



Hybrid ARDL-MIDAS-Transformer time-series regressions for multi-topic crypto market sentiment driven by price and technology factors

Ioannis Chalkiadakis^{1,4} · Gareth W. Peters² · Matthew Ames³

Received: 3 April 2022 / Accepted: 20 March 2023 / Published online: 20 July 2023
© The Author(s) 2023

Abstract

This paper develops a novel hybrid Autoregressive Distributed Lag Mixed Data Sampling (ARDL-MIDAS) model that integrates both deep neural network multi-head attention Transformer mechanisms, and a number of covariates, including sophisticated stochastic text time-series features, into a mixed-frequency time-series regression model with long memory structure. In doing so, we demonstrate how the resulting class of ARDL-MIDAS-Transformer models allows one to maintain the interpretability of the time-series models whilst exploiting the deep neural network attention architectures. The latter may be used for higher-order interaction analysis, or, as in our use case, for design of Instrumental Variables to reduce bias in the estimation of the infinite lag ARDL-MIDAS model. Our approach produces an accurate, interpretable forecasting framework that allows one to forecast end-of-day sentiment intra-daily, with readily attainable time-series regressors. In this regard, we conduct a statistical time-series analysis on mixed data frequencies to discover and study the relationships between sentiment from our custom stochastic text time-series sentiment framework, alternative popular sentiment extraction frameworks (BERT and VADER), and technology factors, as well as to investigate the role that price discovery has on retail cryptocurrency investors' sentiment (crypto sentiment). This is an interesting time-series modelling challenge as it involves working with time-series regression models in which the time-series response process, and the regression time-series covariates, are observed at different time scales. Specifically, a detailed real-data study is conducted where we explore the relationship between daily crypto market sentiment (of positive, negative and neutral polarity) and the intra-daily (hourly) price log-return dynamics of crypto markets. The sentiment indices constructed for a variety of “topics” and news sources are produced as a collection of time-series capturing the daily sentiment polarity signals for each “topic”, namely each particular market or crypto asset. Different sentiment methods are developed in a time-series context, and utilised in the proposed hybrid regression framework. Furthermore, technology factors are introduced to capture network effects, such as the hash rate which is an important aspect of the money

Extended author information available on the last page of the article

supply relating to the mining of new crypto assets, and block hashing for transaction verification. Throughout our real data study, we provide guidance and insights on how to use our hybrid model to combine—in a transparent, non-black-box way—covariates obtained with different time resolutions, how to understand the arising dynamics between these covariates, potentially under the presence of long memory structure, and, finally, successfully leverage these in forecasting applications. The hybrid model developed demonstrated superior performance to alternatives in both in-sample and forecasting application on real data.

Keywords Mixed-data sampling time-series regression (MIDAS) · Transformer deep neural network · Multi-scale resolution data · Natural language processing (NLP) · Text sentiment NLP time-series modelling · Gegenbauer long memory · Econometrics · Time-series

JEL Classification C32 · C36 · C45 · C49 · C51

1 Introduction

Cryptocurrency markets are emerging into mainstream finance with their adoption by institutions taking place on a regular basis, and regulatory frameworks coming into place to better manage the new classes of assets, markets and risks emerging from the re-envisaging of financial models through the Decentralised Finance (DeFi) movement. However, still the overwhelming majority of investors in the crypto space are comprised of retail investors rather than institutions. They are actors in the crypto markets for a variety of reasons; they may be seeking alternatives to the traditional fiat systems that could be overly oppressive in some countries, or may be seeking safe haven from hyper inflationary domestic currencies. Note that the latter may be due to the potential offering of a reliable, censorship-resistant scarce store of wealth, speculation of price activity, convenience of efficient financial transactions, or significantly greater interest bearing accounts from lending and staking platforms. However, there is no denying the fact that the crypto markets are still significantly influenced by sentiment as an entire industry on news and analytics has arisen in the crypto space that mirrors the mainstream financial media outlets such as MSNBC, Reuters and the like. In the online and print media, this includes news brands that have been servicing the crypto community for at least 5–6 years consistently at this stage, and includes well-known news brands in the crypto space such as Cointelegraph, cryptonews, CoinDesk, Bitcoin Magazine, Crypto Reddit, CryptoSlate, CryptoPotato, Coinmarketcap and Cryptoscoop, to name a few of the more widely followed news feeds.

Sentiment analysis, or opinion mining, is an active area of study in the field of natural language processing that analyses people's opinions, sentiments, evaluations, attitudes, and emotions via the computational treatment of subjectivity in text, for instance, see detailed overviews of the field of sentiment analysis in Liu (2012) and Pang and Lee (2008).

In this study, we aim to explore the relationship between crypto market sentiment and intra-daily price. In particular, we seek to study the time-series relationship between the daily sentiment time-series of two major crypto assets, i.e. Bitcoin and Ethereum, to the intra-daily time-series for the price of leading crypto assets and the volatility dynamics of crypto currency markets, on an hourly time resolution. We construct the sentiment time-series based on a collection of curated news articles about Bitcoin and Ethereum that span the last 3 years, which we have collected from widely-read crypto news sources. A range of crypto news sentiment perspectives is then captured by sentiment signals of positive, negative and neutral sentiment polarities for a variety of different cryptocurrency markets and news sources.

To achieve this, we introduce a new approach to cryptocurrency sentiment that is tailored to the crypto space context, and we compare and contrast our proposed methodology with existing sentiment extraction methods such as BERT (an attention-based Transformer sentiment model) and VADER (a rule-based model for online social media text sentiment analysis).

The studies we performed extract sentiment on particular cryptoassets, incorporating in the process the sentiment arising from market opinion, crypto regulation and Decentralised Finance news articles, into a common sentiment index. In addition, we incorporate technology factors related to network and mining efficiency, as well as transaction costs. We finally combine this information into two classes of econometric models. First, we relate sentiment extracted by our rigorous crypto-specific framework, to other less interpretable crypto sentiment methods, using an Autoregressive Distributed Lag (ARDL) modelling framework. Having shown the utility of our proposed sentiment time-series methodology, we next adopt a time-series regression framework that will accommodate the different time scales of the response time-series of daily crypto sentiment, and the covariate time-series of hourly crypto asset prices, volatility and network effects, such as the hash rate.

We will focus our analysis on currency pair markets for the asset exchange rates of BTC/USDT and ETH/USDT extracted from an aggregate price, and obtained from CoinGecko, a leading price aggregator from a variety of centralised cryptocurrency exchanges and distributed DEXs.

1.1 Statistical modelling and application contributions

To undertake the proposed cryptocurrency studies we have adopted and extended a class of time-series regression models known as the mixed data sampling (MIDAS) models which were recently made popular in the econometrics community. They allow one to parametrically accommodate Autoregressive Distributed Lag (ARDL) structures where the response time-series is sampled at a different frequency to the covariate time-series. In this work, we explore and extend the class of MIDAS models to accommodate a few additional key structures:

- First, we incorporate within the MIDAS-ARDL model an infinite-lag structure that is transformed to a finite lag MIDAS-ARDL model via use of the MIDAS-modified classical Koyck transform. We call this the Koyck-MIDAS transform.

