ and add the local mean of Tether price estimated from data over 50 last days up to date t. Under the assumption of locally constant parameters, the 95 % asymptotically valid prediction intervals of $\hat{x}_{t+1}$ are given by $\left[\hat{x}_{t+1} \pm 1.96\sqrt{\hat{w} + \hat{\alpha} x_t^2}\right]$. Figure 1.9 shows a close match between Tether price and its best prediction based on the previous day price. In addition, the computed mean square prediction error is 1.7428×10^{-5} which is very small.

Figure 1.9: One-step ahead out-of-sample Tether price forecasts and prediction intervals

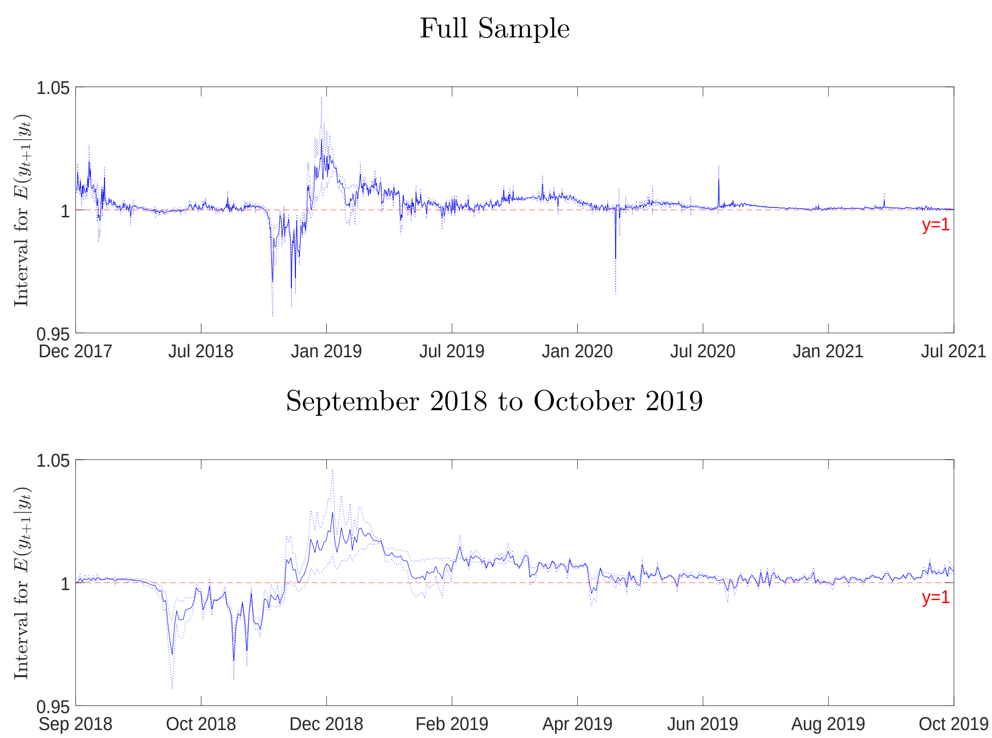

Note: This figure presents the one-step ahead out-of-sample forecasts of Tether prices (solid line) and their 95 % prediction intervals (dashed line). The top panel presents the forecasts over the entire period. The bottom panel shows the forecasts from September 2018 to October 2019.

Next, we examine the goodness of fit of the tvDAR(1) model. For this analysis, we use the rectangular kernel (window) of length 50 observations and perform the Ljung-Box test of white noise on $\hat{\eta}_t$ and $\hat{\eta}_t^2$ by rolling [see Li (1992) and Li and Mak (1994)]. This procedure will show us if the linear and nonlinear serial correlation is fully captured by the tvDAR(1) model in all subsamples, i.e. windows of 50 consecutive dates. Our results reveal that residuals $\hat{\eta}_t$ and $\hat{\eta}_t^2$ are serially uncorrelated in most of these subsamples. More precisely, we reject the null hypothesis of no serial correlation of $\hat{\eta}_t$ only in 1.52% of subsamples, while we reject the null hypothesis of no serial correlation of $\hat{\eta}_t^2$ only in 8.08% of subsamples. The results suggest that the tvDAR model captures well the linear and nonlinear serial correlation in Tether prices.

1.6.2 Epanechnikov Kernel

We now consider a symmetric Epanechnikov kernel producing consistent and asymptotically normally distributed estimates of tvDAR model for hypothesis testing. The top panel of Figure 1.10 plots the estimated time-varying autoregressive coefficient and its confidence band, while the botom panel presents the time-varying estimates of the Lyapunov exponent based on the $\hat{\lambda}_{2,T}(t/T)$ estimator. The bandwidth is $b_T = 50/T$.

The results confirm that when the estimation is performed locally, the Lyapunov exponent displayed in the second panel of Figure 1.10 is negative. Hence the critical validity condition holds at all dates. Moreover, the Lyapunov exponent is on average more negative at the end of the sampling period, confirming that Tether has achieved higher stability at the end of the sampling period.

Furthermore, the dynamics of estimated parameter ϕ in the top panel of of Figure 1.10 remains similar to that of Figure 1.7 and based on the rectangular kernel. These findings confirm that there are few dates when the autoregressive parameter is not statistically different from 1, which suggests strong persistence in Tether prices.

To detect the periods of strong persistence, we test the null hypothesis $H_0 : \phi = 1$ using the series of $\hat{\phi}(t/T)$ and its asymptotically valid 95% confidence intervals obtained from the estimation with the Epanechnikov kernel. The results are reported in Table 1.2.

When $\phi = 1$, Tether prices locally follow a stationary martingale process. Then, the

Figure 1.10: tvDAR with Epanechnikov kernel: Estimation Results

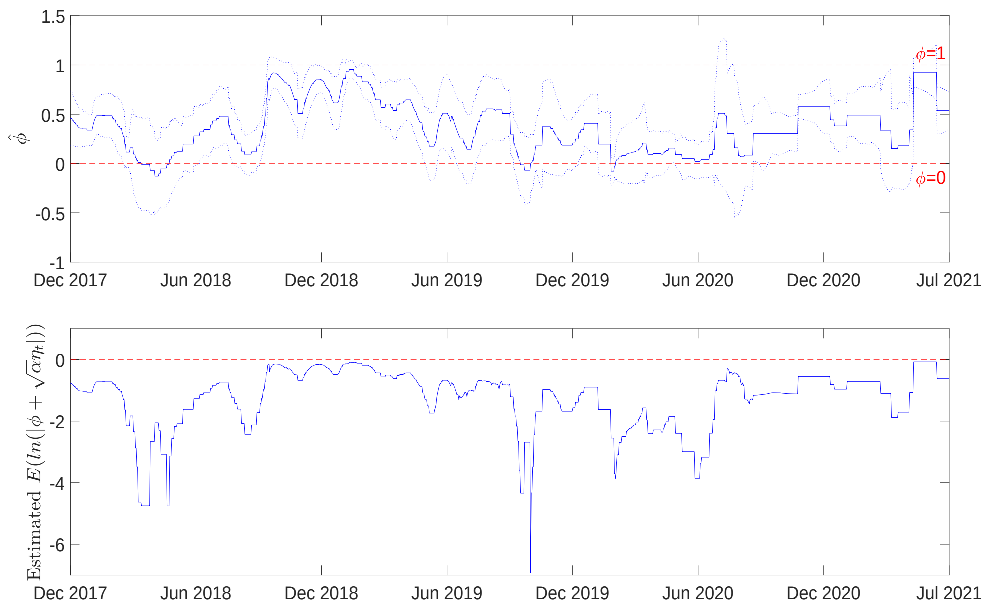

Note: The top panel presents the estimated functional autoregressive coefficient of tvDAR (solid line) and its 95% confidence interval (dashed line). The bottom panel shows the Lyapunov exponent estimated from $\hat{\lambda}_2(t/T)$. Estimation with Epanechnikov kernel and bandwidth of $50/T$.

Table 1.2: Episodes of high persistence in Tether characterized by $\phi = 1$.

Year	*From*	*To*
2018	September 24	October 29
	December 3	December 5
2019	January 11	January 30
	February 9	February 11
	February 16	February 21
2020	July 25	August 8
	August 16	August 19
2021	May 14	June 17

Note: This table reports the time intervals when the value 1 of parameter ϕ is inside the 95% confidence interval.

tvDAR(1) model implies that the price of Tether is also locally efficient because $E(x_t|x_{t-1}) = x_{t-1}$. To results in Table 1.2 are obtained by testing the null hypothesis $\phi = 1$ by using the functional parameter estimates and confidence interval values plotted in the top panel of Figure 1.10. The reported time intervals correspond to the dates when $\phi = 1$ lies inside the asymptotically valid confidence band in Figure 1.10, and the tested hypothesis cannot be rejected at 5% probability level.

Table 1.2 suggest that although Tether price is predictable most of the time, during the reported episodes it exhibits a unit root, while the series remains locally strictly stationary. As explained earlier, the tvDAR estimators remain valid and asymptotically normally distributed. We confirm that the episodes of high persistence in 2018 and 2019 were locally explosive in the sense, they overlapped with a period of high volatility observed in Figure 1.4 that started in September 2018, following the introduction of USD Coin, a stablecoin designed to maintain price equivalence to the U.S. dollar. The episodes of high persistence in 2020 and 2021 were less explosive and associated with smaller increases in Tether price volatility.

1.6.3 Test for Conditional Homoscedasticity

Let us now consider a simple specification test of the tvDAR(1) model (1.5.2) with time-varying parameters. It is based on testing for the constancy of function $\sigma(t/T)$:

$$\sigma(t/T) = \sqrt{\omega(t/T) + \alpha(t/T)x_{t-1,T}^2}, \quad t = 1, ..., T. \tag{1.6.1}$$

To test the null hypothesis $H_0 : \sigma(t/T) = \sigma_0 \; \forall t$, we consider the following test statistics introduced by Chandler and Polonik (2012) for time-varying autoregressive processes:

$$CP_T = \sup_{\alpha \in [0,1]} \sqrt{\frac{T}{\gamma(1-\gamma)}} |\hat{G}_{T,\gamma}(\alpha) - \alpha\gamma|, \tag{1.6.2}$$

where

i) $\hat{G}_{T,\gamma}(\alpha) = \frac{1}{T}\sum_{t=1}^{[\alpha T]} \mathbb{1}(\hat{\epsilon}_t^2 \geq \hat{q}_\gamma^2)$,

ii) $\hat{q}_\gamma^2 = \min\left(q^2 \geq 0 : \frac{1}{T}\sum_{t\in[aT,bT]} \mathbb{1}(\hat{\epsilon}_t^2 > q^2) \leq \gamma\right)$,

iii) $\hat{\epsilon}_t = x_t - \hat{\phi}(t/T)x_{t-1}$.

The process $\hat{G}_{T,\gamma}(\alpha)$ counts the number of squared residuals within the first $(100\times\alpha)$% of the observations that are larger than the empirical quantile of the squared residuals denoted by $\hat{q}_\gamma^2$. The series of squared residuals is constructed by computing $\hat{\epsilon}_t = x_t - \hat{\phi}(t/T)x_{t-1}$ for each period t where $\hat{\phi}(t/T)$ in our case is the corresponding DAR(1) estimate obtained in the previous section.

Chandler and Polonik (2012) show that under the null hypothesis, the test statistics CP_T in equation (1.6.2) converges asymptotically to the supremum of a Brownian bridge. This asymptotic result remains valid when the functional parameters of the model are estimated nonparametrically [see Chandler and Polonik (2012)].

We compute the test statistics CP_T for four different empirical upper γ-quantiles with $\gamma = (0.7, 0.8, 0.9)$, following the approach of Chandler and Polonik (2012). Next, for each value of γ, we calculate the expression $\sqrt{\frac{T}{\gamma(1-\gamma)}}|\hat{G}_{T,\gamma}(\alpha) - \alpha\gamma|$ for probablity levels $\alpha = 0.1, ..., 0.9$ and report the results in Table 1.3.

Table 1.3: The values of $\sqrt{\frac{T}{(\gamma(1-\gamma)}}|\hat{G}_{T,\gamma}(\alpha) - \alpha\gamma|$ for each pair of γ, α coefficients

		α								
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
	0.7	0.922	1.302	2.466	3.328	4.251	5.957	6.759	6.837	3.539
γ	0.8	0.912	1.478	2.390	3.233	3.937	5.471	6.106	6.327	3.509
	0.9	0.562	1.031	1.593	2.155	2.993	3.923	4.485	5.047	3.490

Note: This table presents the values of the argument of the CP_T statistic for different values of α and γ. The rows represent the values of γ. The columns represent the values of α.

Table 1.3 shows that regardless of the choice of γ, the largest value of the argument of CP_T is obtained for $\alpha = 0.8$. By the definition of supremum, the test statistics CP_T should satisfy $CP_T \geq \sqrt{\frac{T}{\gamma(1-\gamma)}}|\hat{G}_{T,\gamma}(\alpha) - \alpha\gamma|$ for all $\alpha \in [0,1]$. Hence, it is safe to conclude that $CP_T \geq 5.047$ is satisfied in the worst case scenario, i.e when $\gamma = 0.9$ [21]. Next, we compare this result with the critical values from the asymptotic distribution of the test statistics [22] and conclude that we can reject the null hypothesis of constant variance function at the 99% confidence level. In other words, we have provided statistical evidence in favor of the alternative that the variance function is varying over time, which consequently justifies our strategy to estimate the model with time-varying parameters.

The tvDAR model is a reliable instrument for monitoring Tether price dynamics for investors and crypto market analysts. The estimates of model parameters can be regularly updated and used for forecasting Tether prices and their volatility. Tether price volatility is predictable as well. The volatility forecasts can be easily computed too, given the estimated parameters of the volatility function and past Tether prices.

The proposed stability measures can provide a timely assessment of Tether price dynamics. They evaluate locally a combined explosive effect of the conditional mean and conditional variance. Hence, they can help assess how close Tether prices are to an explosive episode. The periods of instability are of significant concern for investors. It is also the case for regulators to whom it is essential to avoid a contagion effect that could result from a stablecoin collapse. Tether Limited and other stablecoins issuers can benefit from the proposed stability measures to improve their proactive risk management and maintain investors' confidence.

In the next section, we discuss the application of the proposed methods to other stablecoin price series. We provide further arguments supporting daily monitoring of stablecoins based

[21] Alternatively, we could fix our choice of γ and find the exact value of $\alpha \in [0,1]$ at which the expression $\sqrt{\frac{T}{\gamma(1-\gamma)}}|\hat{G}_{T,\gamma}(\alpha) - \alpha\gamma|$ attains its maximum on this fixed interval. This could slightly improve the precision of the lower bound for the test statistics. For example, when $\gamma = 0.9$, the expression attains its maximum value of 5.106 at $\alpha^* = 0.8012$.

[22] see https://homepages.ecs.vuw.ac.nz/ ray/Brownian/ for the distribution of the supremum of a Brownian Bridge.

on the proposed stability measures to prevent ultimately a crash or collapse, as was the case of Terra USD.

1.7 Robustness and Applications of the Stability Measure

This section estimates the tvDAR model and stability measures from the price series of large-cap stablecoins other than Tether and compares the estimation results. We first fit the tvDAR(1) model to daily closing price series of major stablecoins: USD Coin (USDC), Binance USD (BUSD), Dai (DAi), True USD (TUSD) and Terra USD (UST) with a market capitalization of at least 1 billion USD as of July 2021.[23] Then, we compare the time-varying autoregressive coefficient ϕ of tvDAR estimated from different stablecoins and identify the episodes of local strong persistence by testing the null hypothesis $\phi = 1$.

Next, we compute the stability measure $\hat{\xi}$, which has an asymptotic normal distribution allowing us for testing hypotheses on price instability. We can detect the explosive episodes by testing locally the null hypothesis $\xi \geq 1$.

In the last part of this Section, we discuss further how the proposed stability measure can be used to prevent the run risks.

1.7.1 Comparison of tvDAR Estimations of ϕ

Figures C.2 to C.6 in Appendix C show the estimated functional autoregressive parameters of tvDAR models fitted to the prices of USD Coin (USDC), Binance USD (BUSD), Dai (DAI), True USD (TUSD) and Terra USD (UST). These prices are observed over periods starting from their respective inception dates until July 2021. Hence, the sampling periods are of different lengths, depending on the coin and some of them are as short as 6 months.

[23] Data available online on Yahoo Finance.

The top panels of Figures C.2-C.6 display the stablecoin prices. The estimated time-varying autoregressive coefficients of tvDAR are plotted in the bottom panels of each figure along with their asymptotically valid 95% confidence intervals estimated with the Epanechnikov kernel and bandwidth of (50/T).

From these estimates, we infer that all stablecoins analyzed in this section admit locally autoregressive coefficients equal to one. We find statistical evidence of strong persistence in all stablecoins by testing the null hypothesis $H_0 : \phi = 1$. This can be done by visual inspection, i.e. by checking if value 1 is inside the reported confidence interval. Concerning the timing of the episodes of strong persistence, all stablecoins show increased persistence and some local trends around March 2021 at the end of the sampling period.

The prices of USDC in Figure C.2 display several episodes of strong persistence. For instance, we observe the first episode of strong persistence in USDC in November 2018, soon after its inception in September 2018. During that month, there is evidence of a local trend and relatively high volatility in USDC prices, in line with our observations for Tether. In the case of BUSD [Figure C.3], we fail to reject the null $\phi = 1$ on one day only in March 2021.

The behaviour of DAI is interesting too [Figure C.4]. DAI's prices became more persistent in September 2020 when its local autoregressive coefficient increased above 0.5. This on-chain collateralized stablecoin maintained its high autocorrelation until late 2020 while its price was showing strong local trends and was highly volatile. We also detect high persistence in DAI associated with weaker local trends and lower variance closer to the end of the sampling period.

TUSD and UST are stablecoins with relatively smaller market capitalizations. Figure C.5 shows that TUSD displayed strong persistence over brief episodes in every calendar year. These episodes occurred when the price of TUSD was locally more volatile, similar to what we observed previously in Tether. The value of 1 was inside the confidence interval in TUSD towards the end of the sampling period in June 2021. It was accompanied by a local trend and increased volatility, which nevertheless seem low, as compared to those at the beginning of the sampling period in 2018.

Figure C.6 shows that Terra (UST) had two long-lasting episodes of strong price persistence in 2021 from late January to early March and from May to mid-June. These two episodes were characterized by upward and downward shifts in Terra's prices and close to 1 value of the local autoregressive coefficient over about 82 days, i.e. one-third of its sampling period.

1.7.2 Comparison of the Stability Results

Let us now discuss the locally estimated stability measure $\hat{\xi}$ for the prices of stablecoins examined above. The local estimates of stability measure $\hat{\xi}$ along with their 95% confidence intervals are plotted in the bottom panel of Figure C.1 for Tether and in Figures C.7 - C.11 for other stablecoins in Appendix C. Given that the estimator $\hat{\xi}$ is asymptotically normally distributed, we can easily test the null hypothesis that a stablecoin process x_t is dynamically unstable, $H_0 : \xi \geq 1$ against the alternative where it is dynamically stable, $H_A : \xi < 1$. We use the 95% confidence intervals for ξ to perform the test. Consequently, when the upper bound of the interval is below the unity line, we reject the null of dynamic instability with a significance level of 5%.

All of the stablecoins examined in this section are pegged to the US dollar at a 1:1 ratio. In terms of design, some alternatives resemble Tether more than others. USDC, BUSD, and TUSD rely on the same principles as Tether to stabilize their prices so they can be grouped as off-chain collateralized. On the other hand, DAI is on-chain collateralized and Terra (UST) is the only algorithmic stablecoin in our sample.

Based on the stability tests, we find that there are four distinct subperiods when Tether was unstable: between October 2018 and March 2019, between August and September 2019, between August 2020 and January 2021, and finally in June 2021. Figures C.7 - C.11 show that like Tether, other stablecoins also lacked stability over certain parts of the sampling period. The good news is that in a given period, there was at least one stablecoin that maintained its stability whilst other alternatives did not. Let us now discuss the reasons for Tether lacking stability in the aforementioned subperiods and also compare its stability to that of other stablecoins.

According to our estimations, Tether was unstable for a total of 80 days between October 2018 and March 2019. This period of instability can be explained by Tether showing high persistence and increased variance both of which increase the stability measure by construction. As explained in Section 3.2, there were concerns about the quality of Tether's reserves during this period with legal investigation looming over its business in the United States, which could explain the source of instability. Meanwhile, USDC, which was a fairly new alternative back then, was equally unstable over 39 days. This means that TUSD was the only stable option for investors during this period despite its small market share.

The lack of stability in Tether between August and September 2019 could be linked to increased uncertainty over the ongoing investigation conducted by the New York Attorney General (NYAG) and described in Section 3.2. On the contrary, both alternative coins USDC and TUSD were stable during that period.

Our estimations suggest that Tether remained stable during the initial stages of the COVID-19 pandemic, until the end of August 2020 while other stablecoins did not. The instability in USDC, BUSD, DAI, and TUSD could be attributed to a pronounced increase in their local variances. In the fourth quarter of 2020, Tether (along with DAI and TUSD) was no longer stable, unlike its close substitutes USDC and BUSD.

Tether restored its stability at the beginning of 2021 and maintained it until June. Our estimations show a lack of stability in Tether for almost one month (24 days). Before this, Tether disclosed the composition of its reserves for the first time with a report[24] published on May 13, 2021. Because the report revealed that Tether invested its reserves heavily in commercial papers that are not risk-free, it might have added to the discontent among investors regarding the quality of its reserves, which could explain the lack of stability in the

[24]https://tether.to/en/tether-releases-breakdown-of-its-reserves

subsequent period. Tether would later get rid of these commercial papers and replace them with U.S. Treasury bills to address the concerns of investors.[25]

When the stability of other stablecoins in 2021 is considered, two alternatives to Tether stand out. One of them is USDC and the other one is UST. It appears that USDC performed better than Tether in 2021, as it is the only stablecoin in our sample which remarkably maintained its stability throughout the year. UST (Terra), on the other hand, is the worst-performing stablecoin with 83 days of instability out of 200 estimation days, which means that UST lacked stability about 42% of the time in 2021. Given the fact that UST would eventually collapse in the following year (due to a run from its investors), this evidence of prolonged instability in its price series could have been seen as an early warning signal of its poor performance in a period of distress.

1.7.3 Use of the Stability Measure to Prevent Run Risks

When a stablecoin price is unstable, it is failing to deliver investors the service it has promised. According to the report of the US President's Working Group (2021), a stablecoin's deviation from the expected performance could cause a run on that stablecoin. The collapse of Terra USD in May 2022, which wiped out billions of dollars from the cryptocurrency market, can be given as an example of this failure. Because of the lack of stability in Terra's price combined with the loss of investors' confidence in the coin, people rushed to sell their Terra coins until its price virtually became zero. The proposed stability measure can help stablecoin issuers analyze the overall market reaction towards the coins and address the concerns about their stability before a run becomes a reality.

The collapse of Terra USD also showed that its failure had some spillover effects on other cryptocurrencies and traditional asset prices across the world (Kronick and Zelmer , 2023). The impact on the other markets outside of the blockchain ecosystem was not concerning because Terra USD had a relatively small market cap and no direct link with the rest of the

[25]https://www.reuters.com/technology/tether-says-it-has-completely-eliminated-commercial-paper-reserves-2022-10-13/

financial system. This could raise a question as to what would have happened if any of the top 3 stablecoins including Tether, USD Coin, and Binance USD collapsed. These stablecoins pose an even higher risk to the traditional financial industry due to their reserve compositions. According to their audit reports, all three of these stablecoins hold some form of traditional assets in their reserves including corporate bonds and government bonds including U.S. Treasury bills, and Reverse Repurchase Agreements. Should investors lose their faith in the stability of these stablecoins and start selling them, the issuers would have to liquidate large quantities of reserves to keep the price of their coin stable. In an extreme case of a run, this could lead to a sharp decline in the prices of bonds, which would eventually harm the companies or the institutions that have significant quantities of these bonds on their balance sheets. Since stablecoins are currently non-regulated, there is no incentive for stablecoin companies to internalize this type of negative externalities their businesses are creating for the rest of the financial system. With the proposed stability measure, regulators can monitor stablecoins under different market conditions and determine the reasons for their instability. Going forward, this would help regulators develop a regulatory framework to make sure that stablecoins can maintain or restore their stability at all times without causing any harm to the broader financial system, especially in the case of sudden redemptions.