We study the calibration of these models using Instrumental Variables (IV) for the sentiment signal, constructed based on VADER and deep learning solutions such as BERT. This in turn produces a hybrid time-series model that is denoted as the ARDL-MIDAS-Transformer model, which is an illustration of a class of ARDL-MIDAS-NeuralNet time-series regression models in which the neural network is a Transformer model that combines attention mechanisms with a Feed-Forward Neural-Network. This is used to construct the Instrumental Variables to reduce bias in the estimation of the infinite-lag Koyck-transformed ARDL-MIDAS model. We note that this is a significant contribution in terms of IV design and usage that remains on the interface between classical time-series IV regression and deep learning. There have been deep learning approaches that either jointly perform instruments construction and regression (Singh et al., 2020) or use solely deep learning models in a 2-stage IV regression setting (Hartford et al., 2017; Xu et al., 2020). Yet, such methodologies are distinct to our approach in that we use the expressive power of the complex black-box model while maintaining the transparent, time-series econometrics approach of the 2-stage IV regression: Stage 1 being a regression from the instrument to the treatment, and Stage 2 subsequently regressing the outcome on the treatment, conditioning on the instrument. In our setting, the challenge is not only to find an instrumental variable, but also the fact that the space of the IV is not a standard space, but rather it comes from a document set. Thus, one needs to learn a mapping from an abstract space of text data to a real-valued time-series, which is distinct to what other methods adopt to obtain an IV and is the reason why the Large Language Transformer-based Model was adopted, given it is efficient at learning this mapping. Furthermore, the proposed framework gives a more direct aspect of interpretability to the components of the model and their influence on the response. For instance, in the proposed solution we have a direct interpretation of short-term and long-term dynamic effects of instantaneous or persistent changes in the covariate and the influence this would have on the response (see also Sect. 2.3.1 for an in-depth discussion). This can be achieved through standard equilibrium and transfer function analysis of the class of time-series models developed, which is not achievable or interpretable with other black-box approaches due to their complicated structure. The ability to interpret the influence of the covariate on the response is critical to understanding the practical and statistical relationship between the regression variable and the response. In the context of this work, it gives a direct interpretation of the effect of price on market sentiment, both when there is an instantaneous change of price at time t and how that propagates over time, and when there is a persistent effect of price on sentiment.

- Second, we incorporate long memory structure into the MIDAS-ARDL model class, creating a form of MIDAS-GARDL model where the G stands for the class of long memory structures we incorporate, known as the Gegenbauer long memory polynomial filtertaps. The MIDAS class of models has been a breakthrough in that it allows joint modelling of covariates sampled at different time scales without having to aggregate the high-frequency covariates to match those at a lower-frequency thus losing information content. With our proposed exten-

sion with the long memory filter, we show, in addition, that stylised features of the covariates, like long memory, can be explicitly incorporated into the model as part of model development. Consequently, the MIDAS polynomial choice does not need to be treated during model selection, for example, provided that we know that the high-frequency regressors have certain structure. This structure can inform the choice of MIDAS polynomial in a transparent fashion.

- Third, we study the combination of MIDAS exponential Almon weight functions, Koyck-transformed geometric weight decay structures, and the Gegenbauer polynomial weight function generator in the case studies undertaken in the crypto space.

From an application perspective we make the following additional practical contributions:

- In terms of statistical modelling, we develop an approach to constructing time-series formulations of sentiment to quantify in a single index the market sentiment that is extracted from multiple collections of daily sets of news articles. We propose a novel way to construct the sentiment index, and provide a combining rule to obtain a single index, which is important to parsimoniously summarise the sentiment content from a large collection of different news and text data sources (editorially curated news articles, analysts' reports, social media, GitHub, Discord, Telegram, Twitter etc.), thus facilitating sentiment incorporation into a time-series framework. Details on the construction and text sentiment usage in the cryptospace are presented in Sect. 4.
- We demonstrate that our approach to constructing sentiment time-series is distinct from those that can be derived by popular deep learning, Transformer solutions such as BERT, but also rule-based approaches such as VADER. This is achieved using ARDL time-series regressions methods and formal statistical tests (Sect. 6.1). We remark that at the time of writing this paper, we were not aware of any statistical rigorous study that compares the information content of a classical lexicon-based time-series sentiment and a deep neural network-based sentiment signal. We hope that the fact that we determine meaningful and interpretable differences will encourage researchers also in other disciplines, to start investigating more when to use each method and to find ways to combine the best of both worlds, rather than resorting directly to the most complex but not necessarily effective solution without questioning.
- We analyse the relationships between financial intra-day price signals, technology and network factors related to money supply in cryptocurrencies, and the daily sentiment time-series signal. The goal of the analysis is to enhance understanding of the evolving dynamics between covariates and responses of different time scales, thus facilitating in-sample fitting and out-of-sample forecasting applications (Sect. 6.2). In the application under study in this manuscript, we are interested in forecasting end-of-day sentiment at any point intra-daily, using intra-daily signals from price and technology factors. This kind of setting naturally arises in an NLP sentiment context, where documents arrive for processing at a batch daily rate contrary to the intra-daily price signals obtained from

financial markets and network analytics; nevertheless, many applications would benefit from a forecast of end-of-day sentiment obtained intra-daily. Specifically, regarding forecasting applications, the model can be used to extract information on different forecasting horizons for relationships between processes at different time scales, which is not commonly available; standard time-series or neural network models assume common time scales between response and covariates. Furthermore, this formulation can be extended to other finance and digital finance applications or any other type of markets; for instance, agricultural commodities markets, where one could study the relationship between investor sentiment and price dynamics. Institutional (Commodity Futures Trading Commission, US Department of Agriculture) reports or public news around their release dates could be used to construct sentiment signals, commodity futures or spot prices could be utilised as covariates, and the dynamics of price on sentiment in commodities market could also be studied in a term structure setting. Additional examples include building sophisticated trading strategies; detecting potential market manipulation attempts considering unusual market movements that consistently lead to particular sentiment responses at the end of the forecasting window, e.g. at end of day; forecasting end of day sentiment of a particular news outlet considering, for instance, sentiment from articles published at an hourly rate, journalists' Tweets at a minute rate and articles from more prolific and less prolific journalists. Furthermore, financial time-series often exhibit long memory structure, thus being able to explicitly account for that feature, even at different levels of strength per high-frequency covariate, can significantly improve model estimation.

2 ARDL-MIDAS long-memory time-series regressions

In this section, we present the modelling framework that incorporates four working components: Koyck-transformed infinite-lag Autoregressive Distributed Lag time-series regressions; Mixed Data Sampling (MIDAS) multi-time resolution time-series regressions; natural language text component obtained both via cryptotailored text processing and Transformer deep neural network architectures; and Gegenbauer long memory structure to parametrically capture persistence. The logic of combining these models is to exploit the ability of deep learning architectures to learn higher-order feature structures that can act as inputs to interpretable regression time-series models. We will begin with an overview of the overall regression structure before exploring each component.

The context of this study naturally allows one to explore a range of both Autoregressive Distributed Lag models, as well as MIDAS regression models. There is a subtle difference between such classes of models as explained in detail in Dhrymes et al. (1970). One can not strictly classify a MIDAS model as an autoregressive model in the standard sense as they involve regressors with different sampling frequencies. This can be understood as a consequence of the fact that autoregressive structures implicitly assume that data are sampled at the same frequency in the past.

Instead, MIDAS regressions share some features with distributed lag models but also have unique features we will adopt for part of this study.