1.8 Concluding Remarks

We show that the distributional and dynamic properties of stablecoins have been evolving over the sampling period. We implement local analysis to detect and describe the local explosive patterns, time-varying volatility, and persistence. We model the dynamics of Tether, the largest stablecoin, and provide evidence that the tvDAR(1) model with time-varying coefficients provides a locally good fit to Tether prices and reliable short-term predictions. Our modeling strategy enables us to obtain valid inference even when the tvDAR(1) coefficient ϕ_t of the Tether price is not different from 1. We introduce three simple and testable plug-in stability measures. Two of them are based on the sample Lyapunov exponent and are computed from the parameter estimates. The tvDAR model and proposed stability measures were applied to Tether and compared with the results obtained from other large-cap stablecoin prices.

The tvDAR model can be used by investors, stablecoin issuers, regulators, and other stakeholders for predicting future deviations of Tether prices from the peg and dynamically monitoring its stability. Our results show that Tether often remains relatively stable and predictable. The periods of instability determined from the stability measures can be associated with the effects of the financial ecosystem and various events on crypto markets. Applying the proposed tools to other stablecoins provides similar evidence. An aspect not covered in this paper is the joint dynamic analysis of stablecoins. It would be insightful to jointly model the stablecoin prices and take into account their comovements. Future directions of this work include extending the proposed model to the multivariate setting and studying the dynamic interrelations between the stablecoins. Then, a joint stability measure could be developed for multivariate analysis of stablecoins dynamics.

Chapter 2

Tests for No-Arbitrage in Cryptocurrency Markets

Coauthorship Statement

This chapter is based on joint work titled "Tests for No-Arbitrage in Cryptocurrency Markets" with Dr. Antoine Djogbenou, Dr. Joann Jasiak, and Dr. Razvan Sufana. All authors contributed equally to the conceptualization, data preparation, empirical modeling, analysis, and writing of the manuscript.

A version of this chapter is currently under revise and resubmit at `Economic Modelling`.

2.1 Introduction

There is an ongoing debate concerning the efficiency of the cryptocurrency markets. Several articles have examined the market efficiency of cryptocurrencies such as Bitcoin, Ethereum, Ripple, Litecoin and EOS and provide evidence of time varying market efficiency of cryptocurrency [see, e.g., Tran and Leirvik (2020), Dhawan and Putnins (2020), Liu et al. (2019), Chu et al. (2019), Li et al. (2018)]. While the literature seems to agree that prior to 2017 the cryptocurrency markets were mostly efficient [see, e.g. Urquhart (2016), Vidal-Tomas and Ibanez (2018), Zargar and Kumar (2019)], there is no consensus regarding the period from 2017 to the present. For example, Tran and Leirvik (2020) found evidence of market efficiency over this period, while other authors conclude that the cryptocurrency markets were largely inefficient after 2017.

This paper studies the efficiency of the cryptocurrency markets by testing the no-arbitrage hypothesis in a framework similar to that in Gourieroux et al. (2010), which was developed as a unified approach for pricing bonds, stocks, currencies, and their derivatives and used the Wishart autoregressive (WAR) process (Gourieroux et al. , 2009) as a model for fundamental risk affecting all assets of interest. Wishart processes are tractable modeling tools since they are special cases of affine processes [see Duffie and Kan (1996) and (Duffie et al. , 2003)], and have been used in the financial literature to model multivariate risk, in particular stochastic covariance or correlation matrices, in various applications, such as applications to derivative pricing (Gourieroux and Sufana , 2010) and portfolio choice (Buraschi et al. , 2010).

We extend the applicability of this framework by using it as a general approach for modeling cryptocurrencies traded on multiple exchange platforms. To the best of our knowledge, this paper is the first to analyze the pricing of major cryptocurrencies and their efficiency in such a general framework. We assume that the cryptocurrency assets are driven by a common fundamental risk factor and idiosyncratic risk factors specific to each asset. The dynamics of the fundamental risk factor is modeled using a univariate WAR process, while the idiosyncratic risk factors are assumed to follow a conditional multivariate Gaussian distribution. In the univariate case, the WAR process reduces to an autoregressive gamma process (Gourieroux and Jasiak , 2006), and is a discretized version of the CIR process (Cox et al. , 1981) used in the literature to study short-term interest rates or univariate stochastic volatility.

The simulation method of moments facilitates the estimation of the parameters in the model specification. In particular, this estimation method helps avoid maximizing a complex likelihood function and instead relies on maximizing an objective function that captures the distance between sample moments and the simulated counterparts. The no-arbitrage hypothesis implies a set of pricing restrictions in the form of moment conditions. We test these moment conditions by applying a general specification test.

To account for possible persistence in the autoregressive (AR) component of fundamental risk factor specification, we expand our framework by replacing it with a double-autoregressive (DAR) process. For instance, in Djogbenou et al. (2023), this process provides greater flexibility to handle some non-stable patterns that might occur over time when modeling returns of crypto-assets. In addition to the autoregressive coefficient that allows for capturing the conditional mean dependence, the DAR processes usually incorporate a parameter representing the past dependence in the conditional variance of the considered variable. More details on these processes can be found in Borkovec and Kluppenberg (2001), Ling (2004), Ling (2007), and Ling and Li (2008). The considered extension combines the important features of the univariate WAR and the DAR processes.

We consider the largest five cryptocurrencies by market capitalization on November 16, 2022, represented by Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Binance Coin (BNB), and USD Coin (USDC). For each cryptocurrency, the top six trading pairs by daily volume are selected, and the test of the no-arbitrage hypothesis is applied first globally using the full sample, and then locally using a rolling window approach. Globally, the results show that the null hypothesis of no-arbitrage is rejected for Bitcoin and Ethereum, but is not rejected for Binance Coin and for the two stablecoins, USD Coin and Tether. Locally, we find that there are no arbitrage opportunities for USD Coin and a very limited number of such opportunities for Tether. In contrast, the results show that there are frequent arbitrage opportunities for the three native cryptocurrencies. Specifically, the percentage of days with no arbitrage opportunities is 100% for USD Coin, 94.5% for Tether, 60.5% for Bitcoin, 74% for Ethereum, and 78.5% for Binance Coin.

This paper is organized as follows. Section 2.2 describes the setup, the assumptions, and the no-arbitrage hypothesis. Section 2.3 describes the model estimation, the test for no-arbitrage, and the order condition for identification. Section 2.4 describes the cryptocurrency markets and the data, and presents the empirical results.

2.2 The Model

We use a model for a set of n cryptocurrencies that is similar to the framework described in Gourieroux et al. (2010), and choose the US dollar as the numeraire currency.

2.2.1 Setup and Assumptions

We assume that the information available to investors includes a latent fundamental risk factor, whose value at time t is denoted Σ_t, and some idiosyncratic risk factors, given by the $(n, 1)$ vector $\epsilon_t = (\epsilon_{1,t}, ..., \epsilon_{n,t})'$, where $\epsilon_{j,t}$ is the idiosyncratic risk of cryptocurrency j at time t, for $j = 1, ..., n$.

The latent factor Σ_t is assumed to follow a univariate Wishart autoregressive (WAR) process (Gourieroux et al. , 2009) with one degree of freedom, which is also a special case of autoregressive gamma process (Gourieroux and Jasiak , 2006). The process can be expressed as

$$\Sigma_t = x_t^2, \tag{2.2.1}$$

where x_t is a Gaussian autoregressive process

$$x_t = Mx_{t-1} + \eta_t, \tag{2.2.2}$$

with latent autoregressive coefficient M and latent innovation variance Ω, where $\eta_t \sim N(0, \Omega)$.

The vector ϵ_{t+1} of idiosyncratic risks is assumed to be multivariate Gaussian conditional on $\underline{\Sigma_{t+1}}$, $\underline{\epsilon_t}$ with mean zero and variance matrix $A'\Sigma_{t+1}A$, where $\underline{\Sigma_{t+1}} = (\Sigma_{t+1}, \Sigma_t, ...)$, $\underline{\epsilon_t} = (\epsilon_t, \epsilon_{t-1}, ...)$, and A is a $(1, n)$ vector.

For $j = 1, ..., n$, let $S_{j,t}$ denote the price of cryptocurrency j at time t, and let $y_{j,t+1} = ln\ S_{j,t+1} - ln\ S_{j,t}$ denote its log return for the period $(t, t+1)$. The dynamics of $y_{j,t+1}$ is specified as an affine function of the fundamental and the idiosyncratic risk factors:

$$y_{j,t+1} = f_j + F_j^0 \Sigma_t + \epsilon_{j,t+1}, \tag{2.2.3}$$

where f_j and F_j^0 are scalars, with $F_j^0 \geq 0$ since we expect the risk premium $F_j^0 \Sigma_t$ to be positive.

2.2.2 No-arbitrage Hypothesis

The hypothesis of absence of arbitrage opportunities implies the existence of a stochastic discount factor (sdf), which can be used to price all tradable assets of interest, including cryptocurrencies. We assume that the sdf for the period $(t, t+1)$ is an exponential affine function of the fundamental and the idiosyncratic risk factors:

$$M_{t,t+1} = exp\left(c + C^0 \Sigma_t + \gamma' \epsilon_{t+1}\right), \tag{2.2.4}$$

where c and C^0 are scalars and γ is an $(n, 1)$ vector.

Similar specifications were also considered by Gourieroux and Monfort (2007). The coefficient C^0 and the vector γ capture the sensitivities of the sdf to the factors. As in Gourieroux et al. (2010), we expect $C^0 \leq 0$, implying that the state prices decrease with the fundamental risk.

Under the no-arbitrage hypothesis, we have the following pricing restrictions:

$$E_t\left[M_{t,t+1} S_{j,t+1}\right] = S_{j,t}, \tag{2.2.5}$$

for $j = 1, ..., n$, which are equivalent to

$$E_t \left[M_{t,t+1} exp\ y_{j,t+1} \right] = 1, \tag{2.2.6}$$

for $j = 1, ..., n$, where E_t denotes the expectation conditional on the information set $\underline{\Sigma_t}$, $\underline{\epsilon_t}$. The following section discusses the estimation of the different parameters and how the no-arbitrage hypothesis can be tested.

2.3 Simulated Method of Moments

Originally proposed by McFadden (1989), the simulated method of moments (SMM) is a form of generalized method of moments where the parameter estimators are choosing such that simulated model moments match data moments. Under fairly general regularity conditions, Duffie and Singleton (1993) show the consistency and asymptotic normality of the SMM estimators. Further discussions of the SMM can be found in Gourieroux et al. (1993), Gallant and Tauchen (1996), Gourieroux and Monfort (2002), Davidson and MacKinnon (2004), and Ruge-Murcia (2012).

2.3.1 Model Estimation and Test for No-arbitrage

Let a vector $\theta : P \times 1$ collects the vectorized version of all unknown parameters $M, \Omega, f_1, \ldots,$ $f_n, F_1^0, \ldots, F_n^0, A, c, C^0$, and γ. For any admissible parameter vector θ in a compact set Θ, we can obtain the model-based simulated version of the vector of cryptocurrencies' returns $y_{s+1}^{\theta} = \left(y_{1,s+1}^{\theta}, \ldots, y_{n,s+1}^{\theta} \right)'$, using

$$y_{j,s+1}^{\theta} = f_j + F_j^0 \Sigma_s^{\theta} + \epsilon_{j,s+1}^{\theta}, s = 1, \ldots, S, \tag{2.3.1}$$

where $S = \tau T$ is the number of times the returns are generated and $\tau \geq 1$. The factor Σ_s^θ can be generated using Equation 2.2.1:

$$\Sigma_s^\theta = x_s^2, \tag{2.3.2}$$

and using Equation 2.2.2:

$$x_s^\theta = M x_{s-1}^\theta + \eta_s^\theta, \tag{2.3.3}$$

$s = 1, \ldots, S$. The innovation $\epsilon_{j,s+1}^\theta$ is a typical element of ϵ_{s+1}^θ drawn from a multivariate Gaussian with mean $0_{n\times 1}$ and variance $A'\Sigma_{s+1}^\theta A$ conditional on $\underline{\Sigma_{s+1}^\theta} = (\Sigma_{s+1}^\theta, \Sigma_s^\theta, \ldots)$ and $\underline{\epsilon_s^\theta} = (\epsilon_s^\theta, \epsilon_{s-1}^\theta, \ldots)$, while η_s^θ is drawn from a $N(0, \Omega)$ distribution.

To construct our objective function for estimating the unknown parameters, we note from the no-arbitrage condition in Equation (2.2.6) that

$$E_t \left[y_{j,t-q} M_{t,t+1} \exp \ y_{j,t+1} \right] = y_{j,t-q}, j = 1, \ldots, n, q = 1, \ldots, Q.$$

By the law of iterated expectation, we have:

$$E(y_{j,t+1-q}) = E \left[y_{j,t+1-q} M_{t,t+1} \exp \ y_{j,t+1} \right], j = 1, \ldots, n, q = 1, \ldots, Q. \tag{2.3.4}$$

We can write Equation 2.3.4 as

$$E\underbrace{\begin{pmatrix} y_{1,t} \\ y_{2,t} \\ \vdots \\ y_{n,t} \\ \vdots \\ y_{1,t+1-q} \\ y_{2,t+1-q} \\ \vdots \\ y_{j,t+1-q} \\ \vdots \\ y_{1,t+1-Q} \\ y_{2,t+1-Q} \\ \vdots \\ y_{j,t+1-Q} \end{pmatrix}}_{z_t} = E\underbrace{\begin{pmatrix} y_{1,t}M_{t,t+1}y_{1,t+1} \\ y_{2,t}M_{t,t+1}y_{2,t+1} \\ \vdots \\ y_{n,t}M_{t,t+1}y_{n,t+1} \\ \vdots \\ y_{1,t+1-q}M_{t,t+1}y_{1,t+1} \\ y_{2,t+1-q}M_{t,t+1}y_{2,t+1} \\ \vdots \\ y_{n,t+1-q}M_{t,t+1}y_{n,t+1} \\ \vdots \\ y_{1,t+1-Q}M_{t,t+1}y_{1,t+1} \\ y_{2,t+1-Q}M_{t,t+1}y_{2,t+1} \\ \vdots \\ y_{n,t+1-Q}M_{t,t+1}y_{n,t+1} \end{pmatrix}}_{Z_{t+1}} \tag{2.3.5}$$

So, Equation 2.3.5 is equivalent to $E(z_t) = E(Z_{t+1})$, where z_t and Z_{t+1} are nQ-vectors depending on information up to time t and $t+1$, respectively. Denote Z^{θ}_{s+1}, the simulated version of Z_{t+1}, which replaces $M_{t,t+1}, y_{j,t+1}, y_{j,t+1-q}$, with $M^{\theta}_{s,s+1}, y^{\theta}_{j,s+1}, y^{\theta}_{j,s+1-q}$, where

$$M^{\theta}_{s,s+1} = \exp(c + C^0\Sigma^{\theta}_s + \gamma'\epsilon^{\theta}_{s+1}). \tag{2.3.6}$$

The SMM matches the sample moment

$$g^*_T = \frac{1}{T}\sum_{t=1}^{T} z_t,$$

the sample analog of $E(z_t)$ based on observed data, to

$$g_S(\theta) = \frac{1}{S}\sum_{s=1}^{S} Z_{s+1}^{\theta},$$

the simulation-based sample analog of $E(Z_{t+1})$. The SMM estimator is given by

$$\hat{\theta} = \arg\min_{\theta\in\Theta} G(\theta)'\hat{V}^{-1}G(\theta), \tag{2.3.7}$$

where

$$G(\theta) = g_T^* - g_S(\theta) \tag{2.3.8}$$

with $\hat{V}$, any consistent estimator of the asymptotic variance of $\sqrt{T}G(\theta)$. An heteroskedasticity and autocorrelation consistent (HAC) estimator for V is given by

$$\hat{V} = \frac{1}{T}\sum_{t=1}^{T} u_t u_t' + \frac{1}{T}\sum_{h=1}^{B_T}\sum_{t=h+1}^{T} k\left(\frac{h}{B_T}\right)(u_t u_{t-h}' + u_{t-h}u_t'), \tag{2.3.9}$$

where $k(\cdot)$ is a kernel function, B_T is the bandwidth, and $u_t = z_t - \frac{1}{T}\sum_{t=1}^{T} z_t$, similarly to Newey and West (1987). Under mild conditions, it follows from Lee and Ingram (1991) that the simulated method of moment estimator $\hat{\theta}$ is asymptotically normal with the asymptotic distribution given by

$$\sqrt{T}(\hat{\theta}-\theta) \xrightarrow{d} N\left(0, \left(1+\frac{1}{\tau}\right)BV^{-1}B'\right), \tag{2.3.10}$$

where $B = E\left(\frac{\partial(Z_{t+1}^{\theta})'}{\partial\theta}\right)$ and $\left(1+\frac{1}{\tau}\right)BV^{-1}B'$ corresponds to the asymptotic variance formula based on the optimal weighting matrix $\hat{V}^{-1}$ minimizing the objective function in Equation 2.3.7. The variance in the asymptotic distribution in Equation 2.3.10 shows that the importance of the randomness due to the simulation reduces as τ increases.[1]

To test for no-arbitrage, we test the validity of the moment conditions in Equation

[1]In our application, we set τ to 20, and find this choice to be reasonable after sensitivity analyses, while taking into account the need to have a large number of simulated data.

2.3.5. Therefore, we can test if the identification restrictions generated by the no-arbitrage assumption are valid. In consequence, we can use the Hansen's general specification test defined by

$$W_T = T\left(1 + \frac{1}{\tau}\right) G(\hat{\theta})' \hat{V}^{-1} G(\hat{\theta}), \tag{2.3.11}$$

where

$$W_T \xrightarrow{d} \chi^2(nQ - P), \tag{2.3.12}$$

with P, the number of parameters, nQ, the number of conditions used, and $nQ > P$. See, e.g., Hansen (1982) and Ruge-Murcia (2012) for more details on the general description of the test and the simulated method of moments, respectively.

2.3.2 Order Condition

In this subsection, we discuss the order condition for the SMM estimators. This condition is derived based on the number nQ of moment conditions in Equation 2.3.4 that are used for estimation and the number P of model parameters, where Q denotes the number of lags.

We first note that Σ_s^θ depends on the latent autoregressive coefficient M and the latent innovation variance Ω. Therefore, we can simulate Σ_s^θ based on 2 parameters. Second, if we know Σ_s^θ and the parameters in the $(1, n)$ vector A, then we can simulate ϵ_s^θ. Third, using $2n$ parameters for f_j and F_j^0, $j = 1, ..., n$, we can also simulate the cryptoassets returns $y_{j,s}^\theta$. In addition, we need $n + 2$ parameters for c, C^0, and γ in the expression of the sdf $M_{s,s+1}^\theta$. Therefore, we have to estimate a total of $P = 4n + 4$ parameters.

Since we have P parameters and nQ moment conditions, the number n of cryptoassets needs to be such that $nQ \geq P$, in order to be able to identify the model. Moreover, this inequality needs to be strict, for the degrees of freedom of the asymptotic distribution in 2.3.12 to be strictly positive. This result is summarized in the proposition below.

Proposition 2.3.1 *For the estimation of the unknown parameters, the number n of cryp-*

toassets has to satisfy the order condition:

$$nQ > 4n + 4. \tag{2.3.13}$$

The proposition implies that we need the number of lags $Q \geq 5$, because $Q > \frac{4}{n} + 4$. Also, for $Q = 5$, we need the number of cryptoassets $n \geq 5$.

2.4 Empirical Application

This section describes the cryptocurrency markets and the data, and presents the empirical results.

2.4.1 The cryptocurrency markets and the data

Exchange platforms play an important role in facilitating day-to-day transactions for cryptocurrency investors. Through an exchange platform, investors can transfer funds to get in and out of the cryptomarket, trade one cryptocurrency for another, and store their funds securely in crypto wallets. As cryptocurrency exchange platforms exist only online, investors need to create an account on exchange platform's website to start using their services.

For an investor who wants to actively trade in the cryptomarket, there are mainly two types of exchange platforms they can choose from including centralized exchanges (CEXs) and decentralized exchanges (DEXs). The main difference between these two exchange types is the way the trade orders placed by market participants are carried out. When a trade order is executed under the supervision of a central authority, an exchange platform is considered as centralized. In decentralized exchanges, trade orders are executed with the help of automated algorithms on blockchain technology, which allows users to engage in peer-to-peer transactions without the need for a third party.

Despite the increasing number of decentralized exchanges, the cryptomarket is still vastly dominated by centralized exchanges as far as the daily trade volume is concerned. According to the recent research by Citi,[2] decentralized exchanges account for only 18.2% of the spot-trading volume with Uniswap itself, currently the largest decentralized exchange in the market, making up 70% of the total DEX volume.

For the analysis, we restrict our sample to the data from the five most reliable cryptocurrency exchanges in terms of the exchange score[3] according to coinmarketcap.com as of November 16, 2022. The set of exchange platforms includes Binance, Coinbase Exchange, Kraken, Kucoin, and Binance.US. and each platform in this set has an exchange score of at least 7 (out of 10). These exchange platforms cover a significant part of the overall trading volume and are used to proxy the cryptocurrency markets.[4]

The reason why we select exchange platforms based on their reliability instead of their size is because daily volumes reported by exchange platforms may not be always accurate. In fact, some of the big exchange markets are found to have engaged in wash trading meaning that they fake their trading volume to improve their ranking and popularity in the industry [Cong et al. (2022); Le Pennec et al. (2021)] and manipulate cryptocurrency prices as documented by Chen et al. (2022a).

Table 2.1: The top 5 cryptocurrency exchanges by reliability score as of November 16, 2022

Exchange Platform	# Cryptocurrencies	# Trading Pairs	Fiat Supported
Binance	386	1691	USD, EUR, GBP and +43 more
Coinbase Exchange	229	595	USD, EUR, GBP
Kraken	217	713	USD, EUR, GBP and +4 more
KuCoin	766	1389	USD, EUR, GBP and +45 more
Binance.US	145	309	USD

Source: https://coinmarketcap.com/rankings/exchanges/

All cryptocurrency exchanges in our sample are centralized so their business models are

[2] https://www.coindesk.com/business/2022/10/03/citi-says-decentralized-crypto-exchanges-are-winning-market-share-from-centralized-peers/

[3] Coinmarketcap.com states that the exchange score is calculated as the weighted average of web traffic factor, average liquidity, volume, and the confidence that the volume reported by an exchange is legitimate.

[4] The number of exchanges is large enough to satisfy the order condition in Proposition 2.3.1.

very similar. However, these exchange platforms can still be distinguished from one another when it comes to the number of cryptocurrencies that can be bought and sold, the list of trading pairs available to investors, and the set of fiat currencies supported by the platform. Table 2.1 summarizes the current state of the exchange platforms in our sample with respect to these characteristics.

2.4.1.1 Data Description

We consider the largest five cryptocurrencies by market capitalization including Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Binance Coin (BNB), and USD Coin (USDC). Then, we rank all trading pairs of these cryptocurrencies with respect to their daily volume as of November 16, 2022 and then select the top six pairs for each cryptocurrency. This adds up to a total of 30 different prices series that will be used in the analysis. The complete list of the trading pairs can be found in Table 2.2 along with the names of the exchange platforms each pairs were traded on, the respective daily volumes in US dollars, and their percentage contributions to the total daily volume.

The historical data on daily closing price of each trading pairs is obtained directly from the exchange platforms via their Application Programming Interfaces (APIs). It is worth noting that exchange platforms store and publish the historical data on trading pairs in such a way that the price of the base cryptocurrency is always expressed in terms of the units of the quote cryptocurrency. For consistency, we convert the price of all trading pairs into US dollars by using the corresponding exchange rate between the quote cryptocurrency and the US dollar retrieved from Yahoo Finance.

Table 2.3 provides the summary statistics of daily log returns of cryptocurrencies with respect to trading pairs.

Table 2.2: List of cryptocurrency trading pairs with the highest daily volumes as of November 16, 2022

Cryptocurrency	Trading Pair	Exchange Platform	Daily Volume	Daily Volume (%)
Binance Coin (BNB)	BNB/USDT	Binance	$117,548,065	12.20%
	BNB/BUSD	Binance	$95,634,654	9.92%
	BNB/BTC	Binance	$15,580,100	1.62%
	BNB/ETH	Binance	$5,083,991	0.53%
	BNB/USDT	KuCoin	$3,500,603	0.36%
	BNB/EUR	Binance	$3,182,434	0.33%
Bitcoin (BTC)	BTC/USDT	Binance	$4,738,786,091	13.16%
	BTC/BUSD	Binance	$2,722,430,698	7.56%
	BTC/USD	Coinbase Exchange	$807,460,570	2.24%
	BTC/USDT	KuCoin	$203,840,895	0.57%
	BTC/USD	Binance.US	$166,547,193	0.46%
	ETH/BTC	Binance	$136,816,282	0.38%
Ethereum (ETH)	ETH/USDT	Binance	$1,044,487,266	7.17%
	ETH/USD	Coinbase Exchange	$755,769,692	5.19%
	ETH/BUSD	Binance	$645,680,778	4.43%
	ETH/USDT	KuCoin	$141,319,928	0.97%
	ETH/BTC	Binance	$137,002,090	0.94%
	ETH/USD	Kraken	$91,258,418	0.63%
Tether (USDT)	BTC/USDT	Binance	$4,738,786,091	9.38%
	BUSD/USDT	Binance	$1,151,360,536	2.28%
	ETH/USDT	Binance	$1,044,238,487	2.07%
	USDT/USD	Kraken	$299,184,079	0.59%
	BTC/USDT	KuCoin	$203,840,895	0.40%
	XRP/USDT	Binance	$198,081,627	0.39%
USD Coin (USDC)	USDC/USD	Kraken	$58,226,926	1.30%
	USDC/USDT	Kraken	$33,092,103	0.74%
	USDC/EUR	Kraken	$24,085,577	0.54%
	USDC/USDT	KuCoin	$9,354,817	0.21%
	BTC/USDC	KuCoin	$9,246,101	0.21%
	ETH/USDC	KuCoin	$8,027,523	0.18%

Table 2.3: Summary statistics of daily log returns until November 16, 2022

Trading Pair	Beginning Date	Number of Days	Min	Max	Mean	Std
BNB/BTC	2017-07-14	1958	-0.549	0.657	0.0040974	0.0688291
BNB/BUSD	2019-09-20	1160	-0.527	0.534	0.0021692	0.0558232
BNB/ETH	2017-11-09	1840	-0.545	0.536	0.0026349	0.0595479
BNB/EUR	2020-01-03	1055	-0.588	0.531	0.0028154	0.0571841
BNB/USDT	2017-11-06	1843	-0.529	0.533	0.0026505	0.0603182
BNB/USDT	2019-06-19	1253	-0.531	0.536	0.0016067	0.0551446
BTC/USDT	2017-11-09	1840	-0.449	0.224	0.0004430	0.0412449
BTC/BUSD	2019-09-20	1160	-0.449	0.178	0.0003966	0.0390128
BTC/USD	2017-02-23	2099	-0.491	0.241	0.0012454	0.0421090
BTC/USDT	2018-07-08	1599	-0.448	0.178	0.0005486	0.0386145
BTC/USD	2019-09-24	1156	-0.484	0.178	0.0005489	0.0395931
ETH/BTC	2017-11-09	1840	-0.460	0.227	0.0004411	0.0405444
ETH/USDT	2017-11-09	1840	-0.537	0.255	0.0006837	0.0527709
ETH/USD	2017-02-23	2099	-0.568	0.282	0.0021212	0.0562751
ETH/BUSD	2019-10-21	1129	-0.540	0.235	0.0016527	0.0518516
ETH/USDT	2018-07-08	1599	-0.537	0.234	0.0005253	0.0511201
ETH/BTC	2017-07-14	1958	-0.555	0.247	0.0008752	0.0533127
ETH/USD	2015-10-11	2600	-0.583	0.372	0.0028438	0.0595864
BTC/USDT	2017-08-17	1924	-0.069	0.063	-0.0000061	0.0061191
BUSD/USDT	2019-09-20	1160	-0.050	0.053	-0.0000017	0.0033486
ETH/USDT	2017-11-09	1840	-0.063	0.057	-0.0000065	0.0067026
USDT/USD	2017-03-29	2065	-0.058	0.045	-0.0000001	0.0043200
BTC/USDT	2018-07-08	1599	-0.038	0.037	-0.0000093	0.0042588
XRP/USDT	2018-05-04	1664	-0.029	0.037	-0.0000012	0.0051878
USDC/USD	2020-01-08	1050	-0.010	0.010	-0.0000096	0.0007398
USDC/USDT	2020-01-08	1050	-0.057	0.048	-0.0000216	0.0034762
USDC/EUR	2020-01-08	1050	-0.031	0.042	-0.0000242	0.0061821
USDC/USDT	2020-11-23	730	-0.012	0.010	0.0000028	0.0010906
BTC/USDC	2019-01-07	1416	-0.173	0.149	-0.0000105	0.0080633
ETH/USDC	2019-01-07	1416	-0.167	0.123	-0.0000451	0.0096427

2.4.2 Empirical Results

This section presents the estimation results using the observed daily returns of the cryptocurrencies. Given that the digital currencies can have different returns on different platforms, we treat the cryptocurrencies on different platforms as different assets. Using the data, we answer the following empirical questions. To which extent are the cryptocurrencies priced and affected by the fundamental risk factor? Are individual cryptocurrency markets free of arbitrage opportunity?

To answer these questions, for each cryptocurrency (Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Binance Coin (BNB), and USD Coin (USDC)), we apply the SMM using the corresponsing six trading pair samples ($n = 6$). We choose the number of lags $Q = 5$, so that the values of n and Q satisfy the order condition in Section 3.2.

To implement the SMM, we proceed as follows.

1. Set initial values for the parameters: $M, \Omega, f_1, \ldots,\ f_n, F_1^0, \ldots, F_n^0, A, c, C^0$, and γ. We consider 1000 sets of initial values and select the set that leads to the lowest value of the objective function in Equation 2.3.7.

2. Simulate the fundamental risk factor $\Sigma_{s+1}^{\theta}, s = 1, \ldots, S$, using Equation 2.3.2 and Equation 2.3.3.

3. For $s = 1, \ldots, S$, generate the $(n, 1)$ vector ϵ_{s+1}^{θ} of idiosyncratic risks of cryptoassets from a multivariate Gaussian distribution with conditional mean 0_n and conditional variance $A'\Sigma_{s+1}^{\theta}A$.

4. For $s = 1, \ldots, S$, compute $y_{j,s+1}^{\theta}, j = 1, \ldots, n$ and $M_{s,s+1}^{\theta}$ based on Equations 2.3.1 and 2.3.6, respectively, and using Σ_{s+1}^{θ} and ϵ_{s+1}^{θ}.

5. Find $\hat{\theta}$ from the optimization problem in Equation 2.3.7, based on the generated $y_{j,s+1}^{\theta}, j = 1, \ldots, n$ and $M_{s,s+1}^{\theta}$. Obtain the asymptotic normal interval for the components of $\hat{\theta}$.

6. Perform the no-arbitrage test using Equations 2.3.11 and 2.3.12.

The variance $\hat{V}$ employed in Steps 5 and 6 is computed using a Bartlett kernel with the bandwidth, which we set to $4(T/100)^{2/9}$, as suggested by Newey and West (1994). The kernel $k(\cdot)$ is a decreasing function, which accounts for the the decaying dependence between the observations at t and $t+h$ when h increases.

2.4.2.1 Global Analysis of Individual Cryptocurrency Markets

We first conduct a global analysis using the sample data over the last two years, which consists of 730 days, in order to minimize the impact of changes in the market's structure over multiple years. The model parameters are estimated for all cryptocurrencies and the 95% confidence intervals around the point estimates are constructed based on their asymptotic distribution. The results are presented in Table 2.4. The table shows that the absolute value of the latent autoregressive coefficient M is estimated to be 1 or very close to 1 for Binance Coin, Bitcoin, Ethereum, and USD Coin, implying that the corresponding processes for fundamental risk are nonstationary and, therefore, casting some doubt on the validity of the estimation results.

In response to this finding, we extend our analysis by considering a modified version of the standard model where we assume that x_t in Equation 2.2.1 follows the process:

$$x_t = Mx_{t-1} + \eta_t,$$

with $\eta_t \sim N(0, w + \alpha x_{t-1}^2)$, where w and α are two new parameters. This specification is similar to the one used in Djogbenou et al. (2023) and we refer to it as the DAR version of the model. In this case, the total number of parameters is $4n + 5$ and the order condition in Equation (2.3.13) becomes

$$nQ > 4n + 5.$$

Table 2.4: Standard AR: Global Parameter Estimates

Model Parameters	Binance Coin (BNB)		Bitcoin (BTC)		Ethereum (ETH)		USD Coin (USDC)		Tether (USDT)	
	Point Estimate	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***
M	-0.998	2.133	0.996	1.603	0.999	3.204	1.000***	0.249	0.098***	5.167E-20
Ω	0.052	0.148	0.327	0.276	0.012	0.295	0.034**	0.014	0.046***	4.298E-18
f_1	-3.931***	0.021	0.166	2.088	0.074	4.048	-1.484***	0.012	0.272***	5.566E-17
f_2	-0.213	0.404	-2.764***	0.032	0.189	20.142	0.344	0.494	1.639***	2.992E-17
f_3	-1.509***	0.029	0.162	1.076	0.091	45.232	0.208	0.175	0.247***	7.875E-18
f_4	1.564**	0.701	0.226	2.191	0.143	27.864	0.346	0.389	-0.417***	3.474E-19
f_5	-0.070	0.963	0.322	1.735	0.037	58.394	-0.249	0.219	-0.220***	2.027E-17
f_6	-0.629**	0.254	0.054	0.070	0.003	52.744	-0.218	0.271	-0.102***	2.574E-18
F_1^0	4.644***	0.001	0.762***	0.272	0.659***	0.061	1.605***	0.001	0.004***	6.027E-21
F_2^0	0.885***	0.011	0.789***	0.001	0.006	0.299	0.128***	0.049	0.002***	3.292E-21
F_3^0	0.036***	0.001	0.040	0.144	0.582	0.680	0.092***	0.017	0.007***	9.071E-22
F_4^0	1.123E-08	0.015	0.807***	0.280	0.329	0.417	0.525***	0.041	0.004***	4.551E-23
F_5^0	1.255E-08	0.019	0.230	0.219	0.790	0.883	0.486***	0.024	0.162***	2.206E-21
F_6^0	0.849***	0.006	0.377***	0.010	0.905	0.800	0.440***	0.029	0.185***	2.988E-22
A_1	1.063***	0.047	-0.122	0.185	-0.140***	0.022	0.197***	0.006	7.738***	2.631E-20
A_2	1.010***	0.005	-0.240***	0.007	-0.694***	0.099	-0.522***	0.004	-3.513***	1.616E-21
A_3	-0.307***	0.035	-0.202**	0.091	-0.320*	0.194	-0.058***	0.003	3.889***	5.774E-21
A_4	0.090	0.192	-0.133	0.177	-0.182	0.119	0.107***	0.007	1.540***	1.182E-20
A_5	-2.407***	0.013	-1.072***	0.144	-0.046	0.280	0.183***	0.006	-1.343***	6.985E-21
A_6	-0.013	0.009	0.011	0.034	-0.002	0.259	0.476***	0.010	-1.237***	3.068E-20
c	-2.444***	0.562	-0.376	0.386	-0.002	3.923	-0.312***	0.096	-39.562***	2.708E-17
C^0	-17.371***	0.012	-7.689***	0.028	-3.173***	0.056	-2.408***	0.010	-3.039***	2.923E-21
γ_1	1.529***	0.033	-1.034***	0.005	-0.487***	0.004	-0.780***	0.002	-0.284***	9.117E-20
γ_2	0.060*	0.031	0.174***	0.009	0.073***	0.020	-0.057***	0.005	-0.024***	4.139E-20
γ_3	1.097***	0.010	0.139***	0.007	-0.652***	0.009	-0.195***	0.001	0.406***	4.583E-20
γ_4	5.551***	0.003	-1.703***	0.005	-1.105***	0.005	-1.457***	0.001	1.002***	1.815E-20
γ_5	-0.008	0.075	-0.225***	0.040	-0.584***	0.001	-0.595***	0.002	0.272***	1.582E-20
γ_6	-0.235***	4.118E-04	-0.891***	4.185E-04	-2.421***	4.646E-05	1.035***	0.004	2.607***	1.457E-20

This condition holds when $Q = 5$ and $n = 6$.

The SMM is implemented in a very similar manner with minor changes. In particular, η_s^θ is simulated from a normal distribution, with the difference that Ω is replaced by $w + \alpha(x_{s-1}^\theta)^2$. The point estimates and the 95% confidence intervals for the coefficients of the DAR model are reported in Table 2.5. As seen in the table, for each cryptocurrency, the estimate of the latent autoregressive coefficient M implies a more realistic fundamental risk process.

For both Table 2.4 and Table 2.5, we also note that most estimates of the coefficient F^0 in the cryptoasset return equation are positive. As a result, the corresponding returns (and risk premia) depend positively on the fundamental risk Σ_t. Also, all estimates of the coefficient C^0 in the stochastic discount factor are negative, showing that the state prices decrease with the fundamental risk.

For the global analysis of each cryptocurrency, the result of the no-arbitrage test along with the estimated test statistics $\hat{W}_T$ and the critical value for the 95% confidence level based on Equation 2.3.11 and Equation 2.3.12, respectively, are summarized in Table 2.6 (for the standard AR model) and Table 2.7 (for the DAR model). We consider the results for the DAR model in Table 2.7 as our official test results and present the results in Table 2.6 for an interesting comparison. Even if the test results in Table 2.6 are technically less reliable (for the reason discussed above), these results are the same as those in Table 2.7 for four of the five cryptocurrencies.

The tables show that, for Bitcoin and Ethereum, the null hypothesis that there are no-arbitrage opportunities across trading pairs is rejected in both the standard AR model and the DAR model. These findings suggest that, for these cryptocurrencies, investors could profit by buying and selling different trading pairs of the same cryptocurrency. At the same time, for USD Coin and Tether, the tables indicate that the null hypothesis is not rejected in both models. This is not surprising since these cryptocurrencies are designed to maintain a stable price against the US dollar. Hence, they are often referred to as stablecoins. For Binance Coin, the null hypothesis is rejected in the standard AR model but is not rejected in the DAR model.

Table 2.6: Standard AR: Results of the no-arbitrage test

Cryptocurrency	**Ticker**	$\hat{W}_T$	**Critical Value***	**Test Result**
Binance Coin	BNB	267.268	5.991	*Reject the null*
Bitcoin	BTC	32098.664	5.991	*Reject the null*
Ethereum	ETH	392.598	5.991	*Reject the null*
USD Coin	USDC	4.070	5.991	*Cannot reject the null*
Tether	USDT	2.742	5.991	*Cannot reject the null*

Note: *The value corresponds to the critical value for 95% confidence level.

Table 2.7: DAR: Results of the no-arbitrage test

Cryptocurrency	**Ticker**	$\hat{W}_T$	**Critical Value***	**Test Result**
Binance Coin	BNB	1.165	3.841	*Cannot reject the null*
Bitcoin	BTC	7.573	3.841	*Reject the null*
Ethereum	ETH	4.565	3.841	*Reject the null*
USD Coin	USDC	1.473	3.841	*Cannot reject the null*
Tether	USDT	2.520	3.841	*Cannot reject the null*

Note: *The value corresponds to the critical value for 95% confidence level.

2.4.2.2 Local Analysis of Individual Cryptocurrency Markets

In this section, we perform the analysis locally based on a rolling window approach with the window size equal to 100 days. More specifically, the SMM and the no-arbitrage test are applied locally to each cryptocurrency using only the observations from the last 100 days, with estimation sample ending at $T-200, T-199, \ldots, T-1, T$.

To summarize the results, for each cryptocurrency we calculate the percentage of days

Table 2.5: DAR: Global Parameter Estimates

Model Parameters	Binance Coin (BNB)		Bitcoin (BTC)		Ethereum (ETH)		USD Coin (USDC)		Tether (USDT)	
	Point Estimate	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***	***Point Estimate***	***Standard Deviation***
M	0.142***	0.002	0.133***	0.005	0.021***	1.172E-08	0.161***	5.250E-05	0.204***	3.849E-05
w	8.345E-06	305.244	2.678***	0.001	0.001***	9.168E-05	0.660***	9.041E-05	0.010	1.031E-02
α	0.053***	0.004	0.128***	0.004	0.129***	3.106E-08	0.051***	4.280E-06	0.031***	2.124E-04
f_1	-0.045	7.496	-0.112***	0.020	10.061***	2.726E-08	-4.594***	1.994E-06	0.588***	1.716E-02
f_2	-0.045	13.461	-0.321***	0.003	4.912***	1.007E-10	-4.547***	4.423E-07	0.117***	5.589E-04
f_3	-0.045	6.238	0.107***	0.012	-2.491***	1.625E-13	-15.032***	6.587E-12	0.793***	3.658E-03
f_4	-0.047	1.362	-0.032	0.020	6.370***	3.979E-09	-1.243***	1.918E-07	0.317***	3.067E-04
f_5	-0.045	28.956	-0.394***	0.009	-4.349***	5.139E-14	-1.165***	6.688E-07	0.551***	1.568E-02
f_6	-0.045	28.750	-0.844***	0.002	8.383***	9.848E-09	0.337***	3.228E-05	-0.104***	5.110E-04
F_1^0	0.597***	7.613E-05	0.088*	0.046	3.169***	1.335E-11	7.554E-08	3.313E-07	6.481E-05	2.010E-04
F_2^0	2.26E-04*	1.366E-04	0.113***	0.008	1.128***	6.966E-14	0.481***	8.575E-08	0.523***	6.662E-06
F_3^0	0.422***	6.354E-05	2.480E-08	0.031	0.251***	3.312E-17	0.107***	1.006E-12	0.325***	4.340E-05
F_4^0	1.087***	1.387E-05	0.083*	0.044	0.024***	1.943E-12	0.173***	7.601E-08	1.258E-10	3.650E-06
F_5^0	0.356***	2.937E-04	0.122***	0.025	1.055***	1.532E-17	2.785***	3.974E-07	0.208***	1.823E-04
F_6^0	0.094***	2.917E-04	1.575***	0.001	0.443***	7.781E-12	0.262***	1.967E-05	0.182***	6.115E-06
A_1	-1.347***	0.001	-0.109***	0.004	9.150***	1.351E-08	0.903***	5.700E-05	0.988***	1.196E-04
A_2	-0.304***	0.002	0.192***	0.001	-9.059***	3.355E-09	0.287***	9.275E-06	-1.854***	2.533E-06
A_3	0.311***	0.001	0.069***	0.007	10.792***	2.295E-08	-5.974***	1.994E-04	-0.170***	1.740E-04
A_4	-1.158***	1.586E-04	-0.022***	0.006	13.517***	3.315E-08	-5.416***	2.231E-04	-0.380***	1.711E-04
A_5	0.769***	0.003	-0.161***	0.002	5.152***	3.421E-09	-0.807***	5.963E-05	3.298***	1.936E-05
A_6	-0.829***	0.003	-0.001	0.009	-26.060***	1.673E-08	0.227***	1.471E-05	-1.725***	2.740E-04
c	0.276***	0.005	-6.240***	0.003	-33.330***	2.784E-08	-16.146***	1.111E-05	-6.743***	0.001
C^0	-6.209***	1.320E-07	-5.910***	2.622E-04	-3.436***	1.382E-11	-3.510***	1.273E-06	-0.853***	1.469E-05
γ_1	-0.465***	2.373E-05	-0.688***	2.995E-04	-7.426***	1.908E-08	1.514***	3.238E-05	0.205***	1.520E-04
γ_2	0.027***	5.353E-06	-0.401***	0.001	1.609***	1.889E-08	0.244***	1.031E-05	-0.004***	2.853E-04
γ_3	-0.795***	5.490E-06	-2.446***	1.888E-04	11.006***	2.250E-08	5.557***	2.143E-04	-1.158***	2.620E-05
γ_4	-1.341***	2.041E-05	1.656***	5.918E-05	15.875***	2.818E-08	-6.218***	1.943E-04	-1.112***	5.844E-05
γ_5	-0.666***	1.355E-05	0.177***	4.425E-04	1.641***	1.074E-08	-1.635***	2.894E-05	-0.353***	0.001
γ_6	-1.353***	1.461E-05	-3.239***	3.641E-06	8.000***	5.433E-08	-0.710***	8.150E-06	1.779***	2.654E-04

with no arbitrage opportunities, which is the total number of days for which we do not reject the null hypothesis of no-arbitrage as a percentage of the total number of estimation days, and report it in the last column of Table 2.8 (for the standard AR model) and Table 2.9 (for the DAR model). We consider the results for the DAR model in Table 2.9 as our official test results and present the results in Table 2.8 for an interesting comparison. Even if the test results in Table 2.8 are technically less reliable (for the reason discussed in the global analysis), these results are similar to those in Table 2.9 for all five cryptocurrencies.

The tables show that, for Binance Coin, Bitcoin, and Ethereum, the percentage of days with no arbitrage opportunities varies between 59% and 85%. Therefore, although these cryptocurrencies do not seem to allow arbitrage opportunities on the majority of the days, they do appear to present investors with such opportunities frequently. On the other hand, there are no days with arbitrage opportunities for USD Coin and a very limited number of such days for Tether. As noted above in the global analysis, the absence or limited number of arbitrage opportunities in USD Coin and Tether could be explained by the fact that these cryptocurrencies are stablecoins. These results for USD Coin and Tether could also imply that these cryptocurrencies are locally efficient, which is in line with the findings of Djogbenou et al. (2023).

Table 2.8: Standard AR: Results of no-arbitrage test conducted locally

Cryptocurrency	**Ticker**	**Total Est. Days**	**% of Days with No-arbitrage***
Binance Coin	BNB	200	85%
Bitcoin	BTC	200	59%
Ethereum	ETH	200	60%
USD Coin	USDC	200	100%
Tether	USDT	200	96%

*Note:** Percentage of days with no-arbitrage= 100×(total number of days with no arbitrage opportunity)/(total number of estimation days).

Table 2.9: DAR: Results of no-arbitrage test conducted locally

Cryptocurrency	**Ticker**	**Total Est. Days**	**% of Days with No-arbitrage***
Binance Coin	BNB	200	78.50%
Bitcoin	BTC	200	60.50%
Ethereum	ETH	200	74.00%
USD Coin	USDC	200	100.00%
Tether	USDT	200	94.50%

Note: * Percentage of days with no-arbitrage= 100×(total number of days with no arbitrage opportunity)/(total number of estimation days).

2.5 Concluding Remarks

This paper contributes to the literature on the efficiency of the cryptocurrency markets by considering a general model for cryptocurrencies traded on multiple exchange platforms. The model assumes that the return of each cryptoasset varies in terms of a common fundamental risk factor and a specific idiosyncratic risk factor. The fundamental risk factor is modeled as a univariate Wishart autoregressive process (autoregressive gamma process), while the vector of the idiosyncratic risk factors is assumed to follow a conditional multivariate Gaussian distribution.

We provide an algorithm on how the simulation method of moments could be implemented to facilitate the estimation of the parameters in the model and conduct inferences. To assess the efficiency of cryptocurrency markets, we test the moment conditions implied by the no-arbitrage hypothesis and describe the procedure for its implementation.

The empirical results based on a global analysis show that fundamental risk plays an important role in explaining the dynamics of cryptocurrency returns and suggest the cryptocurrency markets are generally inefficient. A local analysis, where parameters are estimated sequentially using a rolling window, allows to mitigate those findings.

While the native cryptocurrency markets offer more frequent arbitrage opportunities, this inefficiency is drastically reduced in stablecoin markets. The analysis in this paper could be extended by modeling fundamental risk as a matrix that follows a Wishart autoregressive process and by considering alternative risk factors.

Chapter 3

Predictability of Funding Rates

3.1 Introduction

The crypto derivatives market accounts for over 80% of the total trading volume as of July 2025. The market offers a range of instruments including options, traditional futures, and perpetual futures to speculate on and hedge against cryptocurrency price movements. Among them, perpetual futures stand out due to their uniqueness in the crypto ecosystem. Unlike traditional futures, perpetual contracts have no expiry date and do not require the delivery of the underlying cryptocurrency. In the absence of such features, these contracts rely on what is often referred to as a funding fee mechanism to keep the futures and spot markets tethered.

In the funding fee mechanism, traders with either long or short positions are compensated, depending on the state of the market, and the compensation is paid by those on the opposite side of the trade. The amount to be paid is called "funding fee" and is calculated at regular intervals as the funding rate at the time multiplied by the notional value of a trader's position, denominated in units of the underlying cryptocurrency. The funding rate itself is a function of the spread between the spot and futures price of a cryptocurrency, and its sign determines the direction of payments between traders. When the funding rate is positive, long positions are required to make payments to short positions, or vice versa. A comprehensive analysis of the cryptocurrency perpetual futures can be found in De Blasis and Webb (2022).

Although crypto exchanges usually set a target value for each contract, funding rates are observed to fluctuate significantly when market imbalances are large. This makes it compelling for traders to accurately predict funding rates, especially in short horizons. Traders can use funding rate forecasts to calculate the fee that they are expected to pay or receive in the periods to come. This would help them decide whether to keep or exit their positions before the payment becomes due, thereby avoiding substantial costs when the funding rate is predicted to be unusually high. Predictions could also serve as a signal for the prevailing market sentiment about the underlying cryptocurrency in the future. This is related to the fact that a funding rate above (below) its target rate usually reflects bullish (bearish) sentiment for the underlying. Moreover, short-term predictability of funding rates could be leveraged by arbitrageurs to develop and improve their strategies. Funding fee arbitrage strategies are covered in Section 3.