The following notation conventions are adopted to accommodate the various time scales considered in the MIDAS structures. The low frequency time scale is indexed by t for the regression response process $\{y_t, t \in \mathbb{Z}\}$ and the higher frequency time scale denoted by m for the regression covariate time-series processes $\{x_t^{(m)}, t \in \mathbb{Z}\}$ which is observed m -times faster than time scale t , such that for each low frequency period t one has m values of $x_{t-1+1/m_i}^{(m_i)}, x_{t-1+2/m_i}, \dots, x_t$. Here, m_i will denote the i -th high frequency time scale and there may be numerous high-frequency time scales used depending on the covariates utilised. Two lag operators are employed:

- a low-frequency lag operator, which is denoted by L and will be applied as: $LY_t = Y_{t-1}$; and
- a high frequency lag operator $L^{1/m}$ which will apply to time-series observed m -times faster than the t time scale, and which, when applied, produces $L^{1/m}X_t^{(m)} = X_{t-1/m}^{(m)}$.

From this, we will define the following characteristic polynomials for the autoregressive (AR) and distributed lag (DL) time-series components:

$$\begin{aligned} \Phi_p(L) &= 1 - \sum_{j=1}^p \phi_j L^j, \\ \mathbf{x}_t^{(m)} &= [x_{t-1+1/m}, x_{t-1+2/m}, \dots, x_t], \\ \beta_k(L^{1/m}) &= \sum_{j=1}^k \beta_j L^{j/m}, \quad \beta_j = [\beta_{j,0}, \dots, \beta_{j,m}]^T, \\ L^{j/m} \mathbf{x}_t^{(m)} &= [L^{j/m} x_{t-1+1/m}, L^{j/m} x_{t-1+2/m}, \dots, L^{j/m} x_t]^T. \end{aligned} \tag{1}$$

The standard multiple ARDL-MIDAS model would then be given by a regression structure

$$\Phi_p(L)Y_t = \sum_{j=1}^J \beta_k^{(j)}(L^{1/m_j})X_t^{(m_j)} + \epsilon_t, \tag{2}$$

with $\Phi_p(L)$ the standard AR characteristic polynomial expressed in lag operator L at time scale t , and $\beta(L^{1/m_j})$ is the j -th time-series covariate’s characteristic polynomial with MIDAS weight function expressed in lag operator L^{1/m_j} , namely at time scale m_j times faster than t . This MIDAS polynomial is applied to the covariate time-series observed at the time scale of m_j times faster than t . Note that in this notation we have a vector covariate at time t constructed from the m_j sub-time steps.

We wish to extend this model in three important ways:

1. Considering an infinite-lag structure at time scale t with $\tilde{\beta}_\infty(L)$ and developing an ARDL-MIDAS Koyck Transform.

2. Considering an infinite-lag structure at time scale m_j with $\tilde{\beta}_\infty(L, L^{1/m})$ using a fractional integration of the Gegenbauer form, to capture the potential for long memory structure in the regression relationship, where we add the fractional difference operator $(1 - 2uL^{1/m_j} + L^{2/m_j})^{-d}$.
3. Adding a generative embedding model for the construction of high-order covariate feature interactions that can be combined within the multiple ARDL-MIDAS-Gegenbauer time-series regression and includes:
 - Transformer regression structures based on deep neural network multi-head attention architectures;
 - Natural Language crypto-specific covariate feature generation, which may range from time-series on semantics and word distributions to Context-Free Grammar parsing higher-order features.

2.1 Infinite-lag autoregressive distributed lag (ARDL) regressions

In this section, we briefly recall the basic framework of the ARDL regression modelling structure that will be adopted in the studies performed. The concept of distributed lag models is widely studied in econometrics and time-series literature and dates back to early works in the former, e.g. see the following papers on the estimation of such models (Dhrymes et al., 1970; Hannan, 1965; Klein, 1958).

A stylised distributed lag model is a time-series model in which the effect of a covariate on an outcome variable occurs over time. By adding autoregressive lags to a simple distributed lag model, a new model is formed, called an autoregressive distributed lag model (ARDL). A general ARDL(p, k) model is defined as follows either in expanded form or in terms of characteristic AR and DL polynomials:

$$Y_t = \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=0}^k \beta_j X_{t-j} + \varepsilon_t$$

$$\Phi_p(L)Y_t = \mu + \beta_k(L)X_t + \varepsilon_t$$

where ε_t is a stationary white noise error term, $\Phi_p(L)$ and $\beta_k(L)$ are respectively order- p and order- k characteristic polynomials for the AR and DL components expressed with regard to the backshift operator L .

In many cases, it can be challenging to determine the appropriate choice of lag structure for k , so one may set up an infinite-lag model structure by setting $k = \infty$. This model is very similar to an ARMA model, except that the infinite-lag polynomial is applied to the explanatory variable rather than the error term as would be the case in an ARMA structure. As such, this class of models is termed an infinite ARDL model:

$$Y_t = \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=0}^{\infty} \beta_j X_{t-j} + \varepsilon_t. \tag{3}$$

It can be immediately recognised that it is impossible to estimate the coefficients of equation (3) since the number of unknown parameters is infinite, therefore, further assumptions are required to make the problem tractable.

One popular way to solve the infinite-lag distributed model estimation problem is to impose a parametrisation on the relationship of the infinite lag coefficients to map the problem back to a finite parameter space. One may use for instance a geometric distributed lag model. As indicated by the name, these models are based on the geometric distribution. One of the most popular geometric distributed lag models is the Koyck model (Koyck, 1954). Under Koyck’s approach one can assume that all the coefficients of equation (3) have the same sign and decline geometrically, with a specified rate of decay. Then, by taking advantage of the geometric series convergence, the Koyck model turns the infinite coefficients of equation (3) into an equation that includes a finite number of unknown parameters.

Based on the Koyck transformation method, in Eq. (3) we make the substitution of a geometric decaying coefficient relationship given by $\beta_j = \mu\gamma^j$ where $0 < \gamma < 1$. Using the convergence of geometric series, Eq. (3) can now be rewritten as follows:

$$Y_t = \mu' + \phi' Y_{t-1} + \phi'' Y_{t-p-1} + \sum_{i=2}^p \varphi_i Y_{t-i} + \mu X_t + \varepsilon'_t, \tag{4}$$

where $\mu' = (1 - \gamma)\mu$, $\phi' = \phi_1 + \gamma$, $\phi'' = -\gamma\phi_p$, $\varphi_i = \phi_i - \gamma\phi_{i-1}$, and $(\varepsilon_t - \gamma\varepsilon_{t-1}) = \varepsilon'_t \sim \mathcal{N}(0, \sigma^2)$. The operation of the Koyck model to decrease the number of parameters in the model is appreciable; however, whilst the problem of an infinite number of parameters is resolved by this assumed functional parametrisation, the reformulated model then produces a challenge for parameter estimation. In Eq. (4), the error term ε'_t and Y_{t-1} are not independent anymore. Hence, this renewed equation cannot be efficiently solved using conventional techniques that solve the regression models in an unbiased manner.

One common way to resolve this challenge is to introduce the concept of an instrumental variable. In our model, Y_{t-1} should be replaced by an instrumental variable which is independent of ε'_t , whilst capturing the basic dynamic structure of Y_{t-1} . If such an IV can be constructed then the new model can be solved using the conventional regression model estimation techniques, see Stock and Trebbi (2003). One must be careful with the introduction of an IV to resolve this challenge, as, whilst the estimation can be performed, the properties of the resulting estimators will depend on the performance of the constructed IV. In later sections, we will demonstrate how to use deep neural network Transformer-based methods to construct such instrumental variables.