Given the importance for cryptocurrency trading, this study analyzes out-of-sample (OOS) predictability of the next period's funding rate with a focus on BTC/USDT perpetual contracts (the highest-volume Bitcoin pair traded on Binance and Bybit). OOS point forecasts are generated from a set of univariate double autoregressive (DAR) models based on Ling (2007), whose parameters are estimated recursively via expanding and rolling window approaches.

A time-varying parameter version of the DAR process, namely tvDAR(1), was proposed recently by Djogbenou et al. (2023) to model daily stablecoin price series, which, from an empirical modeling perspective, share striking similarities with the funding rates analyzed in this paper. Despite having a long-term target value, neither stablecoin prices nor funding rates are immune to market shocks and fluctuate significantly during periods of market distress. This leads to short-term price fluctuations, inducing conditional heteroskedasticity in their series. Deviations from the target value are usually observed to be temporary for stablecoin prices and funding rates, resulting in mean-reverting behavior. These common characteristics motivate the choice of DAR processes as the basis for the modelling strategy.

The OOS accuracy of each forecast model is measured in terms of Mean Squared Prediction Errors (MSPE) and Mean Directional Accuracy (MDA), and evaluated statistically via the Clark and West (2007) and Pesaran and Timmermann (1992) tests, respectively. Predictability is also examined locally for each funding rate series by estimating the top Lyapunov exponents of DAR models at each point in time over the evaluation period. The link between the stability of a dynamical system and its predictability is established in Abarbanel et al. (1991) in the context of local Lyapunov exponents and applied to observational data in Abarbanel et al. (1992); the smaller the exponent, the more stable the dynamical system, and hence, the greater the predictability of the process. Lyapunov exponent was used in Bask (1996), Dechert and Gencay (1992), Fernández-Rodríguez et al. (2005), Gao et al. (2011), and Resende and Zeidan (2008) to study the dynamics of currency exchange rates.

The study reveals that both Binance and Bybit funding rates were predictable from November 10, 2024 through March 17, 2025, and these results are robust to the estimation window. There is, however, a trade-off between rolling and expanding window estimations, with the former yielding lower MSPE and the latter achieving higher MDA for both funding rate series considered. According to the local predictability analysis, the degree of predictability was not uniform over the entire evaluation period. The funding rates are observed to be structurally less stable and locally less predictable on Bybit during periods of sustained high volatility. On the contrary, the funding rates appear more resilient to such volatile episodes on Binance, showing less heterogeneity over time.

The rest of the paper is organized as follows: Section 3.2 introduces the method used by Binance and Bybit to calculate funding rates; Section 3.3 discusses the related arbitrage strategies; Section 3.4 presents the historical data obtained from these exchanges; Section 3.5 describes the forecast method; Section 3.6 reports the results; and Section 3.7 concludes with a brief discussion.

3.2 Funding Rate Calculation

Conventionally, the funding rate for a given cryptocurrency pair a is determined by exchange b at the end of an n-equally-spaced time interval using the formula below:

$$fr_{a,b,t} = \begin{cases} API_{a,b,t} - d_{a,b} & \text{if } API_{a,b,t} \geq r_{a,b} + d_{a,b} \\ \\ r_{a,b} & \text{if } r_{a,b} - d_{a,b} < API_{a,b,t} < r_{a,b} + d_{a,b} \\ \\ API_{a,b,t} + d_{a,b} & \text{otherwise} \end{cases} \tag{3.2.1}$$

where $API_{a,b,t} = \frac{1}{n}\sum_{m=t-n+1}^{t} PI_{a,b,m}$ is the average premium index at time t with $PI_{a,b,t}$ denoting the difference between the future and spot price of the underlying cryptocurrency pair a. The parameter $d_{a,b}$ is a dampener expressed in percentage points, and $r_{a,b}$ is the default interest rate, both of which can vary across cryptocurrency pairs and exchange platforms, as indicated by the indices.

While the interest rate represents the cost of long positions in neutral markets, the dampener is used by crypto exchanges in order to reduce the effects of short-term volatility, price manipulation, and lack of liquidity on funding fees. The dampener pulls the funding rate back toward its target value when the average premium index is too small, and pushes it away from the extreme values when the index is too large. Technically, using such a dampener limits the responsiveness of funding rates to extreme market imbalances, thereby suppressing the true underlying autoregressive process. Observationally, this creates an almost flat trajectory in historical funding rates when the imbalances are very small, often followed by sudden and explosive spikes.

In this study, the sample is characterized with the indices $a = \{BTC/USDT\}$ and $b \in \{Binance, Bybit\}$ such that $fr_{BTC/USDT,Binance,t}$ and $fr_{BTC/USDT,Bybit,t}$ represent the funding rates for BTC/USDT perpetual futures on Binance and Bybit at time t, respectively. For these contracts, $d_{BTC/USDT,b} = 0.05\%$ and $r_{BTC/USDT,b} = 0.01\%$ for $b \in \{Binance, Bybit\}$, meaning that both Binance and Bybit use the same dampener and default interest rate. These exchanges, however, differ from one another in terms of the frequency at which they calculate the premium index. On Binance, the premium index is calculated every 5 seconds over an 8-hour interval, amounting to n=5,760 data points, whereas Bybit calculates it every minute over the same interval with n=480 data points.

3.3 Funding Fee Arbitrage

In the current state of the market, it is possible for traders to arbitrage from a perpetual contract by taking a market-neutral position in the underlying cryptocurrency. The arbitrage strategy they should apply depends on the funding rate and its sign at the time, and can be either within-exchange or cross-exchange.

Within-exchange arbitrage opportunity follows from the funding fee mechanism, and it is encouraged by crypto exchanges.[1] To profit from this type of arbitrage, traders are supposed to go long in the spot and short in the futures market of the same cryptocurrency pair when its funding rate is positive, or go short in the spot and long in the futures market otherwise. This way, they can collect the corresponding funding fee without exposing themselves to any directional risk.

Cross-exchange arbitrage opportunities, on the other hand, are a result of the fragmented nature of the cryptocurrency derivatives market and present themselves when the funding rates of the exact same perpetual contract are meaningfully different between two exchange platforms. When the funding rate $fr_{a,X,t}$ on exchange X is less than the funding rate $fr_{a,Y,t}$ on exchange Y, traders can pocket the difference by going long on the perpetual contract on exchange X and short on exchange Y. Of course, to eliminate directional risk, the notional

[1] see https://www.binance.com/en/futures/funding-history/perpetual/arbitrage-data for such an arbitrage advertised on Binance.

amount of the contract must be the same on both exchanges. Unlike the within-exchange arbitrage, this type of arbitrage requires transacting only in the futures market.

Whether it is within or across exchanges, the effectiveness and profit of an arbitrage strategy depend heavily on a trader's ability to make predictions for the next period's funding rate. Equipped with accurate forecasts, arbitrageurs could assess the effectiveness of one strategy over another and turn them into actionable insights.

3.4 Sample

Historical data is obtained from the exchange platforms through their API, spanning the period from January 11, 2024,[2] through March 17, 2025. With funding rates being calculated and published every 8 hours for the BTC/USDT perpetuals, this leads to a total of 1296 observations for each exchange. Figure 3.1 below shows the evolution of the funding rates over the sampling period in panels a and b, along with the spot price of the underlying cryptocurrency pair in panel c.

Visually, both funding rate series look quite similar, with Bybit exhibiting slightly larger spikes. Deviations from the default rate of 0.01% are explosive due to anchoring and generally short-lived on both exchange platforms. Such similarities are not surprising and are a consequence of both series following the same underlying cryptocurrency. Also, there is a notable (cor)relation between the BTC/USDT spot prices and the funding rate series in some episodes (Appendix D, Figure D.2).

When BTC displays a strong local trend, deviations look more pronounced and often persist in one direction. For instance, during the bull run leading up to the BTC's halving event in April 2024, funding rates had been significantly higher than usual and at some point, reached as high as 0.088% on Binance and 0.113% on Bybit, rendering long positions

[2]With the introduction of electronically traded funds (ETFs) on January 10, 2024, there is a possibility that the dynamics of BTC might have significantly changed. In fact, Babalos et al. (2025) find, through a wavelet coherence analysis, strong evidence of volatility spillovers from the Grayscale Bitcoin Trust ETF to BTC futures. Therefore, observations before this date are omitted from the analysis.

Figure 3.1: Historical funding rates of the BTC/USDT perpetual futures contracts

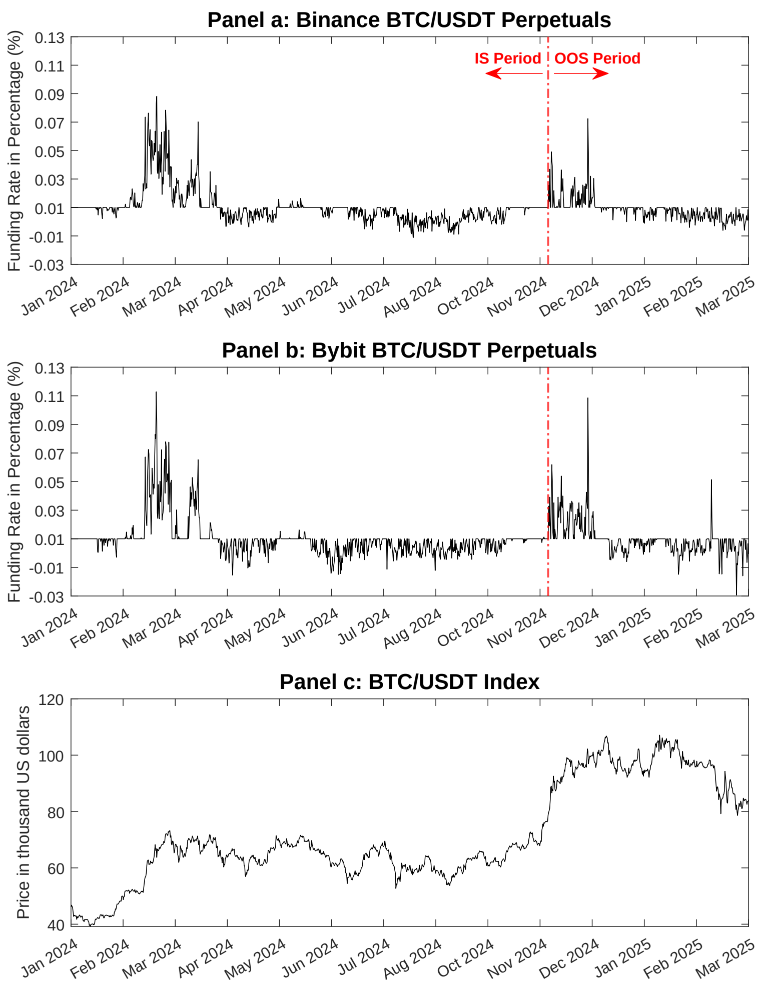

Note: Historical funding rates of the BTC/USDT perpetual futures contracts on Binance (a) and Bybit (b), along with the BTC/USDT spot price index (c). Observations before midnight on November 10, 2024, marked with a red dashed vertical line, constitute the initial in-sample period; observations thereafter are treated as an out-of-sample period.

more costly. As the hype eventually died down, BTC started to fluctuate in a much narrower range of \$52K–\$72K. This was followed by a dramatic decrease in the funding rates, which kept the cost of speculation quite low until the end of October 2024. By mid-November of the same year, BTC started to surge once again, fueled by the growing optimism about the new administration's crypto-friendly agenda. Concurrently, funding rates appear to have increased substantially, creating a trend similar to the previous episode. The occurrence of these repeating patterns presents a great opportunity for predictability analysis and will be exploited in order to split the sample. Observations before midnight on November 10, 2024, marked with a red dashed vertical line on the plots, will be used in the initial in-sample (IS) analysis to perform model selection, and those thereafter will be treated as an OOS period throughout the study.

To examine the distributional characteristics of the funding rates, descriptive statistics are calculated for the full sample as well as for the IS and OOS periods, and documented in Table 3.1. According to these figures, the mean percentage cost of the perpetual contracts was more or less the same on both exchanges, but exhibited greater volatility on Bybit. The table also shows that both funding rate series are highly leptokurtic with positive skewness, providing strong evidence against their normality. The QQ-plots (Appendix D, Figure D.3) confirm the non-normality of the series with a significant divergence observed at the tails. This means that the funding rates exhibit characteristics of heavy tailedness, which is typically observed in financial time series. Given the advantage of the DAR process in handling heavy-tailed data, this further supports its suitability to model the historical funding rates. It is also worth noting here that the statistical properties of the funding rates remain consistent across the subperiods, which is crucial for the robustness of the predictability analysis.

Table 3.1: Descriptive statistics of Binance and Bybit funding rates.

	Obs.	Min	Max	Mean	Median	Std	Kurtosis	Skewness
full sample								
binance	1296	-0.0112	0.0881	0.0098	0.0100	0.0111	15.5394	2.9808
bybit	1296	-0.0294	0.1128	0.0098	0.0100	0.0127	17.0710	3.0157
in-sample (IS)								
binance	912	-0.0112	0.0881	0.0101	0.0100	0.0121	13.8940	2.8692
bybit	912	-0.0154	0.1128	0.0097	0.0100	0.0130	17.0895	3.1920
out-of-sample (OOS)								
binance	384	-0.0061	0.0724	0.0090	0.0100	0.0081	16.4432	2.6671
bybit	384	-0.0294	0.1086	0.0100	0.0100	0.0121	16.8024	2.4814

Note: Each cell reports the statistics of Binance and Bybit funding rate series for the corresponding (sub)samples. Full sample is from January 11, 2024 to March 17, 2025 (1296 observations), in-sample period from January 11, 2024 to November 9, 2025 at 16:00 (912 observations), out-of-sample period from November 10, 2024 at 00:00 to March 17, 2025 at 16:00 (384 observations).

3.5 Forecast Method

The forecast study is conducted in three main parts: model selection, forecasting, and predictability analysis. In the model selection stage, univariate DAR processes of different orders (equation 3.5.1) are fitted to the demeaned series of the funding rates, denoted by $x_t = y_t - \bar{y}$[3], using all $T = 912$ observations within the IS period, as illustrated in Figure 3.1. The optimal forecast model is determined for each series by selecting the smallest lag order p^* that eliminates any serial dependency in the error term η_t. To check for serial dependence, the Ljung-Box test is conducted ex post both on residuals $\hat{\eta}_t$ and squared residuals $\hat{\eta}_t^2$ with a maximum lag of $min(T/5, 10)$ as suggested by Hyndman and Athanasopoulos (2018).

With the selected DAR(p^*) models at hand, quasi-maximum likelihood (QML) estimates of the model parameters are calculated recursively at each point in time over the period $t = T, \ldots, T + 383$ using a rolling window of fixed length $w \in W = \{21, 42, \ldots, 105\}$ as well as with an expanding window. The QML estimates are plugged into equation 3.5.5 to calculate $\hat{x}_{t+1}$, the point forecast for the deviation from the mean at time t+1. These

[3] For ease of notation, y_t will be used interchangably to denote $fr_{BTC/USDT,b,t}$ for $b \in \{\text{Binance, Bybit}\}$ from here on.

forecasted deviations are then used in equation 3.5.6 to obtain the predictions for the one-step-ahead funding rate levels denoted as $\hat{y}_{t+1}$. Each $\hat{y}_{t+1}$ is further transformed into a directional prediction using the function in equation 3.5.7. This yields a total of 384 OOS point forecasts and as many directional predictions for each funding rate series, covering the period from November 10, 2024, at midnight through March 17, 2025, at 16:00.

To gauge the predictability of the funding rates, model-based forecasts are evaluated based on the criteria given in Section 5.4. Statistical significance of the predictive accuracy is assessed by conducting the Clark and West (2007) test on MSPE as well as the Pesaran and Timmermann (1992) test for MDA. In the Clark and West (2007) test, MSPEs from the DAR models are adjusted for estimation bias, and then the adjusted MSPEs are benchmarked against those from the no-change model. The no-change model forecasts $\hat{y}_{t+1} = y_t \; \forall t$, and being able to beat it implies that the funding rates are predictable in terms of forecast errors. For directional predictability, the Pesaran and Timmermann (1992) test is conducted with $m = 3$ given that the funding rates can either go up, or stay constant, or go down in the next period.

Local predictability at time t is captured by the top Lyapunov exponents of the DAR processes calculated as per equations 3.5.11 and 3.5.12. The estimation of the exponents is conducted exclusively based on the expanding window estimates, as rolling counterparts are likely to be more sensitive to noise due to the presence of prolonged periods of mostly constant funding rates. Due to dampening, the funding rates have a bimodal density (Appendix D, Figure D.4), and expanding window estimation is more appropriate for such series to ensure parameter stability, as it retains information from both modes of the distribution.

3.5.1 DAR Model

DAR(p) of Ling (2007) is the extension of the DAR(1) model of Ling (2004) to higher orders and defined by the following equation:

$$x_t = \sum_{i=1}^{p} \phi_i x_{t-i} + \eta_t \sqrt{\omega + \sum_{i=1}^{p} \alpha_i x_{t-i}^2}, \tag{3.5.1}$$

where η_t is an independent and identically distributed (i.i.d.) sequence with mean 0 and variance 1. $\{\phi_i\}_{i=1}^p$ and $\{\alpha_i\}_{i=1}^p$ are autoregressive coefficients capturing the dependence in the conditional mean and conditional variance with $\omega, \alpha_i > 0$.

The model in equation 3.5.1 is flexible in terms of the distribution of errors and relies on a Gaussian-based QML for parameter estimation, as shown in the next subsection. The autoregressive structure in the volatility function allows for autoregressive conditional heteroscedastic (ARCH) effects where larger deviations from the mean process are accompanied by greater volatility. Compared to conventional AR-ARCH models, the DAR(p) model has a larger parameter space and requires only a finite fractional moment on the series x_t to establish the asymptotic normality of the Gaussian QML estimates [see Lin and Zhu (2023) for discussion on DAR-type processes]. In the presence of a unit root, a DAR process with non-zero ARCH effect becomes a stationary martingale with volatility-induced mean-reversion [Gourieroux and Jasiak (2019)].

3.5.2 Parameter Estimation

Let $\Phi = (\phi_1, \ldots, \phi_p)$, $A = (\alpha_1, ..., \alpha_p)$, $z_{1,t} = (x_t, \ldots, x_{t-p+1})'$ and $z_{2,t} = (x_t^2, \ldots, x_{t-p+1}^2)'$. Under Assumption 3.1 in Ling (2007), QML estimators $\hat{\theta} = (\hat{\Phi}, \hat{\omega}, \hat{A})$ are obtained by solving the optimization problem below;

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmax}}\, L_T(\Phi, \omega, A), \tag{3.5.2}$$

where L_T is the Gaussian-based QML function and given by:

$$L_T(\Phi, \omega, A) = \sum_{t=2}^{T} \left(-\frac{1}{2} \ln\left(\omega + A z_{2,t-1}\right) - \frac{1}{2} \frac{\left(x_t - \Phi z_{1,t-1}\right)^2}{\left(\omega + A z_{2,t-1}\right)} \right). \tag{3.5.3}$$

When the error term η_t is normally distributed, $\hat{\theta}_T$ becomes a maximum likelihood estimator and its asymptotic distribution satisfies $\sqrt{n}(\hat{\theta}_T - \theta_0) \xrightarrow{d} N(0, \Omega^{-1})$ such that

$$\Omega = diag\left\{E\left(\frac{z_{1,t}z_{1,t}^{'}}{\omega + Az_{2,t-1}}\right), \frac{1}{2}E\left(\frac{z_{2,t-1}z_{2,t-1}^{'}}{(\omega + Az_{2,t-1})^2}\right)\right\}. \tag{3.5.4}$$

When errors are not normally distributed, the QML estimator is still asymptotically normally distributed if $E(\eta_t^4) < \infty$, and the asymptotic variance is $\Omega^{-1}\Sigma\Omega^{-1}$ with

$$\Sigma = E\left\{\left(\frac{z_{1,t}}{\sqrt{\omega + Az_{2,t-1}}}, \frac{z_{2,t-1}}{(\omega + Az_{2,t-1})}\right)\begin{pmatrix} 1 & E(\eta_t^3) \\ E(\eta_t^3) & E(\eta_t^4) - 1 \end{pmatrix}\begin{pmatrix} \frac{z_{1,t}^{'}}{\sqrt{\omega + Az_{2,t-1}}} \\ \frac{z_{2,t-1}^{'}}{(\omega + Az_{2,t-1})} \end{pmatrix}\right\}.$$

3.5.3 Point Forecasting

Given the assumption $E(\eta_t) = 0$, one-step-ahead point forecasts for the deviations from the mean can be obtained as follows:

$$\hat{x}_{t+1} = \sum_{i=1}^{p} \hat{\phi}_{i,t}\hat{x}_{t-i}, \tag{3.5.5}$$

where $\{\hat{\phi}_{i,t}\}_{i=1}^{p}$ is the series of QML estimates calculated at time t.

Predictions for the future funding rate levels can be obtained easily by adding the forecasted deviations to the corresponding local means:

$$\hat{y}_{t+1} = \hat{x}_{t+1} + \bar{y}_t, \tag{3.5.6}$$

where $\bar{y}_t = \frac{1}{w}\sum_{i=t-w+1}^{t} y_i$ is the (simple) w-day moving average.

3.5.4 Directional Prediction

For a given point forecast $\hat{y}_{t+1}$, the corresponding directional prediction is generated based on the following direction of change (DoC) function;

$$DoC = \begin{cases} 1 & \text{if } \hat{y}_{t+1} - y_t \geq \tau \\ -1 & \text{if } \hat{y}_{t+1} - y_t \leq -\tau \\ 0 & \text{otherwise} \end{cases} \tag{3.5.7}$$

where τ is the threshold that helps determine if the expected change from the current funding rate level is large enough to be categorized as a directional change. For simplicity, this study uses $\tau = 0.0001$, which is the minimum ticker size/resolution used by both Binance and Bybit.

3.5.5 Evaluation Criteria

The performance of model-based predictions is measured using the metrics in equation 3.5.8 and 3.5.9 below;

$$MSPE = \frac{1}{N} \sum_t (y_{t+1} - \hat{y}_{t+1})^2, \tag{3.5.8}$$

$$MDA = \frac{1}{N} \sum_t \boldsymbol{1}_{\{sgn(y_{t+1} - y_t) = sgn(\hat{y}_{t+1} - y_t - \tau)\}}, \tag{3.5.9}$$

where N is the number of forecasted periods.

In terms of forecast errors, the predictability of funding rates requires the MSPE of a DAR model to be statistically significantly smaller than that of the no-change model. However, when forecasting with DAR models, such a direct comparison would not be accurate. This stems from the fact that the no-change model is nested in DAR processes, and as shown in Clark and West (2007), MSPE from a larger model should be adjusted for the upward bias caused by estimating the model parameters.

Let $\hat{y}_{DAR,t+1}$ and $\hat{y}_{nochange,t+1}$ denote the predictions from a DAR model of interest and those from the no-change model, respectively. Then, the adjusted MSPE of a DAR model can be obtained as:

$$MSPE^{DAR}_{adj} = \frac{1}{N}\sum_t (y_{t+1} - \hat{y}_{DAR,t+1})^2 - \frac{1}{N}\sum_t (\hat{y}_{nochange,t+1} - \hat{y}_{DAR,t+1})^2. \tag{3.5.10}$$

Given the relation between the stability and predictability of a dynamical system [Abarbanel et al. (1991); Abarbanel et al. (1992)], the local predictability of the funding rates is evaluated based on the notion of the top Lyapunov exponent of the DAR(p) process provided in equation 2.2 in Ling (2007) as follows;

$$\gamma = \lim_{k\to\infty} \frac{1}{k} \ln \|A_k A_{k-1} \cdots A_1\|, \tag{3.5.11}$$

where $A_t = \begin{bmatrix} \phi_1 + \sqrt{\alpha_1}\,\xi_{1t} & \phi_2 + \sqrt{\alpha_2}\,\xi_{2t} & \dots & \phi_p + \sqrt{\alpha_p}\,\xi_{pt} \\ & I_{p-1} & & 0_{(p-1)\times 1} \end{bmatrix}$ with I_{p-1} being $p-1 \times p-1$ identity matrix and $0_{(p-1)\times 1}$ column vector with (p-1) rows. The innovations ξ_{it} are independent of η_t and normally distributed with mean 0 and variance 1.

The exponent γ is estimated at each point in time by simulating the process with sufficiently large $k = 10^4$. For numerical stability, the matrix A_t will be transformed in each iteration with the QR method of Dieci et al. (2011). This method defines $Z_t = A_t Q_{t-1}$ with $Q_0 = \boldsymbol{I} \in \mathbb{R}^{p\times p}$ (or any orthonormal matrix in general) and applies QR decomposition to establish $Z_t = Q_t R_t$, where Q_t is orthogonal and R_t is upper triangular with positive diagonal elements $(R_t)_{ii}$ for $i = 1, \dots, p$.

After the decomposition, the Lyapunov exponents of the system can be found by

$$\gamma_i = \frac{1}{k}\sum_{j=1}^{k} \ln(R_j)_{ii}, \;\; i = 1, \dots, p \tag{3.5.12}$$

where γ_1 represents the top Lyapunov exponent γ in equation 3.5.11 . A similar approach is used by Valenza et al. (2014) to study heartbeat dynamics with nonlinear autoregressive (NAR) models.

3.6 Results

From the in-sample analysis in Appendix D, the optimal models are found to be DAR(4) and DAR(6) for Binance and Bybit funding rates, respectively. One-step-ahead OOS predictions over the period $t = T + 1, \ldots, T + 384$ are generated from different variants of these models following the steps in Section 5. The model performances are summarized in Table 3.2.

The results show that before adjusting for the estimation bias, only the rolling estimation of the DAR model with certain window lengths produces lower MSPE for Binance funding rates than the no-change model, and no DAR model is able to beat the random walk for Bybit. After the adjustment, all models considered achieve statistically significantly lower MSPE than the no-change model as per the Clark and West test. This means that both Binance and Bybit funding rates are predictable in terms of forecast errors, and the results are robust to the number of observations used in estimating the model parameters. The large differences between the raw and adjusted MSPE values highlight the importance of taking into account the noise introduced into the model forecasts due to parameter estimation. Of the unadjusted MSPEs, about 80-90% for Binance and 72-91% for Bybit are associated with the (upward) estimation bias.

In terms of predicting the direction of the next period's funding rate, the DAR models achieve between 51.82% and 60.68% accuracy for Binance, and between 49.74% and 62.5% for Bybit, depending on the estimation technique and window length used. As per the Pesaran and Timmermann test, each and every one of the forecast models has predictive power with all of the test statistics well above the critical value for 1% significance level. For both funding rate series, expanding window estimation produces the highest MDA by

Table 3.2: OOS performance of forecast models in predicting one-step-ahead funding rates.

Exchange	Evaluation Metric	Forecast Model						
		DAR(4)						no-change (benchmark)
		Expanding Window	Rolling Window					
			w=21	w=42	w=63	w=84	w=105	
Binance	MSPE	46.220	45.160	118.376	39.767	43.370	45.438	45.880
	Adjusted MSPE	9.337 (4.456)	7.399 (4.182)	11.484 (2.626)	4.781 (4.177)	6.298 (4.192)	7.078 (4.459)	NA
	MDA	60.68 (9.838)	51.82 (6.441)	53.13 (7.484)	49.22 (5.487)	52.60 (7.172)	52.86 (7.269)	23.96

Exchange	Evaluation Metric	Forecast Model						
		DAR(6)						no-change (benchmark)
		Expanding Window	Rolling Window					
			w=21	w=42	w=63	w=84	w=105	
Bybit	MSPE	109.962	282.219	111.920	127.355	127.184	102.894	93.169
	Adjusted MSPE	30.870 (3.025)	24.446 (2.73)	29.894 (2.964)	30.994 (2.78)	30.630 (2.812)	28.919 (3.128)	NA
	MDA	62.50 (10.353)	50.26 (5.168)	49.48 (5.133)	52.08 (6.342)	49.74 (5.236)	52.60 (6.719)	22.66

Note: Each cell reports the static measures of the model performances calculated for the entire evaluation period, spanning from November 10, 2024 through March 17, 2025. MSPEs are multiplied by 10^6, and MDAs are expressed as percentages to improve readability and make comparisons across models easier. MSPEs of the DAR models are adjusted for estimation bias using equation 3.5.10, and the corresponding test statistics are calculated as per Clark and West (2007) and provided in the parentheses below. The statistical significance of directional accuracy is evaluated based on the Pesaran and Timmermann (1992) test statistic reported in the parentheses below each MDA estimate.

far, whereas rolling estimation appears to be a better strategy if the objective is to minimize forecast errors.

The predictability of the funding rates is further supported by the results from the local analysis. According to the estimated values in Figure 3.2 (panels b and d), the top Lyapunov exponents are never positive for Binance and almost always negative for Bybit, except for one positive exponent observed on January 29, 2025. A negative exponent indicates that the funding rate was dynamically stable against small perturbations at the time of the estimation, which is necessary for its predictability. In contrast, a positive exponent is a sign of chaos, meaning that small perturbations would not contract but rather expand exponentially. When this is the case, forecasting the funding rates becomes problematic, given that the long-term behavior of chaotic systems cannot be predicted (Wolf et al. , 1985). However, the evidence for chaos in the Bybit series should be taken with a grain of salt, considering that it was only a one-time occurrence with the Lyapunov exponent estimated to be borderline positive ($\hat{\gamma} = 0.0064$). To be exact, this event is detected to have taken place on January 29, 2025, and driven mainly by persistence in the conditional mean with $\sum_{i=1}^{6} \hat{\phi}_i = 1.111$ (Appendix D, Figure D.5).

The evolution of the Lyapunov exponents suggests that the degree of predictability is not homogeneous over time, particularly for the funding rates on Bybit. There are essentially two subperiods where the local exponents become visibly less negative for the exchange, implying a deterioration in predictability: one from the beginning of the OOS sampling period through the end of November 2024, and the other from mid-February 2025 onwards. Although volatility was high on both exchanges during these periods, it was more persistent on Bybit with much larger clusters (Appendix D, Figure D.5). The Lyapunov exponent captures this difference in volatility dynamics by showing larger variation for the historical rates on the exchange than those on Binance. Greater variability in the local estimates of the exponent suggests stronger heterogeneity in the degree of predictability and greater sensitivity to shocks.

Figure 3.2: Local estimates of the top Lyapunov exponent γ

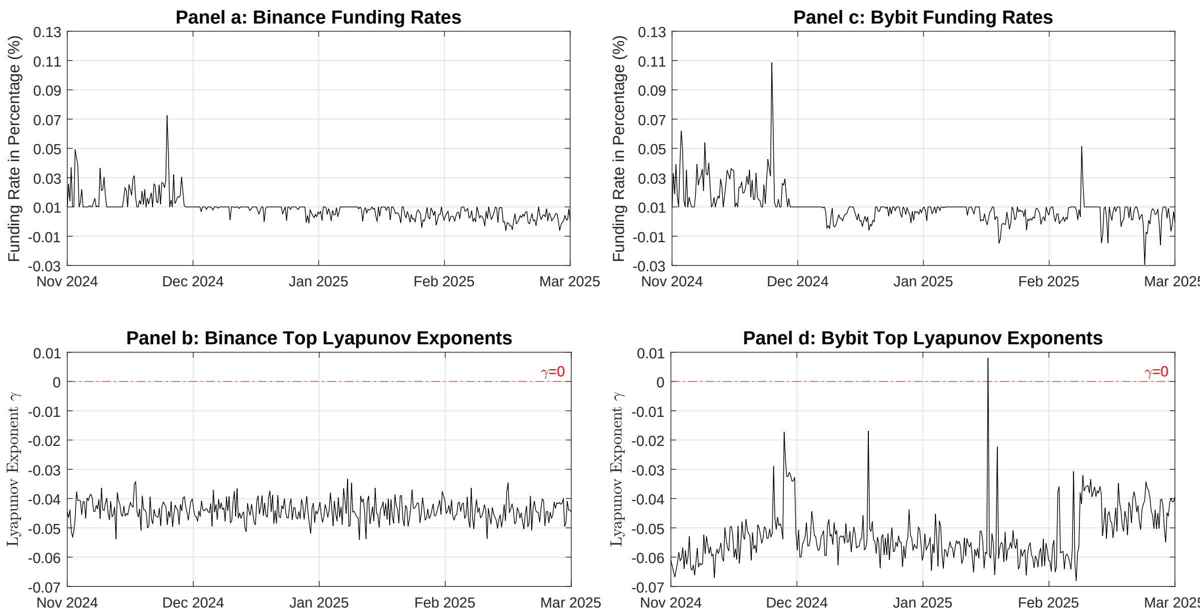

Note: This figure illustrates historical funding rates of BTC/USDT perpetual contracts on Binance (a) and Bybit (c) along with the local estimates of the top Lyapunov exponent γ in panels b and d, respectively. The local estimates of γ are obtained at each point in time by simulating the process in equation 3.5.11 with $k = 10^4$.

3.7 Concluding Remarks

The funding fee mechanism plays a vital role for crypto exchanges in keeping the futures and spot markets of a cryptocurrency aligned. At the center of this mechanism are funding rates, which are contract-specific and vary across exchange platforms. Funding rate levels and their signs have direct implications for the profitability of a trader's position and the effectiveness of arbitrage strategies. Motivated by that, this study examines the OOS predictability of funding rates, particularly that of the BTC/USDT perpetual contracts on Binance and Bybit.

Throughout the analysis, historical funding rates are modeled as a DAR(p) process due to their similarity to the stablecoin dynamics documented in Djogbenou et al. (2023). Point forecasts and directional predictions for the next period's funding rates are computed from forecast models optimized via in-sample analysis, whose parameters are estimated recursively by a rolling window of various lengths as well as an expanding window for comprehensive analysis. Predictability is examined globally using two static metrics, MDA and MSPE, and locally with the help of the top Lyapunov exponents of the DAR processes.

The outcome of the analysis is twofold. First and foremost, OOS point forecasts from DAR models and the corresponding directional predictions are found to be statistically more accurate than the no-change model. This result is also shown to be invariant to the estimation technique used, providing strong evidence for the predictability of the funding rates. Second, the predictability of the historical rates is not only reinforced by the results from the local analysis but also shown to be time-varying, particularly on Bybit. Relatively low predictability is detected for the exchange during periods of heightened volatility and sudden rate jumps, whereas Binance appears to be more resistant to such shocks. This suggests that only sufficiently large and persistent volatility clusters give rise to a significant heterogeneity in the degree of predictability of the funding rates.

This study tackles the predictability of funding rates in the context of autoregressive modeling. Future studies can exploit the (cor)relation between the future and/or spot price of a cryptocurrency and its various funding rates by jointly modeling them. Such predictions are anticipated to be more accurate and better suited for tasks where precision is of the utmost importance. Also, the focal point of this study is the econometric analysis of funding rate predictability. Future studies can evaluate the economic value of funding rate forecasts by designing an arbitrage exercise similar to those described in Section 3.3.

Conclusion

The global crypto market continues its long-term upward growth trend. As of January 2026, the market is up by approximately 2370% from the low of $135B during the early COVID-19 period. Along with this global trend, the market has been experiencing periods of sharp decline and increased volatility due to idiosyncratic events and global macroeconomic uncertainty. In the absence of a universal framework, these local trends have been exacerbated by a shift in the market sentiment associated with the regulatory changes [4] introduced by local governments. In such a dynamic environment, time series methods developed in this thesis can help investors, traders, and policymakers understand the prevailing market structure and make informed decisions.

The methodology developed in Chapter 1 provides the tools to make a comprehensive analysis of stablecoin dynamics and observe their evolution over time under different market conditions. The tvDAR model can be easily re-estimated as new data become available, and local estimates can be utilized along with historical data for inference, price and volatility forecasting, as well as for the dynamic analysis of price stability. When implemented locally, the no-arbitrage test in Chapter 2 would allow practitioners to check if cryptocurrency markets are efficient at any point in time and assess the overall (in)efficiency in a given period. From the underlying model, it is also possible to determine whether fundamental

[4]While cryptocurrency transactions are still banned in China, there have been several crypto-friendly developments in the United States under the Trump administration; 1) the introduction of Bitcoin ETF in January 2024, 2) the establishment of the Strategic Bitcoin Reserve and United States Digital Asset Stockpile in March 2025, and 3) finally the enactment of the Guiding and Establishing National Innovation for U.S. Stablecoins Act (GENIUS Act) in July 2025. In Europe, crypto markets are regulated under the Markets in Crypto-Assets (MiCA or MiCAR) framework.

risk has a significant impact on cryptocurrency returns or not. The evidence in Chapter 3 suggests that univariate DAR models can also be used by cryptocurrency futures traders to generate baseline predictions for one-step-ahead funding rates of Bitcoin perpetual contracts. Besides, simulating the top Lyapunov exponent of the optimal forecast model locally would enable them to quantify the degree of predictability at the time of analysis.

In summary, this dissertation leverages time series analysis to shed light on the state of cryptocurrency markets from late 2017 through early 2025 and addresses some of the present challenges by proposing statistically robust techniques that are adaptable to varying market conditions. Moving forward, future studies can use the proposed methods as a benchmark and measure their effectiveness against alternative methods, including machine learning.

Bibliography

Abarbanel, H. D. I., Brown, R., and Kennel, M. B. (1991). *Local Lyapunov exponents computed from observed data*, Journal of Nonlinear Science, 2, 343–365.

Abarbanel, H. D. I., Brown, R., and Kennel, M. B. (1992). *Variation of Lyapunov exponents on a strange attractor*, Journal of Nonlinear Science, 1, 175–199.

Adachi, M., Da Silva, P. B. P., Born, A., Cappuccio, M., Czák-Ludwig, S., Gschossmann, I., Pellicani, A., Plooij, M., Paula, G., and Philipps, S. (2022). *Stablecoins? Role in crypto and beyond: Functions, risks and policy*, Macroprudential Bulletin, European Central Bank, Vol. 18.

Allen, F., Gu, X., and Jagtiani, J. (2022). *Fintech, cryptocurrencies, and CBDC: Financial structural transformation in China*, Journal of International Money and Finance, 124(C).

Andrews, D. (1987). *Consistency in nonlinear econometric models: A generic uniform law of large numbers*, Econometrica, 55, 1465–1471.

Babalos, V., Bouri, E., and Gupta, R. (2025). *Does the introduction of US spot Bitcoin ETFs affect spot returns and volatility of major cryptocurrencies?* Quarterly Review of Economics and Finance, 102, 102006.

Bask, M. (1996). *Dimensions and Lyapunov exponents from exchange rate series*, Chaos, Solitons & Fractals, 7(12), 2199–2214.

Baumöhl, E., and Vyrost, T. (2020). *Stablecoins as a crypto safe haven? Not all of them!*, ZBW–Leibniz Information Centre for Economics, Kiel–Hamburg.

Bank of England (2022). *Financial stability in focus: Cryptoassets and decentralised finance*, Financial Policy Committee.

Bianchi, D., Rossini, L., and Iacopini, M. (2022). *Stablecoins and cryptocurrency returns: What is the role of Tether?*, Working Paper, University of Milan.

Billingsley, P. (1999). *Convergence of Probability Measures*, Wiley, New York.

Blanchard, O., and Watson, M. (1982). *Bubbles, rational expectations, and financial markets*, in Wachtel, P. (ed.), *Crisis in the Economic and Financial Structure*, Lexington Books.

Borkovec, M. (2000). *Extremal behavior of the autoregressive process with ARCH(1) errors*, Stochastic Processes and Their Applications, 85, 189–207.

Borkovec, M., and Kluppenberg, C. (2001). *The Tail of the Stationary Distribution of the Autoregressive Process with ARCH(1) Errors*, Annals of Applied Probability, 11, 1220–1241.

Buraschi, A., Porchia, P., and Trojani, F. (2010). *Correlation risk and optimal portfolio choice*, The Journal of Finance, 65, 393–420.

Catalini, C., and de Gortari, A. (2021). *On the Economic Design of Stablecoins.* Available at SSRN: `https://ssrn.com/abstract=3899499` or `http://dx.doi.org/10.2139/ssrn.3899499`.

Chandler, G., and Polonik, W. (2012). *Mode identification of volatility in time-varying autoregression*, Journal of the American Statistical Association, 107(499), 1217–1229.

Chandler, G., and Polonik, W. (2017). *Residual empirical processes and weighted sums for time-varying processes with applications to testing for homoscedasticity*, Journal of Time Series Analysis, 38, 72–98.

Chen, M., Li, D., and Ling, S. (2014). *Non-stationarity and quasi-maximum likelihood estimation on a double autoregressive model*, Journal of Time Series Analysis, 35, 189–202.

Chen, J., Lin, D., and Wu, J. (2022). *Do cryptocurrency exchanges fake trading volumes? An empirical analysis of wash trading based on data mining*, Physica A: Statistical Mechanics and its Applications, 586, 126405.

Chen, M., Qin, C., and Zhang, X. (2022). *Cryptocurrency price discrepancies under uncertainty: Evidence from COVID-19 and lockdown nexus*, Journal of International Money and Finance, 124(C).

Chu, J., Zhang, Y., and Chan, S. (2019). *Payments on Digital Platforms: Resiliency, Interoperability and Welfare*, Int. Rev. Financ. Anal., 64, 221–231.

Clark, T. E., and West, K. D. (2007). *Approximately normal tests for equal predictive accuracy in nested models*, Journal of Econometrics, 138, 291–311.

Cline, D. B. H., and Pu, H. H. (2004). *Stability and the Lyapunov exponent of threshold AR–ARCH models*, Annals of Applied Probability, 14, 1920–1949.

Cong, L. W., Li, X., Tang, K., and Yang, Y. (2019) *Crypto Wash Trading*, Technical Report, National Bureau of Economic Research, 30783.

Cox, J., Ingersoll, J., and Ross, S. (1981). *The relation between forward prices and futures prices*, Journal of Financial Economics, 9, 321–346.

Dahlhaus, R. (2000). *A likelihood approximation for locally stationary processes*, Annals of Statistics, 28, 1782–1794.

DAHLHAUS, R., RICHTER, S., and WU, W. (2019). *Towards a general theory for nonlinear locally stationary processes*, Bernoulli, 25, 1013–1044.

DAVIDSON, R. and MACKINNON, J. (2004). *Econometric Theory and Methods.* Oxford University Press.

DAY, W. (1976). A reform of the European currency snake. *IMF Economic Review*, 23, 580–597.

DE BLASIS, R., and WEBB, A. (2022). *Arbitrage, contract design, and market structure in Bitcoin futures markets*, Journal of Futures Markets, 42(3), 492–524.

DECHERT, W. D., and GENCAY, R. (1992). *Lyapunov exponents as a nonparametric diagnostic for stability analysis*, Journal of Applied Econometrics, 7(S1), S41–S60.

DHAWAN, A. and PUTNINS, T. (2020). *A New Wolf in Town? Pump-and-Dump Manipulation in Cryptocurrency Markets*, Working Paper, U of Technology, Sydney.

DIECI, L., JOLLY, M. S., and VAN VLECK, E. S. (2011). *Numerical techniques for approximating Lyapunov exponents and their implementation*, Journal of Computational and Nonlinear Dynamics, 6(1), 011003.

DJOGBENOU, A. INAN, E., and JASIAK, J. (2023). *Time-Varying Coefficient DAR Model and Stability Measures for Stablecoin Prices: An Application to Tether*, Journal of International Money and Finance, Volume 139, Article 102946, December 2023.

DUFFIE, D., FILIPOVIC, D., and SCHACHERMAYER, W. (2003). *Affine processes and applications in finance*, Annals of Applied Probability, 13, 984–1053.

DUFFIE, D., and KAN, R. (1996). *A yield factor model of interest rates*, Mathematical Finance, 6, 379–406.

DUFFIE, D., and SINGLETON, K. J. (1993). *Simulated moments estimation of markov models of asset prices*, Econometrica, 61, 929–952.

European Central Bank (2022). *Decrypting financial stability risks in crypto-asset markets*, Financial Stability Review, https://www.ecb.europa.eu/pub/financial-stability/fsr/special/html/ecb.fsrart202205_02~1cc6b111b4.en.html.

Fan, J., Farmen, M., and Gijbels, I. (1998). *Local maximum likelihood estimation and inference*, Journal of the Royal Statistical Society, Series B, 60, 591–608.

Fernández-Rodríguez, F., Sosvilla-Rivero, S., and Andrada-Félix, J. (2005). *Testing chaotic dynamics via Lyapunov exponents*, Journal of Applied Econometrics, 20(7), 911–930.

Force, ECB Crypto-Assets Task (2020). *Stablecoins: Implications for monetary policy, financial stability, market infrastructure and payments*, Occasional Paper Series, European Central Bank, Vol. 247.

Froot, K., and Obstfeld, M. (1991). *Intrinsic bubbles: The case of stock prices*, American Economic Review, 81, 1189–1214.

Gallant, A. R., and Tauchen, G. (1996). *Which Moments to Match?*, Econometric Theory, 12(4), 657–681.

Gao, J., Hu, J., Tung, W.-W., and Zheng, Y. (2011). *Multiscale analysis of economic time series by scale-dependent Lyapunov exponent*, Quantitative Finance, 13(2), 265–274.

Gourieroux, C. (1997). *ARCH Models and Financial Applications*, Springer, New York.

Gourieroux, C., and Jasiak, J. (2006). *Autoregressive gamma processes*, Journal of Forecasting, 25, 129–152.

Gourieroux, C., and Jasiak, J. (2019). *Robust analysis of the martingale hypothesis*, Econometrics and Statistics, 9, 17–41.

Gourieroux, C., and Monfort, A. (2002). *Simulation-based econometric methods*, Oxford university press.

Gourieroux, C., and Monfort, A. (2007). *Econometric specification of stochastic discount factor models*, Journal of Econometrics, 136, 509–530.

Gourieroux, C., Monfort, A., and Renault, E. (1993). *Indirect Inference*, Journal of Applied Econometrics, 8(S1), S85–S118.

Gourieroux, C., Jasiak, J., and Sufana, R. (2009). *The Wishart Autoregressive process of multivariate stochastic volatility*, Journal of Econometrics, 150, 167–181.

Gourieroux, C., Monfort, A., and Sufana, R. (2009). *International money and stock market contingent claims*, Journal of International Money and Finance, 29, 1727–1751.

Gourieroux, C., and Sufana, R. (2010). *Derivative pricing with Wishart multivariate stochastic volatility*, Journal of Business and Economic Statistics, 28, 438–451.

Griffin, J. M., and Shams, A. (2020). *Is Bitcoin Really Untethered?* The Journal of Finance, 75(4).

Hansen, L. P. (1982). *Large sample properties of generalized method of moments estimators*, Econometrica, 50, 1929–1954.

Ho, A., Darbha, S., Gorelkina, Y., and Garcia, A. (2022). *The relative benefits and risks of stablecoins as a means of payment*, Staff Discussion Papers, Bank of Canada, Vol. 21.

Hyndman, R. J., and Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*, OTexts.

International Monetary Fund (2022). *Cryptic connections: Spillovers between crypto and equity markets*, Global Financial Stability Notes.

Kronick, J., and Zelmer, M. (2023). *Stablecoins: Sailing without a rudder*, C.D. Howe Institute Commentary, Vol. 637.

LeBaron, B. (1994). *Chaos and nonlinear forecastability in economics and finance*, Philosophical Transactions of the Royal Society A, 348, 397–404.

Lee, B.-S., and Ingram, B. F. (1991). *Simulation Estimation of Time-series Models*, Journal of Econometrics, 47, 195–205.

Le Pennec, G., Fiedler, I., and Ante, L. (2021). *Wash trading at cryptocurrency exchanges*, Finance Research Letters, 43, 101982.

Li, W. K. (1992). *On the asymptotic standard errors of residual autocorrelations in nonlinear time series modeling*, Biometrika, 79, 435–437.

Li, W. K., and Mak, T. K. (1994). *On the squared residual autocorrelations in nonlinear time series*, Journal of Time Series Analysis, 15, 627–636.

Li, Y., and Mayer, S. (2022). *Money creation in decentralized finance: A dynamic model of stablecoin and crypto shadow banking*, Working Paper.

Li, T., Shin, D., and Wang, B. (2018). *Cryptocurrency Pump-and-Dump Schemes*, available at SSRN.

Lin, Y., and Zhu, Q. (2023). *On vector linear double autoregression*, Journal of Time Series Analysis, 45(3), 376–397.

Ling, S. (2004). *Estimation and Testing Stationarity for Double Autoregressive Models*, JRSS Series B, 66, 63–78.

Ling, S. (2007). *A Double AR(p) Model: Structure and Estimation?, Statistica Sinica*, Vol. 17, No. 1., 161–175.

Ling, S. and Li, D. (2008). *Asymptotic Inference for a Nonstationary Double AR(1) Model*, Biometrika , 95, 1, pp. 257–263.

Liu, F., Li, D., and Kang, X. (2018). *Sample path properties of an explosive double autoregressive model*, Econometric Reviews, 37, 484–490.

Liu, Y., Tsyvinski, A., and Wu, X. (2019) *Common Risk Factors in Cryptocurrency*, Technical Report, National Bureau of Economic Research, 25882.