2.2 Fractionally integrated mixed data sampling (MIDAS) models for long memory regressions

In the MIDAS regression model context, the objective is to accommodate such distributed lag structures but the time-series of observations and regressors acting as lagged predictors are no longer sampled or observed at the same temporal resolution.

It is assumed that the covariate variables of the ARDL-type model are now observed at higher frequency than the covariates in an AR-type regression.

In Ghysels et al. (2005) and Ghysels et al. (2006), such classes of problems were studied and the resulting modelling framework was named the MIDAS class of distributed lag regressions. The MIDAS structures have been extended broadly by others such as Ghysels et al. (2007), Andreou et al. (2013), and Andreou et al. (2011). This is a suitable model structure with established estimation techniques that is capable of treating mixed frequency ARDL-type regression models.

Let the variable Y_t represent the crypto sentiment constructed at a daily sampling frequency and let X_t represent an explanatory factor such as crypto asset prices, technology factors such as hash rate which are sampled m times faster than Y_t (as an example, when Y_t is daily, $X_t^{(m)}$ is sampled hourly in the 24 h, 7 day per week crypto markets giving $m = 24$). Suppose that Y_t is available once between $t - 1$ and t . Using the MIDAS model, we want to project Y_t onto a history of lagged observations of $X_{t-j/m}^m$. A simple MIDAS regression model is defined as below:

$$Y_t = \mu + \beta_1 B(L^{1/m}; \psi) X_t^{(m)} + \varepsilon_t^{(m)}, \tag{5}$$

where $B(L^{1/m}; \psi) = \sum_{k=0}^K B(k; \psi) L^{k/m}$, as before $L^{1/m}$ is a lag operator on fractional time scale such that $L^{1/m} X_t^{(m)} = X_{t-1/m}^{(m)}$, μ and β_1 are the unknown parameters of the model, and $\varepsilon_t^{(m)}$ is the error term. In Eq. (5), the lag coefficients in $B(k; \psi)$ are parametrised as a function of a low dimensional vector of the parameters ψ .

One of the concepts in introducing the MIDAS models is to take advantage of the lag polynomials. In Ghysels et al. (2005), the authors use lag polynomials to avoid the parameter proliferation problem, and to reduce the cost of estimation of a variety of finite basis models are considered—we will focus on the Exponential Almon family. By applying some modifications to the Almon lag, Ghysels et al. (2005) introduce the Exponential Almon lag defined as:

$$B(k; \psi) = \frac{\exp(\psi_1 k + \dots + \psi_Q k^Q)}{\sum_{k=1}^K \exp(\psi_1 k + \dots + \psi_Q k^Q)},$$

where K is the number of lags required in Eq. (5).

In this work, we are interested in working with the potential for a strong persistence in the ARDL-MIDAS regression structure, which we will demonstrate can be achieved through the introduction of a long memory component in the model. A stationary time-series process $Y \equiv \{Y_t\}_{t=1:T}$ is said to be a long memory stationary process if the following condition (Beran, 1994) holds in terms of the divergence of the autocorrelation function for Y_t and Y_{t+j} at lag j :

$$\lim_{n \rightarrow \infty} \sum_{j=-n}^n |\rho(j)| \rightarrow \infty, \tag{6}$$

where

$$\rho(j) = \frac{\text{Cov}(Y_t, Y_{t+j})}{\sqrt{\text{Var}(Y_t), \text{Var}(Y_t + j)}}. \tag{7}$$

We will parametrise processes with this property into the ARDL-MIDAS model through a fractional difference operator that will admit an infinite-lag Gegenbauer functional polynomial series generator that can be combined with the ARDL-MIDAS models previously presented. We introduce, for the first time we believe, this new class of lag basis functions that incorporate long memory features to the MIDAS structure to produce a family of fractional-MIDAS basis functions that can produce long memory regression effects in the distributed lag factors at the high-frequency time scale m .

Definition 1 (*Gegenbauer MIDAS basis*) Consider the fractional long-memory Gegenbauer weight functional form given by

$$B(L^{1/m}; \psi) = \sum_{j=0}^{\infty} \psi_j L^{j/m} = (1 - 2uL^{1/m} + L^{2/m})^{-d},$$

with ψ_j given by Gegenbauer polynomial functions

$$\psi_j = \sum_{q=0}^{\lfloor j/2 \rfloor} \frac{(-1)^q (2u)^{j-2q} \Gamma(d - q + j)}{q!(j - 2q)! \Gamma(d)}, \tag{8}$$

where $|u| < 1$, $d \in (0, 1/2)$ and $\lfloor j/2 \rfloor$ represents the integer part of $j/2$. Furthermore, the Gegenbauer polynomials satisfy the recursive calculation given by

$$\psi_j = 2u \left(\frac{d-1}{j} + 1 \right) \psi_{j-1} - \left(2 \frac{d-1}{j} + 1 \right) \psi_{j-2}, \tag{9}$$

where $\psi_0 = 1$, $\psi_1 = 2du$ and $\psi_2 = -d + 2d(1 + d)u^2$.

Remark 1 The following remarks characterise this class of Gegenbauer-MIDAS coefficient functions that we introduce:

- The class of fractional Gegenbauer-MIDAS basis functions allows one to control the strength of the long-memory in the process through selection of d and u .
- If $u = 1$ one produces an ACF function that is strictly positive and decays with a hyperbolic decay rate.
- For $|u| < 1$ the ACF will oscillate between positive and negative with period dictated by d and a hyperbolic envelope decay rate.
- As $d \uparrow 0.5$ the strength of long-memory increases, the slower the hyperbolic decay of the coefficients in the MIDAS weight function will decay and therefore the longer the past of $X_t^{(m)}$ will influence the current regression response.

We will work with the extended MIDAS framework often termed in the econometrics literature as the Multi-MIDAS regression structure. In this model, one may adopt multiple covariates with different time scales as follows, for the J -variate case:

$$Y_t = \mu + \sum_{i=1}^J \beta_{1,i} B(L^{1/m_i}; \psi_i) X_{t,i}^{(m_i)} + \sum_{i=1}^J \varepsilon_t^{(m_i)}, \tag{10}$$

where one has J covariates each sampled at time scales $\{m_1, m_2, \dots, m_J\}$ with associated driving white noise processes for each time scale. We will consider the real data case studies the situation where $J = 2$ and $m_1 = 24$ and $m_2 = 1$. Often we will assume that we subsume all the driving noise processes for the regression, in the case of i.i.d. Gaussian errors, into one driving noise process given by $\varepsilon_t := \sum_{i=1}^J \varepsilon_t^{(m_i)}$.

2.3 Koyck infinite-lag ARDL-MIDAS(p, ∞, K, m) regressions

An ARDL-MIDAS model can be expressed in numerous ways. In this work, we have built upon the approach of Ghysels et al. (2004) and we have extended it by incorporating the infinite-lag ARDL(p, ∞) Koyck transform model with the MIDAS structure. Using Eqs. (4) and (10), we produce a variation of the classical ARDL(p, ∞) time-series regression, with the normalised $\sum_k^K B(k; \psi) L^{k/m}$ MIDAS weight function, given in the ARDL-MIDAS(p, ∞, K, m) context as follows:

$$\begin{aligned} Y_t &= \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=0}^{\infty} \beta_j B(L^{1/m}; \psi) X_{t-j}^{(m)} + \varepsilon_t^{(m)} \\ &= \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=0}^{\infty} \sum_{k=0}^K \beta_j B(k; \psi) L^{k/m} X_{t-j}^{(m)} + \varepsilon_t^{(m)} \\ &= \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=0}^{\infty} \sum_{k=0}^K \beta_j B(k; \psi) X_{t-j-k/m}^{(m)} + \varepsilon_t^{(m)}, \end{aligned} \tag{11}$$

which can be written in compact form as

$$\Phi_p(L) Y_t = \mu + \tilde{\beta}_{\infty}(L, L^{1/m}) X_t^{(m)} + \varepsilon_t^{(m)}, \tag{12}$$

with the double polynomial given by

$$\tilde{\beta}_{\infty}(L, L^{1/m}) = \sum_{j=0}^{\infty} \sum_{k=0}^K \beta_j B(k; \psi) L^{j+k/m}. \tag{13}$$

We will now introduce for this class of models a variation of the classical Koyck transform that we will denote as the MIDAS-Koyck Transform which we will use to refactor this model into a parsimonious parametrisation as detailed in the following proposition.