Lyons, R. K., and Viswanath-Natraj, G. (2020). *What keeps stablecoins stable?*, NBER Working Paper No. 27136.

McFadden, D. (1989). *A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration*, Econometrica: Journal of the Econometric Society, 995–1026.

Nelson, D. B. (1990). *Stationarity and persistence in the GARCH(1,1) model*, Econometric Theory, 6, 318–334.

Newey, W. K., and West, K. D. (1987). *A Simple, Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix*, Econometrica 55, 703–708.

Newey, W. K., and West, K. D. (1994). *Automatic Lag Selection in Covariance Matrix Estimation*, The Review of Economic Studies, 61(4), 631–653.

Pesaran, M. H., and Timmermann, A. (1992). *A simple nonparametric test of predictive performance*, Journal of Business & Economic Statistics, 10(4), 461–465.

President's Working Group on Financial Markets (2021). *Report and recommendations on stablecoins*, https://home.treasury.gov/news/press-releases/jy0454.

Resende, M., and Zeidan, R. M. (2008). *Expectations and chaotic dynamics: Empirical evidence on exchange rates*, Economics Letters, 99(1), 33–35.

Ruge-Murcia, F. (2012). *Estimating Nonlinear DSGE Models by the Simulated Method of Moments: With an Application to Business Cycles*, Journal of Economic Dynamics and Control, 36, 914–938.

Sprott, J. C. (2003). *Chaos and Time-Series Analysis*, Oxford University Press.

Sprott, J. C. (2014). *Numerical calculation of largest Lyapunov exponent*, Working Paper, University of Wisconsin.

Tran, V.L., and Leirvik, T. (2020) *Efficiency in the markets of crypto-currencies*, Finance Research Letters 35 101382

Urquhart, A. (2016) *The Inefficiency of Bitcoin*, Economic Letters, 148, 80–82.

Valenza, G., Citi, L., and Barbieri, R. (2014). *Estimation of instantaneous complex dynamics through Lyapunov exponents: A study on heartbeat dynamics*, PLoS ONE, 9(8), e105622.

Vidal-Tomas, D., and Ibanez, A. (2018). *Semi-Strong Efficiency of Bitcoin*. Finance Res. Lett. https://doi.org/10.1016/j.frl.2018.013

Wang, G., Ma, X., and Wu, H. (2020). *Are stablecoins truly diversifiers, hedges, or safe havens?*, Research in International Business and Finance, 54, 101225.

Wolf, A., Swift, J. B., Swinney, H. L., and Vastano, J. A. (1985). *Determining Lyapunov exponents from a time series*, Physica D: Nonlinear Phenomena, 16(3), 285–317.

Zakoian, J. M. (1994). *Threshold heteroskedastic models*, Journal of Economic Dynamics and Control, 18, 931–955.

Zargar, F.N., and Kumar, D. (2019). *Informational Inefficiency of Bitcoin:A study based on high-frequency data*, Res. Int. Bus. Finance 47, 344–353.

Appendix A

Technical Appendix for Paper 1

A.1 Proof of Proposition 1

Because of the symmetry of the density of η, the Lyapunov exponent is an even function of ϕ. Hence we can suppose that $\phi > 0$ to find the expression of the Lyapunov exponent, and then replace ϕ by $|\phi|$.

For $\phi > 0$ we have:

$$\lambda(\phi, \alpha) = E \ln|\phi + \eta\sqrt{\alpha}| = \int \ln|\phi + \eta\sqrt{\alpha}|\psi(\alpha)d\alpha.$$

Let us assume that $\eta \sim U_{[-1,1]}$. Then, its density is $\psi(\eta) = \frac{1}{2}1_{\eta\in[-1,1]}$.

We have:

$\lambda(\phi, \alpha) = \frac{1}{2}\int_{-1}^{1} \ln|\phi + \eta\sqrt{\alpha}|d\alpha.$

To complete the proof, we first write $\lambda(\phi, \alpha)$ without the absolute value. Next, we compute the integral for $\phi/\sqrt{\alpha} > 1$ and $\phi/\sqrt{\alpha} < 1$.

Step 1: We observe that:

$$\begin{aligned} \phi + \eta\sqrt{\alpha} > 0 \quad &\Longleftrightarrow \quad \eta > -\phi/\sqrt{\alpha}, \\ \phi + \eta\sqrt{\alpha} < 0 \quad &\Longleftrightarrow \quad \eta < -\phi/\sqrt{\alpha}. \end{aligned}$$

Hence,

$$\lambda(\phi,\alpha) = \frac{1}{2}\int_{-1}^{1} 1_{\eta>-\phi/\sqrt{\alpha}} \ln(\phi+\eta\sqrt{\alpha})d\eta + \frac{1}{2}\int_{-1}^{1} 1_{\eta<-\phi/\sqrt{\alpha}} \ln(-\phi-\eta\sqrt{\alpha})d\eta$$

Step 2: Let us now examine the two cases:

a) If $\phi/\sqrt{\alpha} > 1 \iff -\phi/\sqrt{\alpha} < -1$, we get:

$\lambda(\phi,\alpha) \;= \frac{1}{2}\int_{-1}^{1}\ln(\phi+\eta\sqrt{\alpha})d\eta + \frac{1}{2}0 = \frac{1}{2}\frac{1}{\sqrt{\alpha}}\int_{-1}^{1}\ln(\phi+\eta\sqrt{\alpha})d(\sqrt{\alpha}\eta)$

$= \frac{1}{2\sqrt{\alpha}}\int_{\phi-\sqrt{\alpha}}^{\phi+\sqrt{\alpha}}\ln(u)du$ with change of variable $u = \phi+\eta\sqrt{\alpha}$

$= \frac{1}{2\sqrt{\alpha}} u\ln(u) - u|_{\phi-\sqrt{\alpha}}^{\phi+\sqrt{\alpha}}$

$= \frac{1}{2\sqrt{\alpha}}\left[(\phi+\sqrt{\alpha})\ln(\phi+\sqrt{\alpha}) - (\phi+\sqrt{\alpha}) - (\phi-\sqrt{\alpha})\ln(\phi-\sqrt{\alpha}) + (\phi-\sqrt{\alpha})\right]$

b) If $\phi/\sqrt{\alpha} < 1 \iff -\phi/\sqrt{\alpha} > -1$, we get:

$\lambda(\phi,\alpha) = \frac{1}{2}\int_{-\phi/\sqrt{\alpha}}^{1}\ln(\phi+\eta\sqrt{\alpha})d\eta + \frac{1}{2}\int_{-1}^{-\phi/\sqrt{\alpha}}\ln(-\phi-\eta\sqrt{\alpha})d\eta$

$= \frac{1}{2}\frac{1}{\sqrt{\alpha}}\int_{0}^{\phi+\sqrt{\alpha}}\ln(u)du + \frac{1}{2}\int_{\phi/\sqrt{\alpha}}^{1}\ln(-\phi+\eta\sqrt{\alpha})d\eta$ with the change of variable $u = \phi+\eta\sqrt{\alpha}$

$= \frac{1}{2\sqrt{\alpha}}\int_{0}^{\phi+\sqrt{\alpha}}\ln(u)du + \frac{1}{2\sqrt{\alpha}}\int_{0}^{-\phi+\sqrt{\alpha}}\ln(u)du$

$= \frac{1}{2\sqrt{\alpha}}\left[(\phi+\sqrt{\alpha})\ln(\phi+\sqrt{\alpha}) - (\phi+\sqrt{\alpha}) - [(-\phi+\sqrt{\alpha})\ln(-\phi+\sqrt{\alpha})] + (-\phi+\sqrt{\alpha})\right]$

Step 3: By putting the two expressions in a) and b) together we get for $\phi > 0$:

$$\lambda(\phi,\alpha) = \frac{1}{2\sqrt{\alpha}}\left[(\phi+\sqrt{\alpha})\ln(\phi+\sqrt{\alpha}) - (\phi+\sqrt{\alpha}) - |\phi-\sqrt{\alpha}|\ln|\phi-\sqrt{\alpha})| + |\phi-\sqrt{\alpha}|\right]$$

The general expression of the Lyapunov Exponent without the sign constraint on ϕ is:

$$\lambda(\phi,\alpha) = \frac{1}{2\sqrt{\alpha}}\left[(|\phi|+\sqrt{\alpha})\ln(|\phi|+\sqrt{\alpha}) - (|\phi|+\sqrt{\alpha}) - ||\phi|-\sqrt{\alpha}|\ln||\phi|-\sqrt{\alpha})| + ||\phi|-\sqrt{\alpha}|\right]$$

We observe a non-differentiability in $\phi = \pm\sqrt{\alpha}$. Moreover, for $\phi = 0$, we get a zero value of the Lyapunov exponent:

$$\lambda(\phi,\alpha)=\frac{1}{2\sqrt{\alpha}}\left[\sqrt{\alpha}\ln\sqrt{\alpha}-\sqrt{\alpha}\ln\sqrt{\alpha}+\sqrt{\alpha}\right]=0$$

A.2 Proof of Proposition 2

We need to show that when $T\to\infty$, then $\int\ln[\hat{\phi}_T+\eta\sqrt{\hat{\alpha}_T}]\psi_0(\eta)d\eta\to\int\ln(\phi_0+\eta\sqrt{\alpha_0})\psi_0(\eta)d\eta$ in probability if $(\hat{\phi}_T,\hat{\alpha}_T)\to(\phi_0,\alpha_0)$ in probability, and if condition A.6 of Proposition 2:

$$\exists\delta>0,\text{such that}\int sup_{\substack{\phi_0-\delta<\phi<\phi_0+\delta\\ \alpha_0-\delta<\alpha<\alpha_0+\delta}}|\ln|\phi+\eta\sqrt{\alpha}||\psi_0(\eta)d\eta<\infty\quad\text{(a.1)}$$

holds.

Proof of convergence:

If $(\hat{\phi}_T,\hat{\alpha}_T)\to(\phi_0,\alpha_0)$ in probability, then they also converge in distribution. It follows from the Skorokhod theorem that up to a change of probability space, we can assume that the almost sure (a.s.) convergence also holds [Billingsley (1999)]. Therefore, if g is a continuous function of (ϕ,α), we have $g(\hat{\phi}_T,\hat{\alpha}_T)\to g(\phi_0,\alpha_0)$ a.s. and in distribution in that new space. Then, $g(\hat{\phi}_T,\hat{\alpha}_T)\xrightarrow{\text{d}}g(\phi_0,\alpha_0)$ in the initial space and also in probability because the limit is constant, we get the "in probability" version of the continuous mapping theorem. Therefore, we need only the condition ensuring that $g(\phi,\alpha)=\int\ln|\phi+\eta\sqrt{\alpha}|\psi_0(\eta)d\eta$ is continuous. Condition (A.6) ensures the continuity of integral function g, which follows from the dominated convergence theorem.

A.3 Proof of Proposition 3

To show that the estimator $\hat{\lambda}_{2,T}=G_T(\hat{\theta}_T)$ converges in probability to $\lambda_0=G(\theta_0)$, we first establish a general lemma on the convergence of a function $G_T(\cdot)$ of a sequence of estimators

$\hat{\theta}_T$ that converges in probability to an unknown parameter θ_0. Second, we apply it to the DAR model and Lyapunov estimator $\hat{\lambda}_{2,T}$.

Step 1: This step outlines and proves the general lemma.

Lemma:

Let us consider a sequence $G_T(\theta)$ of stochastic functions of θ, $\theta \in \Theta$, and a sequence of estimators $\hat{\theta}_T$. We assume that:

i) Θ is compact and θ_0 is in the interior of Θ.

ii) $\hat{\theta}_T$ tends in probability to θ_0.

iii) $G_T(\theta)$ tends in probability to a limit $G(\theta)$, $\forall \theta \in \Theta$.

iv) Sufficient Lipschitz condition for stochastic equicontinuity:

There exists a stochastic sequence B_T with $B_T = O_p(1)$ and an increasing function $h : [0, \infty) \to [0, \infty)$ continuous at zero, with $h(0) = 0$ and such that for all $\tilde{\theta}$, $\theta \in \Theta$, $|G_T(\tilde{\theta}) - G_T(\theta)| \leq B_T h(d(\tilde{\theta}, \theta))$.

Step 2: Then, $\hat{G}_T(\hat{\theta}_T)$ tends in probability to $G(\theta_0)$.

Proof:

We have :

$$\begin{aligned} |G_T(\hat{\theta}_T) - G(\theta_0)| &= |G_T(\hat{\theta}_T) - G_T(\theta_0) + G_T(\theta_0) - G(\theta_0)| \\ &\leq |G_T(\hat{\theta}_T) - G(\theta_0)| + |G_T(\theta_0) - G(\theta_0)| \\ &\leq B_T h[d(\hat{\theta}_T, \theta_0)] + |G_T(\theta_0) - G(\theta_0)|. \end{aligned}$$

We know that if $X_T \overset{\mathrm{P}}{\to} 0$, $Y_T \overset{\mathrm{P}}{\to} 0 => X_T + Y_T \overset{\mathrm{P}}{\to} 0$, i.e. the sum of $o_p(1)$ is $o_p(1)$.

Under condition iii) $|G_T(\theta_0) - G(\theta_0)| = o_p(1)$. It remains to be shown that $B_T h[d(\hat{\theta}_T, \theta_0)]$ is $o_p(1)$.

We have:

$$[B_T < M \text{ and } h[d(\hat{\theta}_T, \theta_0)] < \epsilon/M] => [B_T h[d(\hat{\theta}_T, \theta_0)]] < \epsilon]$$

$$\iff \Big((B_T < M) \cap [h[d(\hat{\theta}_T, \theta_0)] < \epsilon/M]\Big) \subset [B_T h[d(\hat{\theta}_T, \theta_0)] < \epsilon]$$

Consider the complement: $(A \cap B)^c = A^c \cup B^c$. We get:

$$((B_T > M) \cup (h[d(\hat{\theta}_T, \theta_0)] > \epsilon/M) \supset [B_T h[d(\hat{\theta}_T, \theta_0)] > \sqrt{\epsilon}]$$

It follows that:

$$\begin{aligned} P[B_T h[d(\hat{\theta}_T, \theta_0)] > \epsilon] &\leq P[(B_T > M) \ \cup \ (h[d(\hat{\theta}_T, \theta_0)] > \epsilon/M)] \\ &\leq P[B_T > M] + P[h[d(\hat{\theta}_T, \theta_0)] > h^{-1}(\epsilon/M)], \end{aligned}$$

because $P[A \cup B] \leq P(A) + P(B)$.

Then, for any ϵ, we can choose a value of M and a number of observations T sufficiently large to get $P[B_T h[d(\hat{\theta}_T, \theta_0)] > \epsilon]$ arbitrarily small. Therefore, $B_T h[d(\hat{\theta}_T, \theta_0)]$ tends to zero in probability. QED

Then, the lemma can be applied with $G_T(\theta) = \frac{1}{T}\sum_{t=2}^{T} g(x_t, x_{t-1}; \theta)$ and $g(x_t, x_{t-1}; \theta) = \ln|\phi + \frac{x_t - \phi x_{t-1}}{\sqrt{\omega + \alpha x_{t-1}^2}}\sqrt{\alpha}|$. Under assumptions A.1, A.4, A.7, conditions i), ii), iii) are satisfied. For example:

$$G_T(\theta) \xrightarrow{\text{P}} E_0 g(x_t, x_{t-1}; \theta),$$

by the weak law of large numbers applied to the transformation $g(x_t, x_{t-1}; \theta)$ of the ergodic stationary process (x_t). Assumption A.8 corresponds to condition iv) of the lemma.

A.4 Proof of Proposition 4

This proof has two parts. The first part shows the convergence in probability of the proposed estimator of stability measure $\hat{\phi}_T^2 + \hat{\alpha}_T$ to the true value $\phi_0^2 + \alpha_0$. The second part establishes the convergence in distribution of $\hat{\xi}_T$.

a) Proof of convergence

We need to show that $\hat{\phi}_T^2 + \hat{\alpha}_T \to \phi_0^2 + \alpha_0$ in probability if $(\hat{\phi}, \hat{\alpha}) \to (\phi_0, \alpha_0)$ in probability when $T \to \infty$.

This is a consequence of the "in probability" version of the continuous mapping theorem given in Appendix A.2

b) Proof of Normality

The Taylor series expansion pre-multiplied by $\sqrt{T}$ implies:

$$\sqrt{T}\left[(\hat{\phi}_T^2+\hat{\alpha}_T)-(\phi_0^2+\alpha_0)\right]=\begin{pmatrix}2\phi_0\\ 1\end{pmatrix}'\sqrt{T}\begin{pmatrix}\hat{\phi}_T-\phi_0\\ \hat{\alpha}_T-\alpha_0\end{pmatrix}+o_p(1)$$

$$=A'\sqrt{T}\begin{pmatrix}\hat{\phi}_T-\phi_0\\ \hat{\alpha}_T-\alpha_0\end{pmatrix}+o_p(1)$$

where $A'=[2\phi_0\ \ 1]$. We get the asymptotic normal distribution of $\hat{\xi}_T$:

$$\sqrt{T}(\hat{\xi}_T-\xi_0)\sim N(0,V_\xi),$$

where $V_\xi = A'\Omega^* A$. The matrix Ω^* is: $\Omega^* = diag(\Sigma^{-1}, V(\hat{\alpha}))$ where $\Sigma = E_0(y^2/(\omega_0 + \alpha_0 y^2)$ given in Section 4.3.1. Matrix $V(\hat{\alpha}) = (E_0\frac{1}{(\omega_0+\alpha_0 y^2)^2})/\tilde{V}_0(y^2)$ and $\tilde{V}_0(y^2) = \tilde{E}_0(y^4) - (\tilde{E}_0 y^2)^2$. In this formula, $\tilde{E}_0$ denotes the expectation of variables $y^4\frac{1}{(\omega_0+\alpha_0 y^2)^2}/E_0\frac{1}{(\omega_0+\alpha_0 y^2)^2}$ and $y^2\frac{1}{(\omega_0+\alpha_0 y^2)^2}/E_0\frac{1}{(\omega_0+\alpha_0 y^2)^2}$ [see, Section 4.3.1 and Ling (2004) for the variance estimator formula].

Appendix B

Simulation Results for Paper 1

Figure B.1: The Lyapunov exponent.

We use the analytical expression established in Proposition 1 and plot the Lyapunov exponent $\lambda = E(ln(|\phi + \sqrt{\alpha}\eta|))$ for different values of parameters $\phi \in [-1, 1]$ and $\alpha \in [0, 1]$ and uniformly distributed errors $\eta \sim U[-1, 1]$. We consider $\phi = \{-1, -0.8, -0.6, \ldots, 0.6, 0.8, 1\}$ and $\alpha = \{0, 0.1, 0.2, 0.3, \ldots, 0.8, 0.9, 1\}$. Figure B.1 shows that the Lyapunov exponent λ remains less than zero and it is not everywhere differentiable.

Figure B.2: Asymptotic density of stability measure estimators $\hat{\lambda}_{2,T}$ and $\hat{\xi}_T$ in finite sample.

We consider the error density $\eta \sim N(0, 1)$ and the true parameter values $\phi_0 = 0.7$, $\alpha_0 = 0.5$, $\omega_0 = 0.01$. Note that these parameter values are close to those estimated from Tether prices in our application. The kernel-smoothed densities of stability measure estimators are based on $4,000$ simulations. Figure B.2 plots the density of the Lyapunov exponent estimator $\hat{\lambda}_{2,T}$ and the density of stability measure $\hat{\xi}_T$ when the model parameters are estimated from a sample of size T and plugged in. The results in panel (a) show that for the selected parameter values, the estimated $\hat{\lambda}_{2,T}$ is below zero for sample sizes $T = 50$ or $T = 100$, implying valid inference. Panel (b) of Figure B.2 shows that the density of the estimated alternative stability measure $\hat{\xi}_T = \hat{\phi}_T^2 + \hat{\alpha}_T$ is centered at $\xi_0 = \phi_0^2 + \alpha_0 = 0.99$, which is the true value of ξ.

Figure B.3: Asymptotic densities of DAR estimators in finite sample.

Figure B.3, shows the simulation based density of estimators $\hat{\phi}_T$, $\hat{\alpha}_T$ and $\hat{\omega}_T$ from samples of size $T = 50$ and $T = 100$ obtained with 4,000 replications. The three panels in Figure B.3 provide evidence that the parameters are close to normally distributed and their densities are centered at the true parameter values. The accuracy is improved when the sample size increases from $T = 50$ to $T = 100$ and the variance of estimators diminishes.

Figure B.1: Lyapunov exponent $\lambda = E(ln(|\phi + \sqrt{\alpha}\eta|))$ as a function of ϕ and α for $\eta \sim U[-1, 1]$

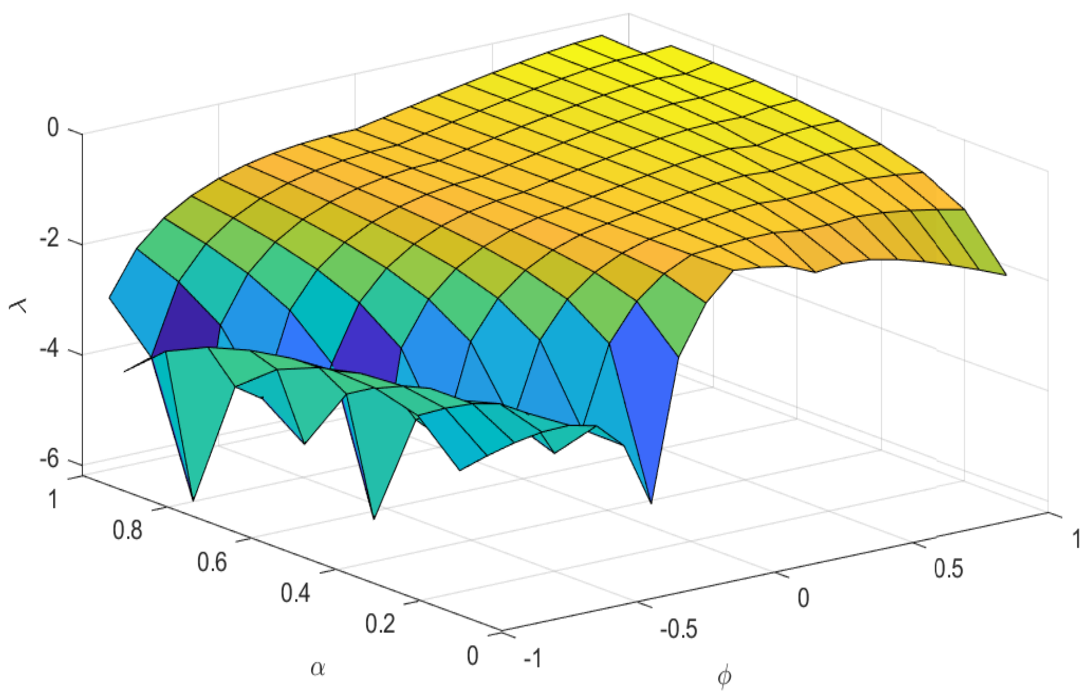

Note: This figure presents the Lyapunov exponent as a function of parameters $\phi \in [-1, 1]$ and $\alpha \in [0, 1]$ and uniformely distributed errors: $\eta \sim U[-1, 1]$.