Proposition 2.1 (MIDAS-Koyck Transform) *Consider the time-series model given by the ARDL-MIDAS(p, ∞, K, m) model specified as follows:*

$$\Phi_p(L)Y_t = \mu + \tilde{\beta}_\infty(L, L^{1/m})X_t^{(m)} + \epsilon_t^{(m)}. \tag{14}$$

Then the modified Koyck transform applied to this model uses the modified geometric decay characteristic polynomial at time scale t given as follows:

$$\tilde{\beta}_\infty(L, L^{1/m}) := \beta_0 B(L^{1/m}; \psi) \sum_{j=1}^\infty \gamma^j L^j = \frac{\beta_0 B(L^{1/m}; \psi)}{1 - \gamma L}, \tag{15}$$

for $\gamma \in (0, 1)$, which can transform this ARDL-MIDAS(p, ∞, K, m) into the simplified ARDL-MIDAS($p, 1, K, m$) given by

$$Y_t = \beta'_0 + \phi' Y_{t-1} + \phi'' Y_{t-1-p} + \sum_{i=2}^p \varphi_i Y_{t-i} + \beta_0 \sum_{k=0}^K B(k; \psi) X_{t-k/m}^{(m)} + \epsilon_t^{(m)'} \tag{16}$$

where $\mu' = (1 - \gamma)\mu$, $\phi' = \beta_1 + \gamma$, $\phi'' = -\gamma\beta_p$, $\varphi_i = \beta_i - \gamma\beta_{i-1}$, and $(\epsilon_t^{(m)} - \gamma\epsilon_{t-1}^{(m)}) = \epsilon_t^{(m)'} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$.

Proof The derivation of this MIDAS-Koyck transform is a basic extension of the standard Koyck transform approach with a MIDAS component applied. This proceeds to transform the ARDL-MIDAS(p, ∞, K, m) model as follows:

$$\begin{aligned} Y_t &= \beta_0 + \sum_{i=1}^p \beta_i Y_{t-i} + \beta_0 B(L^{1/m}; \psi) \sum_{j=0}^{+\infty} \gamma^j X_{t-j}^{(m)} + \epsilon_t^{(m)}, \\ Y_{t-1} &= \beta_0 + \sum_{i=1}^p \beta_i Y_{t-1-i} + \beta_0 B(L^{1/m}; \psi) \sum_{j=0}^{+\infty} \gamma^j X_{t-1-j}^{(m)} + \epsilon_{t-1}^{(m)}. \end{aligned} \tag{17}$$

If one then multiplies the second row with the geometric decay rate γ :

$$\gamma Y_{t-1} = \gamma \beta_0 + \gamma \sum_{i=1}^p \beta_i Y_{t-1-i} + \beta_0 B(L^{1/m}; \psi) \sum_{j=0}^{+\infty} \gamma^{j+1} X_{t-1-j}^{(m)} + \gamma \epsilon_{t-1}^{(m)}, \tag{18}$$

and subtracts the expressions in Eqs. (17) (for Y_t) and (18) one then obtains:

$$Y_t - \gamma Y_{t-1} = (1 - \gamma)\beta_0 + \sum_{i=1}^p \beta_i (Y_{t-i} - \gamma Y_{t-1-i}) + \beta_0 B(L^{1/m}; \psi) X_t^{(m)} + \epsilon_t^{(m)} - \gamma \epsilon_{t-1}^{(m)},$$

which results in

$$\begin{aligned} Y_t - \gamma Y_{t-1} &= (1 - \gamma)\beta_0 + \beta_1 Y_{t-1} - \beta_p \gamma Y_{t-1-p} + \sum_{i=2}^p (\beta_i - \beta_{i-1} \gamma) Y_{t-i} \\ &\quad + \beta_0 B(L^{1/m}; \psi) X_t^{(m)} + \epsilon_t^{(m)} - \gamma \epsilon_{t-1}^{(m)}, \end{aligned}$$

which then gives the desired result after some changes of variable. □

Remark 2 One can make the following remarks about this ARDL-MIDAS(p, ∞, K, m) model transformed via the MIDAS-Koyck Transform to an ARDL-MIDAS($p, 1, K, m$) model:

- This specification of the infinite-lag structure allows one to specifically accommodate a type of m -period seasonal structure for the regressors at time scale m which will be consistent with a period t seasonal pattern for the high-frequency covariates. This has an advantage over other approaches to constructing infinite-lag structures due to the fact that it does not require assumptions of knowledge of the slower time scale process y_t at times between $t - 1$ and t .
- One should consider the use of an Instrumental Variable to replace the term Y_{t-1} to attempt to break the correlation that would be present between this variable and the transformed regression error term $\epsilon_t^{(m)}$.

2.3.1 The evolution and interplay of dynamics between MIDAS covariates and the response variable

The significance of the proposed econometrics model based on the MIDAS formulation becomes apparent when one tries to investigate and interpret the dynamics between the covariates at different frequencies and the response variable. Specifically, one may deploy standard time-series analysis methods to quantify the impact of the covariates at different time scales on the response variable at time t ; the mechanism that will facilitate this analysis are the dynamic multipliers, which are defined as follows. Beginning with Eq. (12), the model can be equivalently written in the following form, provided that $\Phi_p(L)$ is invertible:

$$\begin{aligned}
 Y_t &= \Phi_p(L)^{-1} \mu + \Phi_p(L)^{-1} \tilde{\beta}_\infty(L, L^{1/m}) X_t^{(m)} + \Phi_p(L)^{-1} \epsilon_t^{(m)} \\
 &= \Phi_p(L)^{-1} \mu + D(L) X_t^{(m)} + \Phi_p(L)^{-1} \epsilon_t^{(m)} \\
 &= \Phi_p(L)^{-1} \mu + \sum_{s=0}^{\infty} \delta_s X_{t-s}^{(m)} + \Phi_p(L)^{-1} \epsilon_t^{(m)}
 \end{aligned}
 \tag{19}$$

where $\tilde{\beta}_\infty(L, L^{1/m}) = \Phi_p(L)D(L)$. Then the impact of different lags of the covariate X on the response Y is given by:

$$\begin{aligned}
 m_0 &= \frac{\partial Y_t}{\partial X_t} := \delta_0, \\
 m_1 &= \frac{\partial Y_t}{\partial X_{t-1}} := \delta_1, \\
 &\vdots \\
 m_s &= \frac{\partial Y_t}{\partial X_{t-s}} := \delta_s,
 \end{aligned}
 \tag{20}$$

where $\delta_s, s = 0, 1, \dots$ can be obtained by matching coefficients in the system of equations $\tilde{\beta}_\infty(L, L^{1/m}) = \Phi_p(L)D(L)$. Note here that $s = g(p, j, K)$, for some function $g(\cdot)$ determined by the previous system of equations, which means that the time scale of the lags is also considered in the measurement of the impact on Y .