Figure B.2: Densities of plug-in stability measure estimators in finite sample

(a) Density of the Lyapunov exponent estimator $\hat{\lambda}_{2,T}$

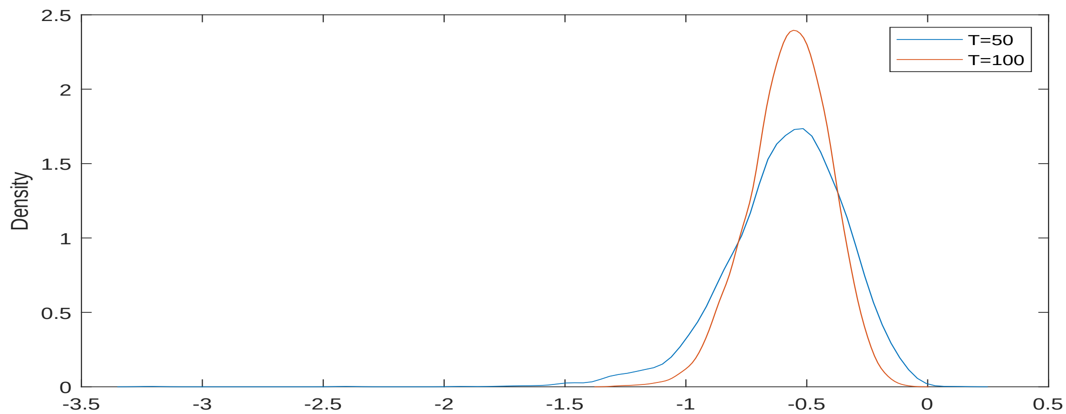

(b) Density of stability measure estimator $\hat{\xi}_T$

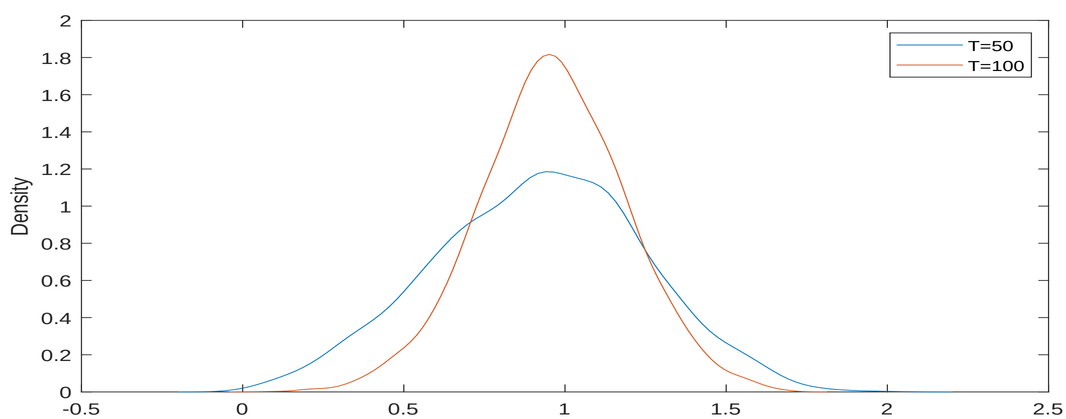

Note: This figure shows the densities of plug-in stability measure estimators in finite sample. The top panel shows the density of the Lyapunov exponent estimator $\hat{\lambda}_{2,T}$, based on T=50 (blue line) and T=100 (red line). The bottom panel presents the density of stability measure estimator $\hat{\xi}_T$, based on T=50 (blue line) and T=100 (red line). All results are based on $4,000$ simulations.

Figure B.3: Densities of estimators of ϕ, α and ω

(a) Density of $\hat{\phi}_T$

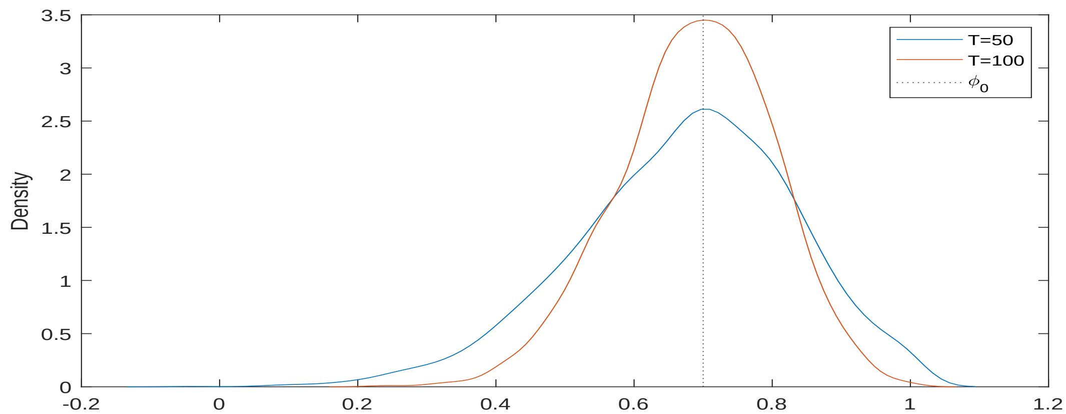

(b) Density of $\hat{\omega}_T$

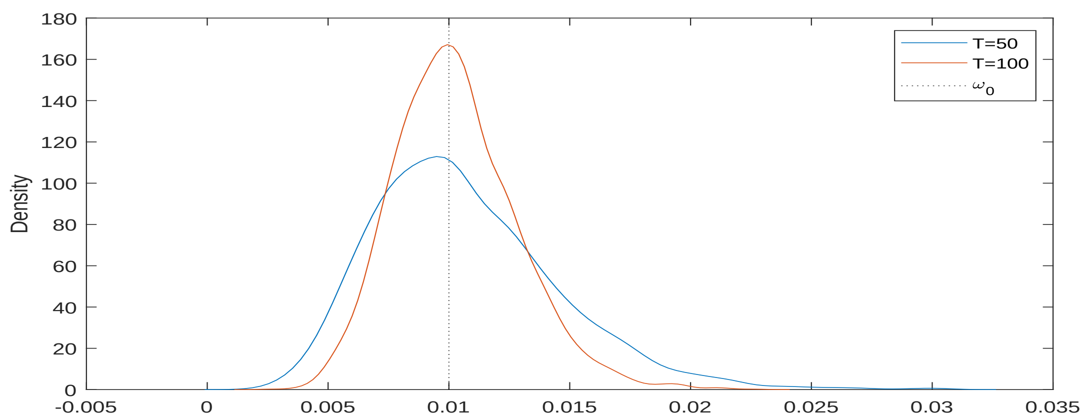

(c) Density of $\hat{\alpha}_T$

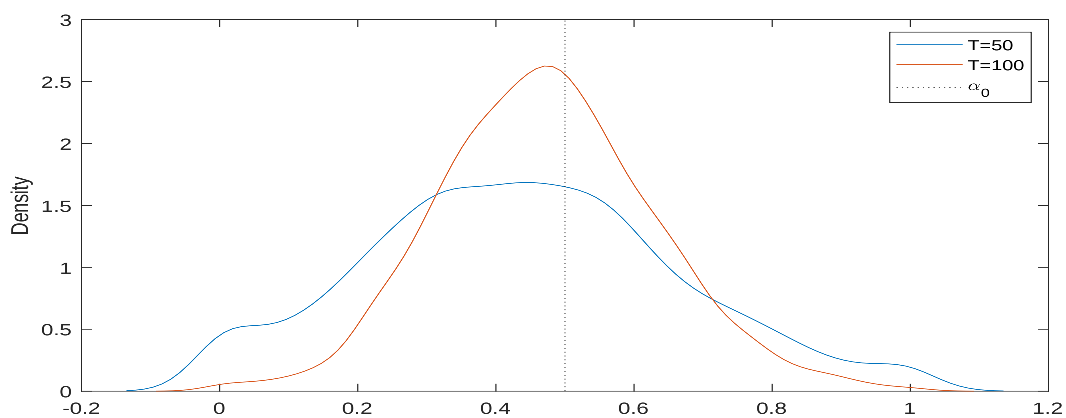

Note: This figure shows the densities of DAR parameter estimators in finite sample for T=50 (blue line) and T=100 (red line) based on 4,000 simulations. The three panels displays the densities of $\hat{\phi}_T, \hat{\omega}_T, \hat{\alpha}_T$, respectively.

Appendix C

Additional Empirical Results for Paper 1

C.1 Additional Empirical Results on Tether

Figure C.1 presents the autoregressive coefficient of the tvDAR model fitted to Tether prices (top panel), square root of conditional variance of tvDAR (middle panel), and stability measure $\hat{\xi}$ (bottom panel) along with the 95% confidence level. The shaded areas indicate local instability when the confidence interval includes value 1 (red line) and the null hypothesis $\xi \geq 1$ cannot be rejected. The stability measure in the bottom panel reveals that the conditional heteroskedasticity effect contributed the third instability period in September 2020 in addition to the high persistence in Tether price. This period can also be linked to higher local volatility in the data observed in Figure 1.4. There is no specific event, we can associate with this movement. We note that the proposed tvDAR model and stability measure allow for capturing these local dynamics.

The measure of stability ξ plotted in the bottom panel of Figure C.1 is more conservative than the Lyapunov exponent λ. We can have $\xi \geq 1$ while $\lambda < 0$ so that valid inference is still possible. When the confidence interval includes value 1 (red line), the null hypothesis

$\xi \geq 1$ cannot be rejected. This should be interpreted as a warning to investors, regulators or stablecoin issuers to be cautious during these periods.

Figure C.1: Tether: tvDAR parameters $\phi(t/T)$, square root of conditional variance and stability measure $\xi(t/T)$

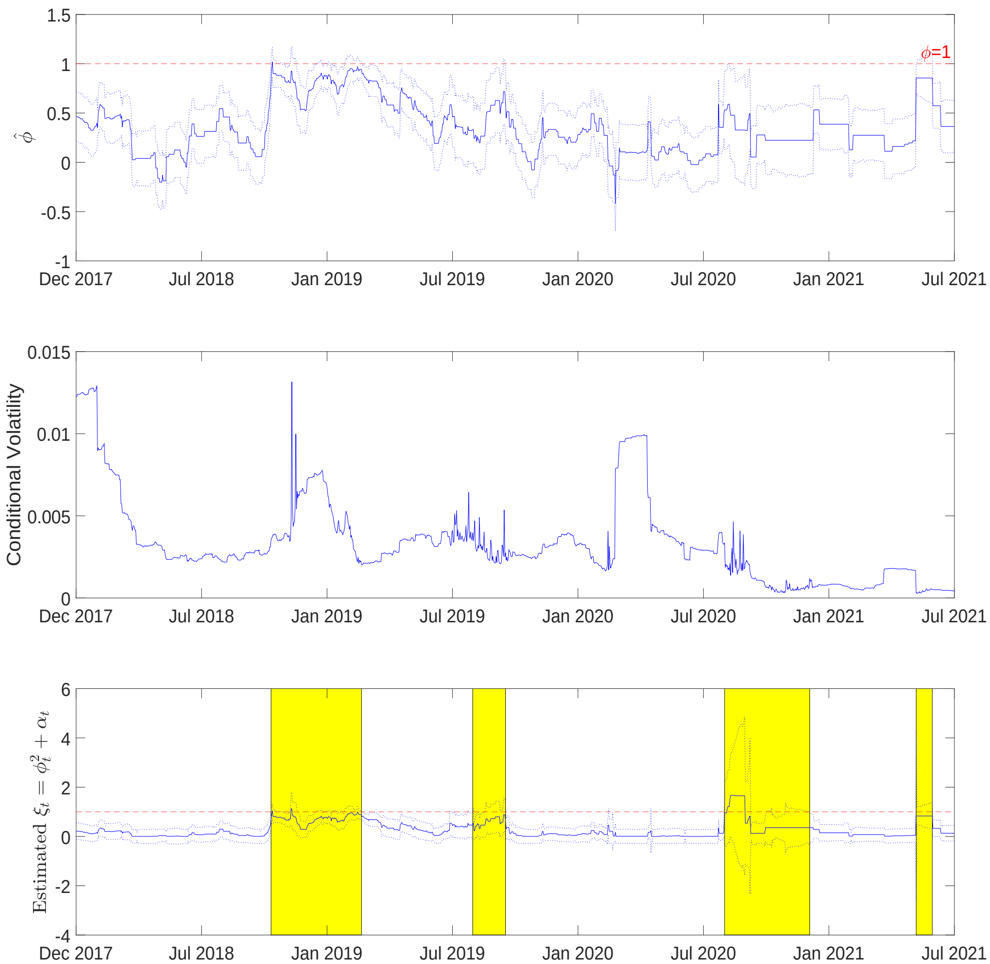

Note: The first panel presents the estimated tvDAR autoregressive coefficient and its 95% confidence interval. The second panel shows the estimated square root of tvDAR conditional variance. The TvDAR estimates are based on Epanechnikov kernel, bandwidth = (50/T). The third panel presents the estimates of stability measure ξ with its 95% confidence interval. Rolling window of 50 observations. Shaded areas of local instability.

C.2 Robustness Check: Empirical Results for Other Stablecoin Prices

C.2.1 Estimated coefficient $\phi(.)$ of tvDAR and other stablecoin prices

Figure C.2: USD Coin (USDC), T=1028

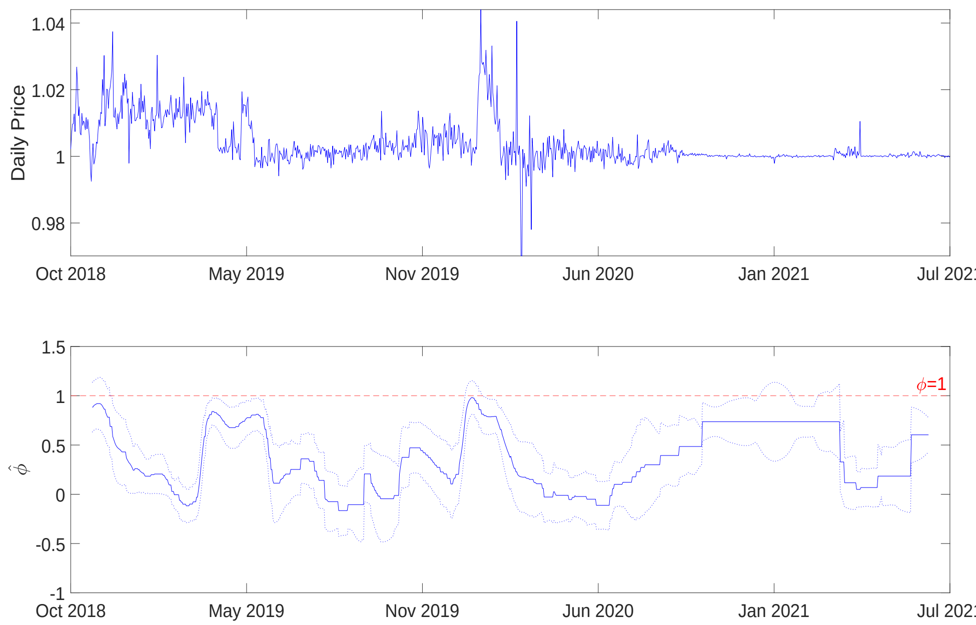

Note: The first panel shows the daily closing price series. The second panel presents the estimated autoregressive coefficient of tvDAR and its 95% confidence interval (Epanechnikov kernel, bandwidth =50/T).

Figure C.3: Binance USD (BUSD), T=618

Note: The first panel shows the daily closing price series. The second panel presents the estimated autoregressive coefficient of tvDAR and its 95% confidence interval (Epanechnikov kernel, bandwidth =50/T).

Figure C.4: DAI (DAI), T=618

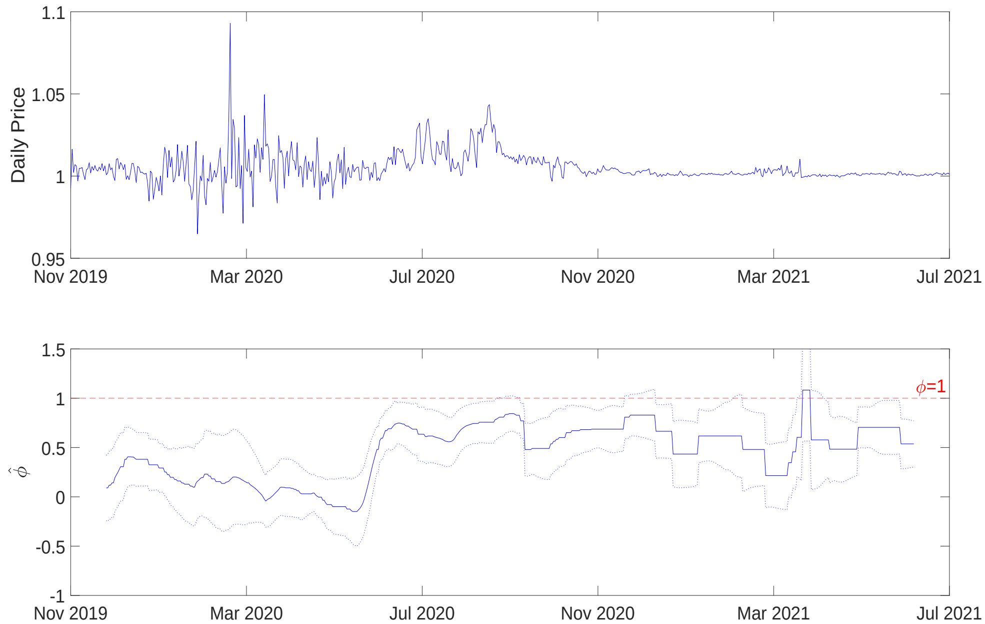

Note: The first panel shows the daily closing price series. The second panel presents the estimated autoregressive coefficient of tvDAR and its 95% confidence interval (Epanechnikov kernel, bandwidth =50/T).

Figure C.5: True USD (TUSD), T=1244

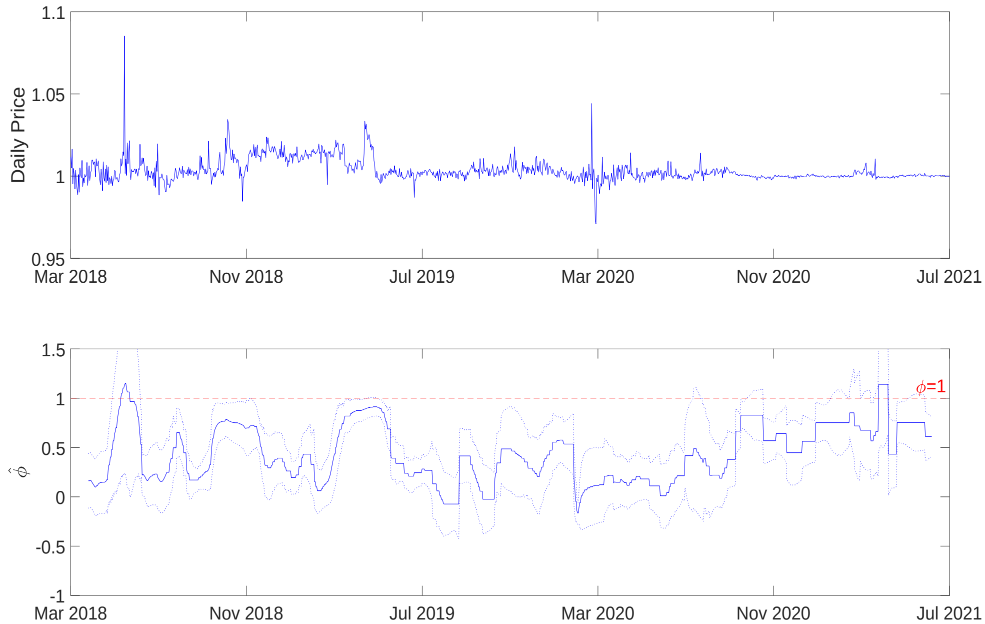

Note: The first panel shows the daily closing price series. The second panel presents the estimated autoregressive coefficient of tvDAR and its 95% confidence interval (Epanechnikov kernel, bandwidth =50/T).

Figure C.6: Terra USD (UST), T=249

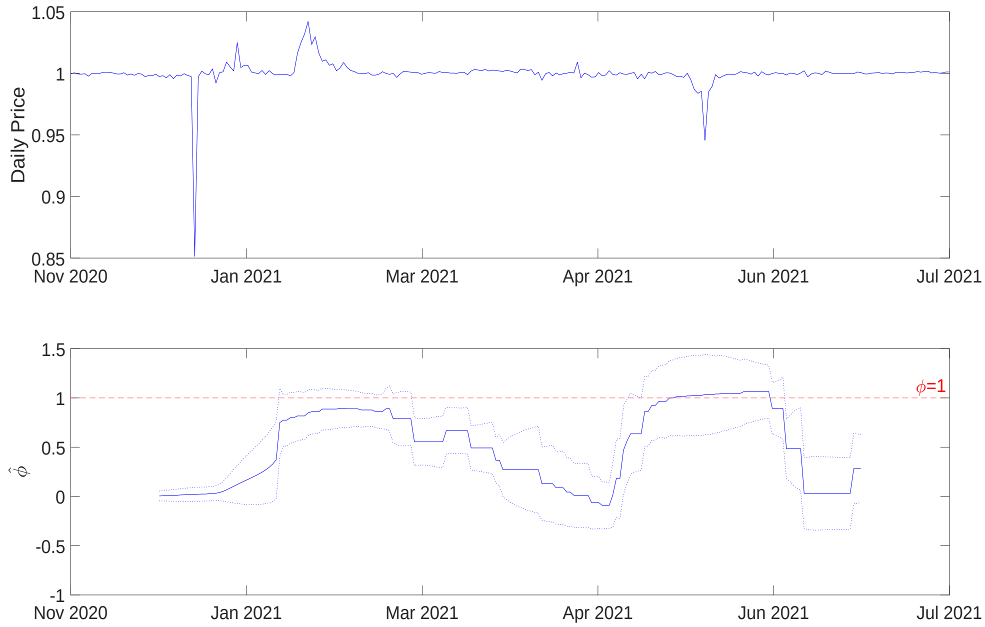

Note: The first panel shows the daily closing price series. The second panel presents the estimated autoregressive coefficient of tvDAR and its 95% confidence interval (Epanechnikov kernel, bandwidth =50/T).

C.2.2 Estimated Stability Measure $\hat{\xi}$ for Top Cap Stablecoins

We use the stability measure $\hat{\xi}$ to assess the stability of other processes of stablecoin prices. Based on the results of Proposition 4, we estimate ξ by rolling and construct a confidence interval for ξ. Value one (red line) inside the confidence interval is a warning of local instability.

Figure C.7: Stability measure $\hat{\xi}_t$ for USDC, T=1028

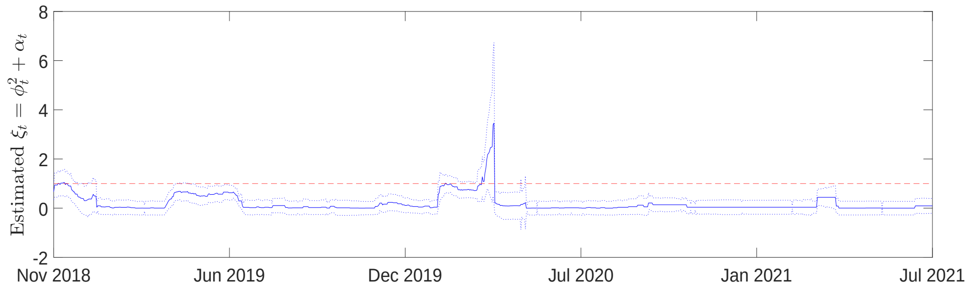

Note: This figure shows the estimates of stability measure $\hat{\xi}$ based on a 50-day rolling window and its 95% confidence interval.

Figure C.8: Stability measure $\hat{\xi}_t$ for BUSD, T=681

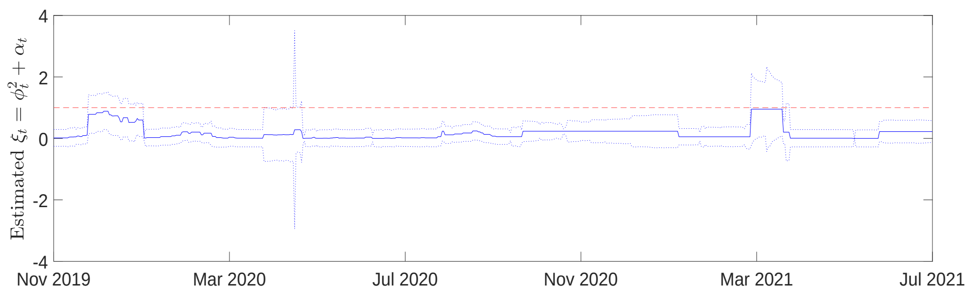

Note: This figure shows the estimates of stability measure $\hat{\xi}$ based on a 50-day rolling window and its 95% confidence interval.

Figure C.9: Stability measure $\hat{\xi}_t$ for DAI, T=618

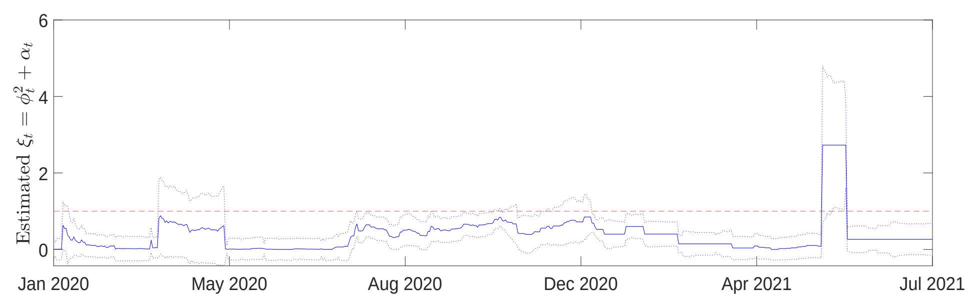

Note: This figure shows the estimates of stability measure $\hat{\xi}$ based on a 50-day rolling window and its 95% confidence interval.

Figure C.10: Stability measure $\hat{\xi}_t$ for TUSD, T=1244

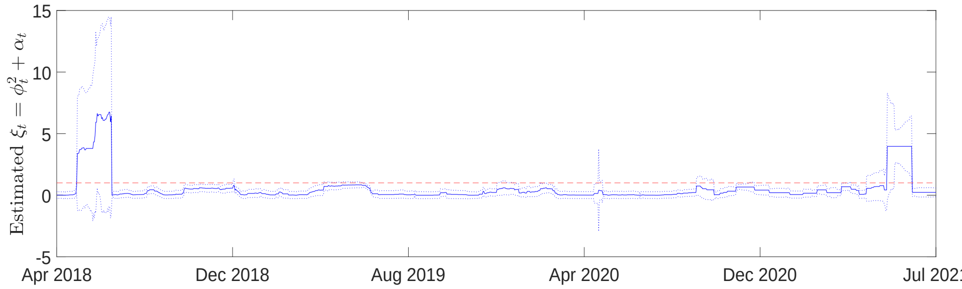

Note: This figure shows the estimates of stability measure $\hat{\xi}$ based on a 50-day rolling window and its 95% confidence interval.

Figure C.11: Stability measure $\hat{\xi}_t$ for Terra (UST), T=249

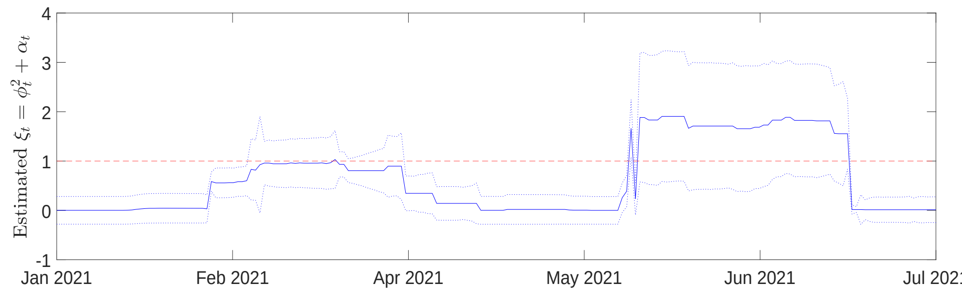

Note: This figure shows the estimates of stability measure $\hat{\xi}$ based on a 50-day rolling window and its 95% confidence interval.

Appendix D

Supplementary Analysis for Paper 3

D.1 In-sample Analysis

In-sample analysis is conducted using the first $T = 912$ observations to select the optimal DAR(p^*) process for each funding rate series and assess the goodness of fit. The optimal lag order p^* of the DAR process is determined via standard residual analysis, where serial dependence in the error terms of various DAR models of $p \geq 1$ is inspected by applying the Ljung-Box test on the standardized residuals $\hat{\eta}$ and squared residuals $\hat{\eta}^2$ from the candidate models. The results are summarized in Table D.1.

Table D.1: Ljung-Box Test on the standardized residuals from DAR(p) models

	DAR(p) Model					
	p=1	p=2	p=3	p=4	p=5	p=6
Binance						
$\hat{\eta}$	1	1	1	**0**	1	1
$\hat{\eta}^2$	0	0	0	**0**	0	0
Bybit						
$\hat{\eta}$	1	1	1	1	1	**0**
$\hat{\eta}^2$	0	0	0	0	0	**0**

Note: Rejection (1) or no rejection (0) of the null hypothesis of no autocorrelation. Optimal lag for the test is chosen as $min(T/5, 10)$ as per Hyndman and Athanasopoulos (2018) "Forecasting: Principles and Practice".

According to the Ljung-Box test results, p=1 would be enough for both funding rate series to remove any serial dependence in the squared residuals, but larger lags are required for the independence of the residual levels. To ensure that both the residuals and squared residuals are serially independent at the same time, the DAR(p) process should be fitted with p=4 for Binance and p=6 for Bybit. Parameter estimates of the selected models are provided in Table D.2.

Table D.2: Estimation of the optimal DAR(p^*) models for the period $t = 1, \ldots, 912$.

	Model Paramaters												
	ϕ_1	ϕ_2	ϕ_3	ϕ_4	ϕ_5	ϕ_6	ω	α_1	α_2	α_3	α_4	α_5	α_6
Binance													
Estimates	0.487	0.254	0.150	0.104			4.88E-06	0.101	0.064	0.037	0.022		
Standard error	(0.041)	(0.043)	(0.042)	(0.035)			(1.19E-06)	(0.071)	(0.068)	(0.059)	(0.044)		
Likelihood	4651.777												
Bybit													
Estimates	0.387	0.182	0.135	0.066	0.059	0.113	8.17E-06	0.216	0.062	0.030	0.022	0.028	0.016
Standard error	(0.046)	(0.044)	(0.043)	(0.042)	(0.041)	(0.035)	(1.43E-06)	(0.079)	(0.057)	(0.047)	(0.042)	(0.041)	(0.029)
Likelihood	4411.721												

Note: Model parameters are estimated using the observations up to and including November 09, 2024 at 16:00 (first 912 observations). Standard deviations are derived from the asymptotic distribution in Section 5.1.1 without the normality assumption and reported in parentheses.

To evaluate the goodness of fit of the selected models, the likelihood-ratio test is conducted against DAR(1) as the restricted model, and the results are reported in Table D.3. In addition, the fitted and actual funding rate series are plotted together in Figure D.1 for visual inspection.

Table D.3: Likelihood Ratio (LR) Test against DAR(1).

	dof	pvalue	test-statistics	critical value
Binance	6	0	211.062	12.592
Bybit	10	0	143.379	18.307

Note: Testing of the null hypothesis, restricted model in favor of the alternative unrestricted model.

Both selected models pass the LR test with the test statistics greater than the critical value for 1% significance level of the corresponding χ^2 distributions. This evidence further supports the argument that the higher-order DAR processes are a much better fit for the historical funding rates. From the graphical representations, it can be easily seen that the fitted values do a satisfactory job of capturing the deviations in the funding rates, even during periods of high volatility.

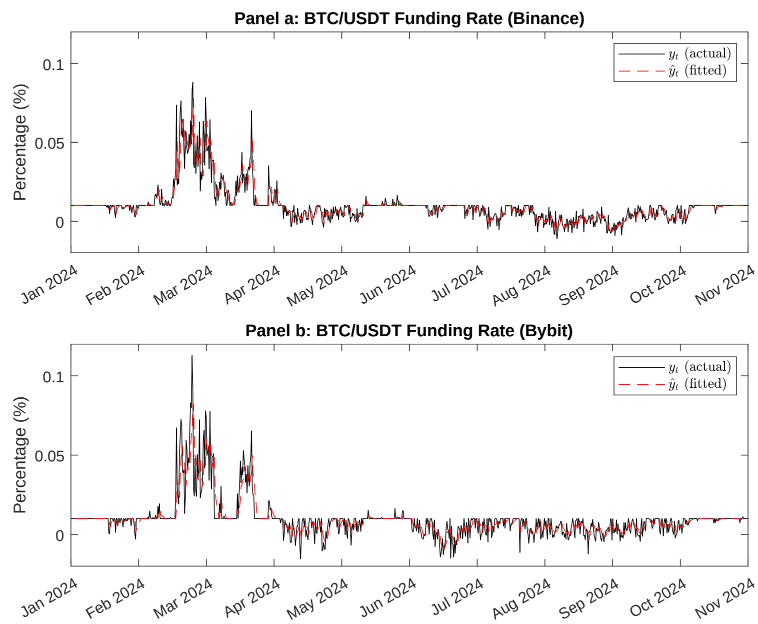

Figure D.1: Fitted and actual values of the Binance (a) and Bybit (b) funding rates.

D.2 Supporting Material

Figure D.2: Correlation between the BTC/USDT exchange rate and its funding rates.

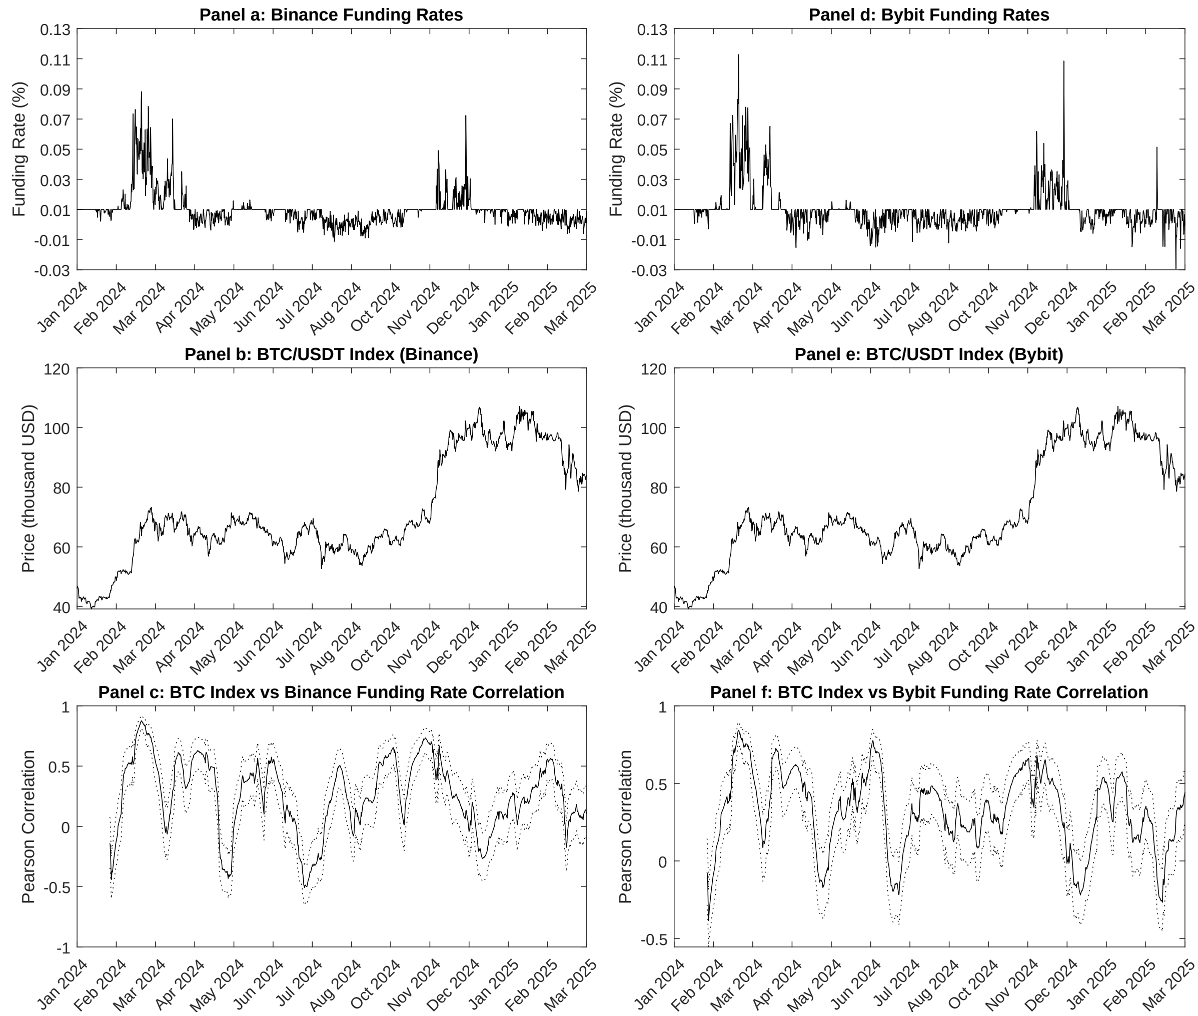

Note: This figure illustrates the historical funding rates of BTC/USDT perpetual contracts on Binance (a) and Bybit (d), the spot prices of BTC (b, e), and the local Pearson correlation coefficients along with 95% confidence intervals in panels c and f, respectively. The Pearson correlation coefficients are calculated at each point in time by rolling with a window of 84 observations, corresponding to 4 weeks of data.

Figure D.3: Quantile-Quantile (QQ) plots of Binance and Bybit funding rates.

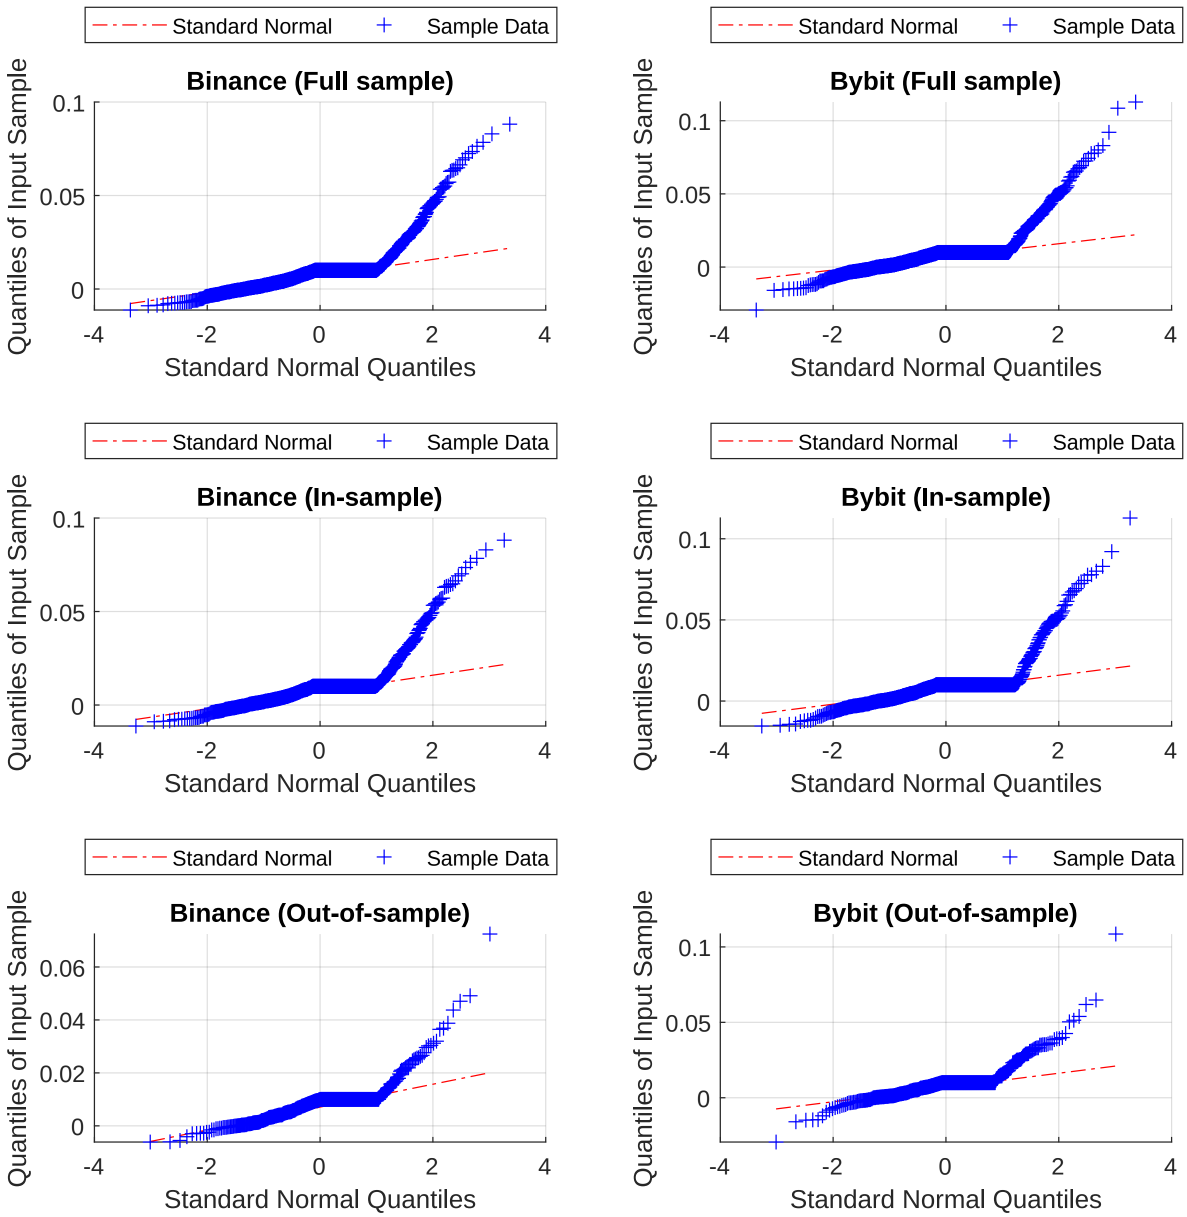

Note: This figure presents Quantile-Quantile (QQ) plots of Binance funding rates on the left-hand side and those of Bybit on the right-hand side. Full sample spans from January 11, 2024 to March 17, 2025 (1296 observations), in-sample period from January 11, 2024 to November 9, 2025 at 16:00 (912 observations), out-of-sample period from November 10, 2025 at 00:00 to March 17, 2025 at 16:00 (384 observations).

Figure D.4: Kernel densities of the historical funding rates.

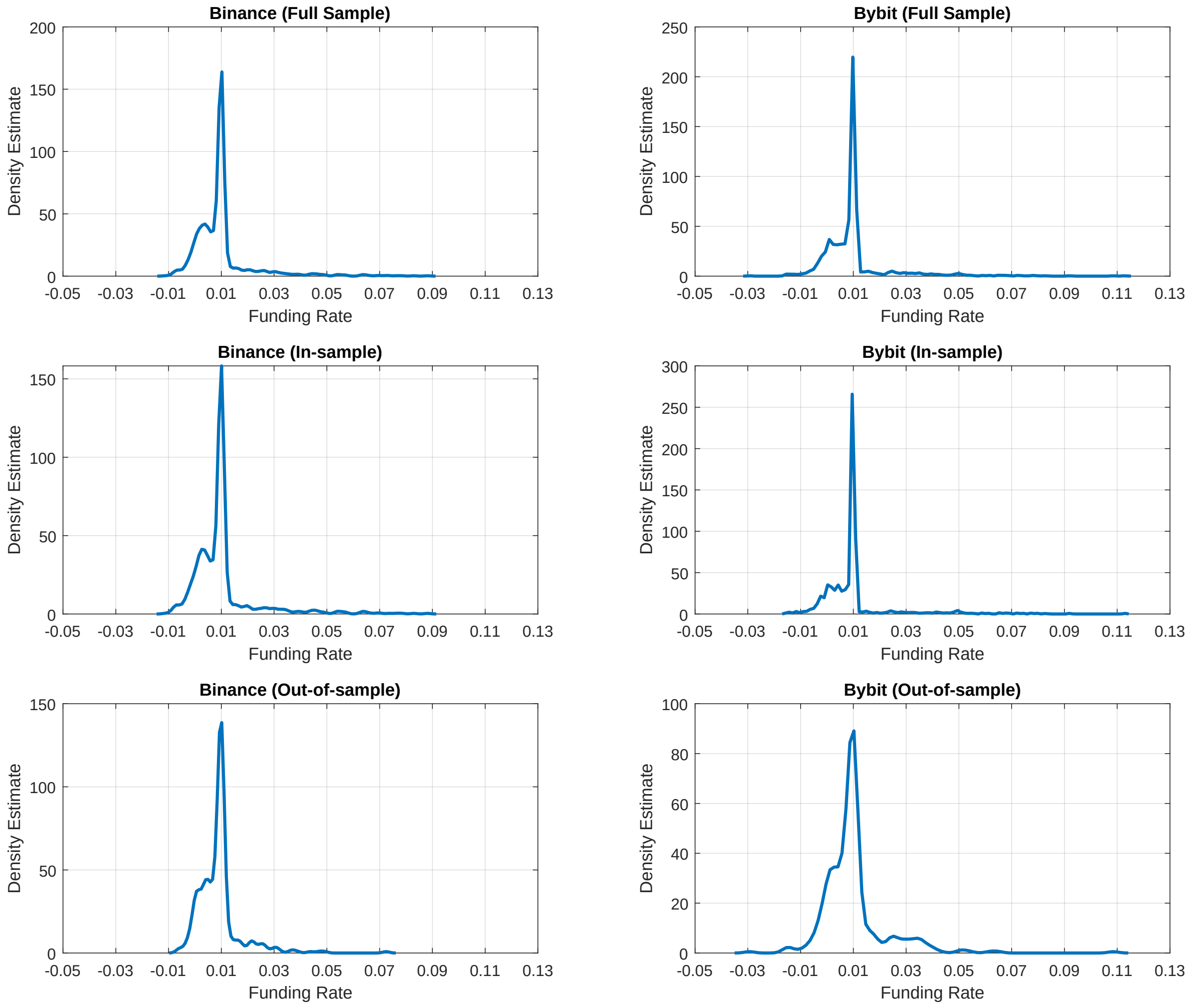

Note: This figure plots kernel densities of the funding rates, with Binance on the left-hand side and Bybit on the right. Full sample spans from January 11, 2024 to March 17, 2025 (1296 observations), in-sample period from January 11, 2024 to November 9, 2025 at 16:00 (912 observations), out-of-sample period from November 10, 2025 at 00:00 to March 17, 2025 at 16:00 (384 observations).

Figure D.5: Evolution of the model parameters and the conditional variances.

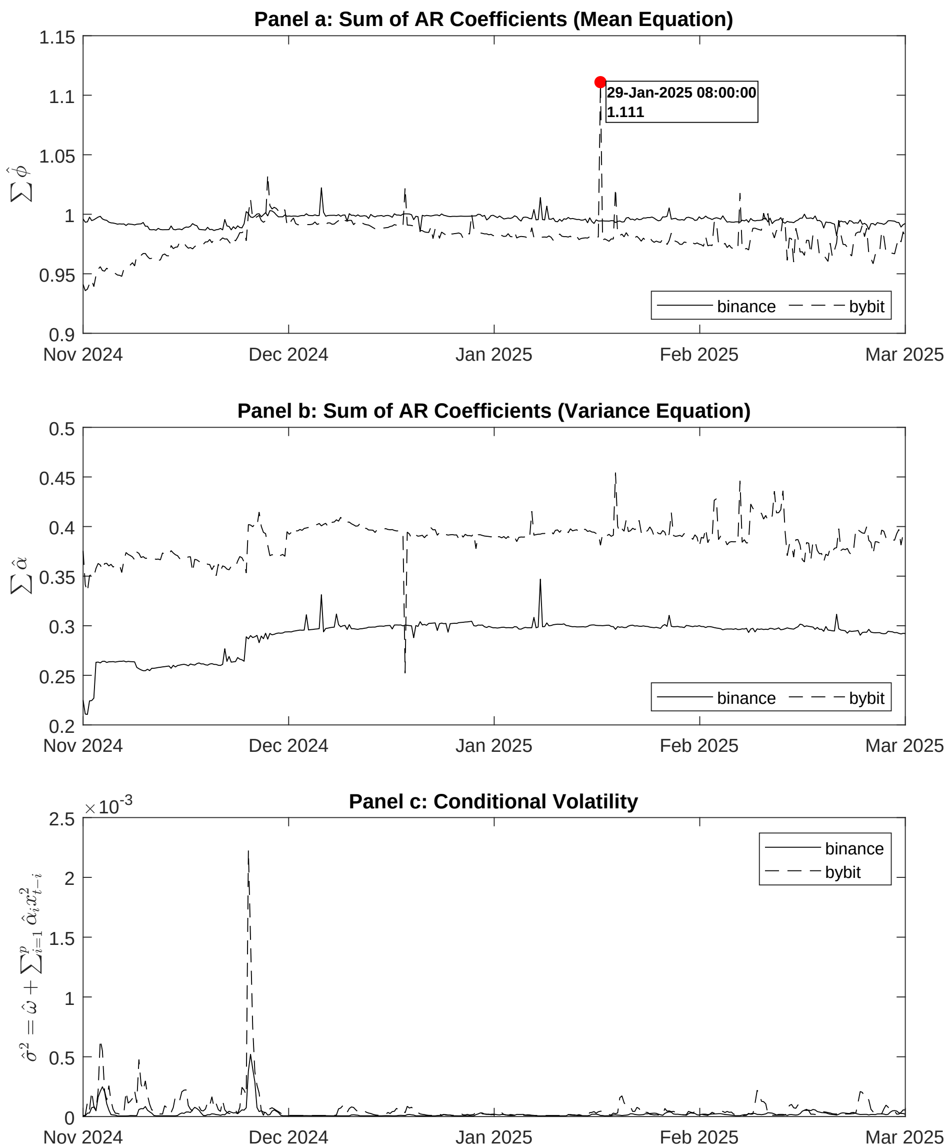

Note: This figure shows the evolution of the model parameters (a, b) and conditional variances (c) over the evaluation period, with Binance represented by a solid line and Bybit by a dashed line. For each exchange, these quantities are calculated based on the expanding window estimates.