The interpretation behind δ_s is that we can investigate not only how much a change in X affects Y but also when that effect occurs and whether it is instantaneous or it occurs over time. In particular, the short-run multiplier δ_0 expresses the immediate impact of a unit change in X_t at time t to the change of Y_t at time t , while the interim multipliers $\delta_s, s \geq 1$, show the response of Y_t to a unit change in X_{t-s} at time $t - s$. On the other hand, the long-run cumulative effect of X on Y measures how much Y will eventually change in response to a permanent change in X as $t \rightarrow \infty$. Assuming a long-run equilibrium condition, namely that changes in X do not cancel out, e.g. $X_{t-s} = X_{t-s+1} = \dots = X_t = X$ and $Y_{t-s} = Y_{t-s+1} = \dots = Y_t = Y$, then changes to $X_{t-s}, X_{t-s+1}, \dots, X_t$ lead to cumulative marginal effects on Y given by:

$$m_T = \sum_{s=0}^{\infty} \delta_s.
 \tag{21}$$

2.3.2 Gegenbauer-MIDAS Koyck transform

The ARDL-MIDAS structure is especially suitable for the application considered, where, if one lines up the time reference t with one of the leading markets morning period, for instance in Europe or Korea, then one would expect a periodic daily structure, when referenced to a US-based timezone. Based on this result, we can then construct the following specialised corollary models incorporating the Gegenbauer long memory coefficient functions within the infinite-lag ARDL-MIDAS model.

Corollary 2.1 (Gegenbauer-MIDAS Koyck transform) *Consider the time-series model given by the ARDL-Gegenbauer-MIDAS(p, ∞, ∞, m) specification as follows:*

$$\Phi_p(L)Y_t = \mu + \beta_\infty(L)(1 - 2uL^{1/m} + L^{2/m})^{-d}X_t^{(m)} + \varepsilon_t^{(m)}.
 \tag{22}$$

Then, the modified Koyck transform applied to this fractionally integrated MIDAS model produces

$$\begin{aligned}
 Y_t &= \beta'_0 + \phi' Y_{t-1} + \phi'' Y_{t-p-1} + \sum_{i=2}^p \varphi_i Y_{t-i} + \beta_0(1 - 2uL^{1/m} + L^{2/m})^{-d} X_t^{(m)} + \varepsilon_t^{(m)'} \\
 &= \beta'_0 + \phi' Y_{t-1} + \phi'' Y_{t-p-1} + \sum_{i=2}^p \varphi_i Y_{t-i} + \beta_0 \sum_{j=0}^{\infty} \sum_{q=0}^{\lfloor j/2 \rfloor} \frac{(-1)^q (2u)^{j-2q} \Gamma(d-q+j)}{q!(j-2q)! \Gamma(d)} X_{t-j/m}^{(m)} + \varepsilon_t^{(m)'},
 \end{aligned}$$

where $|u| < 1$, $d \in (0, 1/2)$, $\mu' = (1 - \gamma)\mu$, $\phi' = \beta_1 + \gamma$, $\phi'' = -\gamma\beta_p$, $\varphi_i = \beta_i - \gamma\beta_{i-1}$, and $(\varepsilon_t^{(m)} - \gamma\varepsilon_{t-1}^{(m)}) = \varepsilon_t^{(m)'} \sim \mathcal{N}(0, \sigma^2)$.

3 Hybrid ARDL-MIDAS-NeuralNet time-series regressions

To complete our framework, we now present a general hybrid time-series regression structure that allows one to incorporate into the ARDL-MIDAS time-series structure additional components obtained from a deep neural network architecture. We will consider a generic example at first using a feed-forward neural network and then we will specialise, for the NLP sentiment time-series context, to Transformer multi-head attention mechanisms.

3.1 Hybrid ARDL-MIDAS-FFNN time-series regressions

At this point, it suffices to capture the idea where we extend the ARDL-MIDAS regression to the hybrid Feed-Forward Neural Network (FFNN) version (ARDL-MIDAS-FFNN), with neural network depth n (number of computation layers) and additional covariate time-series denoted by $\{S_t^{(m)}\}$ that are observed at frequency m times faster than t ,

$$\Phi_p(L)Y_t = \sum_{j=1}^J \beta_k^{(j)}(L^{1/m_j})X_t^{(m_j)} + \left\langle \beta_q(L^{1/m}), (z^{(n)} \circ z^{(n-1)} \circ \dots \circ z^{(1)})(S_t^{(m)}) \right\rangle + \varepsilon_t, \tag{23}$$

where $z^{(l)}$, $1 \leq l \leq n$ denotes the l -th hidden network layer of dimension $q_m + 1 \in \mathbb{N}$ and $\beta_q(L^{1/m})$ is the MIDAS Distributed Lag (DL) operator at time resolution m times faster than t for a single neuron output layer. Note that this can trivially be generalised for each output layer neuron, but for clarity of notation it suffices to consider this simple case for the model specification.

This *readout* transformed covariate time-series is then combined with a MIDAS DL structure to produce an additional higher-order interaction component that enhances the linear non-interaction terms of the ARDL-MIDAS model through the *readout lag operator parameters* in polynomial $\beta_q(L^{1/m})$. This component can then be used either as additional trend structure, or, as we will do, as Instrumental Variables to reduce bias in the infinite lag ARDL-MIDAS setting.




Fig. 1 Example of a simple 3-layer Deep Neural Network architecture for a Feed-Forward Neural Net

In this architecture, for a given activation function in the FFNN denoted by $\psi : \mathbb{R} \rightarrow \mathbb{R}$, the l -th hidden network layer with activation function ψ is a map

$$z^{(l)} : \{1\} \times \mathbb{R}^{q_{l-1}} \rightarrow \{1\} \times \mathbb{R}^{q_l}, \quad z \rightarrow z^{(l)}(z) = \left(1, z_1^{(l)}(z), \dots, z_{q_l}^{(l)}(z)\right) \quad (24)$$

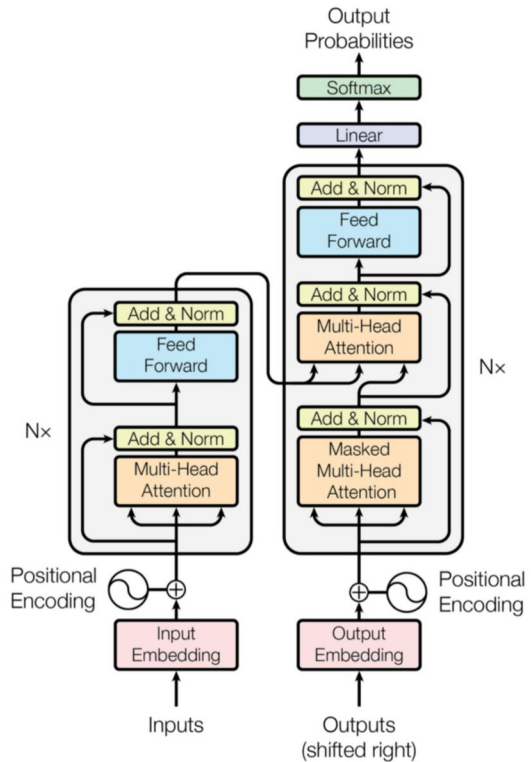
with *hidden neurons* $z_j^{(l)}, 1 \leq j \leq q_l$, being described by

$$z_j^{(l)}(z) = \psi(\alpha_j^{(l)}, z), \quad (25)$$

for given network parameters $\alpha_j^{(l)} \in \mathbb{R}^{q_{l-1}+1}$.

For instance, in the case of the FFNN we could set $n = 3, q_1 = 16, q_2 = 12$ and $q_3 = 10$ and we would have the output $(z^{(n)} \circ z^{(n-1)} \circ \dots \circ z^{(1)})(S_t^{(m)})$ that is projected as a time-series into the regression via the MIDAS lag operator and depicted in this case by the architecture given in Fig. 1.

Fig. 2 The Transformer model architecture that forms the fundamental building block of BERT (Devlin et al., 2019) for NLP tasks such as sentiment analysis (figure from “Attention Is All You Need” by Vaswani et al., 2017)



3.2 Hybrid ARDL-MIDAS-Transformer time-series regressions

In the context of Natural Language Processing and sentiment analysis it will be meaningful to also consider a class of deep neural network architectures known as “Transformers”. This is a specific neural network model which has proven to be especially effective for common natural language processing tasks such as sentiment analysis, see discussion in Vaswani et al. (2017). In particular, the Transformer model makes use of multiple attention mechanisms (Vaswani et al., 2017) which are effectively incorporated as encoders in sequence-to-sequence architectures (Sutskever et al., 2014), which aim to capture the sequential nature of text data. A thorough and detailed account of a Transformer is beyond the scope of this manuscript; suffice to say we will think of it as a more complex projection function of the input time-series generically denoted by $\{S_t^{(m)}\}$ to produce a transformed output time-series denoted by $\{T(S_t^{(m)})\}$, in which the transformation is comprised of significantly more components than the FFNN case. The Transformer-based model deployed in the study of the current manuscript and which has been the state-of-the-art in NLP, is BERT (Devlin et al., 2019), which is benefiting from multiple attention modules (“heads”), the basic format of one of which is illustrated diagrammatically in Fig. 2. A detailed description of BERT’s architecture is available in Tenney et al. (2019a,

b). Having this component will allow us to develop an ARDL-MIDAS-Transformer model as follows:

$$\Phi_p(L)Y_t = \sum_{j=1}^J \beta_k^{(j)}(L^{1/m_j})X_t^{(m_j)} + \langle \beta_q(L^{1/m}), T(S_t^{(m)}) \rangle + \epsilon_t. \quad (26)$$

To briefly explain the mapping inside the Transformer which we denoted generically by $T(\cdot)$, we may summarise it conceptually as follows. A Sequence-to-Sequence (“Seq2Seq”) architecture is a neural net system configuration, which comprises two components: an Encoder and a Decoder. The Encoder takes the input sequence and maps it into a higher dimensional feature space (n -dimensional vector). That abstract vector is then fed into the Decoder which turns it into an output sequence. The output sequence can be in another language, symbols, or a copy of the input. An initial, more intuitive, popular choice for this type of model is Long-Short-Term-Memory (LSTM)-based models. With sequence-dependent data, the LSTM modules can give meaning to the sequence while remembering (or forgetting) the parts they find important (or unimportant). Sentences, for example, are sequence-dependent since the order of the words is crucial for understanding the sentence, hence LSTM models are a natural choice for this type of data. Therefore, a very basic choice for the Encoder and the Decoder of a Sequence-to-Sequence model could be a single LSTM for each of the two components.

The Transformer is then integrating the attention mechanism into the Seq2Seq modelling, effectively looking at an input sequence and deciding at each step which other parts of the sequence are important. In contrast to LSTM-based encoders, for each input that the attention-based Encoder reads, the attention mechanism takes into account at the same time several other inputs that precede and follow the current input and decides which ones are important by attributing different weights to those inputs. The Decoder will then take as input the encoded sentence and the weights provided by the attention mechanism.

Hence, BERT (BERT: Bidirectional Encoder Representations from Transformers) is an architecture for transforming one sequence into another one via an Encoder and Decoder, but it differs from previous sequence-to-sequence models because it does not imply any Recurrent Networks (GRU, LSTM, etc.), but rather is solely based on Transformer modules. The BERT Transformer architecture provides significant improvements in Natural Language Tasks such as sentiment extraction. We use it in a manner to facilitate effective Instrumental Variable design to reduce bias in the estimation of our infinite-lag ARDL-MIDAS-Koyck-transformed Gegenbauer long memory time-series regression models.

3.3 Deep neural networks for instrumental variable design to reduce estimation bias

In this section, we will explain why we have chosen to extend our ARDL-MIDAS structure to an ARDL-MIDAS-NN or ARDL-MIDAS-Transformer type time-series regression. Consider the case of the infinite lag ARDL-MIDAS model parametrised

under a Koyck Transform as in Proposition 2.1, then we know that the transformed model given by

$$\Phi_p(L)Y_t = \mu + \tilde{\beta}_\infty(L, L^{1/m})X_t^{(m)} + \varepsilon_t^{(m)} \tag{27}$$

is modified via the proposed MIDAS-Koyck Transform to produce, for $\gamma \in (0, 1)$ a transform of the ARDL-MIDAS(p, ∞, K, m) model into the simplified ARDL-MIDAS($p, 1, K, m$) given by

$$Y_t = \beta'_0 + \phi'Y_{t-1} + \phi''Y_{t-1-p} + \sum_{i=2}^p \varphi_i Y_{t-i} + \beta_0 \sum_{k=0}^K B(k; \psi) X_{t-k/m}^{(m)} + \varepsilon_t^{(m)'} \tag{28}$$

In this model the classical Ordinary Least Squares OLS estimator will be biased by the fact that Y_{t-1} and $\varepsilon_t^{(m)'} = \varepsilon_t^{(m)} - \varepsilon_{t-1}^{(m)}$ are no longer independent, as a result of the Koyck Transformation. In this case it is standard practice to replace the regression variable Y_{t-1} with an Instrumental Variable which should capture similar information to Y_{t-1} but remain uncorrelated with white noise $\varepsilon_t^{(m)'}$. This is where we use the Neural Network structure to build such an Instrumental Variable and we obtain the regression model given by

$$Y_t = \beta'_0 + \phi'\tilde{Y}_{t-1} + \phi''\tilde{Y}_{t-1-p} + \sum_{i=2}^p \varphi_i \tilde{Y}_{t-i} + \beta_0 \sum_{k=0}^K B(k; \psi) X_{t-k/m}^{(m)} + \varepsilon_t^{(m)'}. \tag{29}$$

where we select the instrumental variable as

$$\tilde{Y}_t = \langle \beta_q(L^{1/m}), (z^{(n)} \circ z^{(n-1)} \circ \dots \circ z^{(1)})(S_{t-1}^{(m)}) \rangle. \tag{30}$$

We will illustrate this in more detail below in the context of Natural Language processing sentiment time-series, where we will consider state-of-the-art Transformer models for the IV construction.

4 Natural language processing for sentiment time-series structures

In this section, we discuss the specifics of the application studied in this paper, where the response time-series, denoted by $\{Y_t\}$, for our ARDL-MIDAS-Transformer model is a novel construction of an entropy-based sentiment time-series from a corpus of news articles specific to particular financial assets, in this case from the cryptocurrency market. We will also discuss how to construct the instrumental variable time-series $\{\tilde{Y}_t\}$ to replace $\{Y_t\}$ in the infinite-lag ARDL-MIDAS-Transformer model as discussed in Sect. 3.3. This will be obtained from what is known in machine learning as generative embedding feature learning and we will use specific mechanisms for this in the NLP sentiment context: Transformer models as presented in the previous section [BERT (Devlin et al., 2019)] and semantic- and grammar-based rules for sentiment extraction such as VADER [Valence Aware Dictionary for sEntiment Reasoning (Hutto et al., 2014)].

Pre-trained Transformer models such as BERT have achieved state-of-the-art performance on natural language processing tasks and have been adopted as feature extractors for solving downstream tasks such as question answering, natural language inference, and sentiment analysis. The current state-of-the-art Transformer based pre-trained models consist of dozens of layers and millions of parameters. While deeper and wider models yield better performance, they also need large GPU/TPU memory modules and a significant amount of text corpus for fine-tuning. For example, BERT-large is trained with 335 million parameters, and requires at least 16 GB of GPU memory to fine-tune (24 GB is the minimum recommended). The larger size of these models limits their applicability in time- and memory-constrained environments. Furthermore, there is an entire discipline now emerging to attempt to simplify and understand the layers of such complex architectures, see examples such as Jiao et al. (2020).

In the context of sentiment extraction, alternative methods have been developed based on tailored Sentiment Lexicons or Context-Free Grammar (CFG) parsing; a popular example explored in this work is the VADER approach and its associated micro-blogging sentiment lexicons. A substantial number of sentiment analysis approaches rely greatly on an underlying sentiment (or opinion) lexicon. A sentiment lexicon is a list of lexical features (e.g. words) which are generally labelled according to their semantic orientation as either positive or negative.

Manually creating and validating such lists of opinion-bearing features, while being among the most robust methods for generating reliable sentiment lexicons, is also one of the most time-consuming. For this reason, much of the applied research leveraging sentiment analysis relies heavily on preexisting manually constructed lexicons (Loughran & McDonald, 2011; Pennebaker et al., 2007; Zhang & Liu, 2017) in which words are categorized into binary classes (i.e., either positive or negative) according to their context-free semantic orientation. We have developed a tailored lexicon for cryptocurrency markets taking into account the jargon and colloquialisms that arise in cryptocurrency text that is unique to this domain. We will construct the target regression time-series $\{Y_t\}$ using our novel sentiment extraction framework, described below. Then, we will utilise both BERT and VADER to construct the Instrumental Variable time-series $\{\tilde{Y}_t\}$.

Therefore, in this section we briefly introduce how to transform processed text tokens into a time-series of distributions, in the process explaining what is known in the natural language processing context as the text embedding representation. Note that we are specifically interested in producing text embeddings with the aim to incorporate them in time-series regression models. Few approaches to constructing sentiment indices readily admit a time-series construction. A recent example is Hassani et al. (2020), who construct a sentiment scoring rule based on the difference between the number of positive and negative words in Tweets, which is an approach significantly different from ours; for an in-depth description of our framework we refer the interested reader to Chalkiadakis et al. (2021) where we employ the sentiment time-series construction to COVID-19 sentiment analysis, and Chalkiadakis et al. (2020) where we study statistical causality in crypto markets.

4.1 Distributional sequential text data embedding for response time-series $\{Y_t\}$

The embedding framework we construct is based on the well-known *bag-of-words model* (BoW), which is commonly applied in natural language processing (NLP) and information retrieval (Harris, 1954). The idea behind BoW in NLP is to represent a segment of text as a collection (“bag”) of words without considering the order in which they appear in the text. Instead, we are now setting BoW into a time-series context, and present a novel online formulation that allows us to incorporate the text-based sentiment index into a time-series system. Furthermore, in this way we avoid the computational limitations of BoW, stemming from having to manipulate large sparse matrices whose size depends on the number of distinct document words and corpus size, and may well be in the order of hundreds of thousands.

We begin by introducing some basic notation: t denotes a “token”, i.e. a linguistic unit of one or more characters (a word, a number, a punctuation character etc), \mathcal{V} is the *vocabulary*, namely a finite set of tokens that is valid in the language, and \mathcal{D} is a *dictionary* ($\mathcal{D} \subseteq \mathcal{V}$), i.e. a finite set of tokens, which we consider adequate to express the topic under study. We will work with n -grams, where n denotes the number of tokens in the text processing unit we consider, namely a set of n consecutive terms.

The time-series embedding is defined by the 3-ary relation $\mathcal{R} \subseteq \mathcal{V} \times \mathbb{D} \times \tilde{\mathcal{N}}$, where $\mathbb{D} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^p\}$, $\mathcal{D}^j \subseteq \mathcal{V}$ is a set of dictionaries each of size q_j , and $\tilde{\mathcal{N}} = \{\mathbb{N}^{q_1}, \mathbb{N}^{q_2}, \dots, \mathbb{N}^{q_p}\}$. To compute the members of $\tilde{\mathcal{N}}$ for each element of \mathcal{R} we use the following equation, which defines \mathcal{R} :

$$\hat{\gamma}_N^{j,l}(\tilde{v}_N, \mathcal{D}^{j,l}) = \begin{cases} \frac{m_N^{j,l}}{n \times N}, & r_m(\tilde{v}_N, \mathcal{D}^{j,l}) = 1 \\ 0, & \text{otherwise} \end{cases} \tag{31}$$

where $\tilde{v}_N = \{\tilde{v}_{wN}\}_{w=1:n}$, $m_N^{j,l} = |\{v' : v' \in \tilde{v}_N\} \cap \{\mathcal{D}^{j,l}\}|$, and $\mathcal{D}^{j,l}$ denotes a dictionary token $l \in \{1, \dots, q_j\}$, for dictionary $j \in \{1, \dots, p\}$, and

$$r_m(\tilde{v}_N, \mathcal{D}^{j,l}) = \begin{cases} 1, & m_N^{j,l} \geq m_{\min} \\ 0, & \text{otherwise,} \end{cases} \tag{32}$$

where N is the index of the current timestep, in n -gram “time” which indexes n -grams in our setting. Therefore, at each N we have a vector of dimension q_j which is the embedding of the n -gram at N . In this construction, the condition in Eq. (32) restricts the count of any token of \mathcal{D}^j which is in n -gram v_{1N}, \dots, v_{nN} at timestep N to be at least m_{\min} .

To capture the time-dependent nature of text, we note that the total number of observed tokens increases as we shift the n -gram towards the end of the text. Therefore, we want to recursively extract proportions of the dictionary tokens within the n -gram at time N . To account for this effect we apply the following transformation at each N :

$$\tilde{\gamma}_N^{j,l}(\cdot) = \begin{cases} \frac{\sum_{i=1}^{N-1} m_i^{j,l} + m_N^{j,l}}{M_N}, & r_m(\cdot) = 1 \\ 0, & \text{otherwise,} \end{cases} \tag{33}$$

where $m_N^{j,l}$ is the count of token l in dictionary \mathcal{D}^j at timestep N , and M_N is the total count of tokens we have observed up to timestep N which satisfy $r_m(\cdot) = 1$.

It is important to point out at this stage that the support of the distribution of proportions is restricted by the condition in Eq. (32). Tokens with count less than m_{\min} will be excluded from M_N , and consequently the support of the distribution. To construct the time-series for the current study, we set $n = 20$ and $m_{\min} = 1$.

4.2 Converting sequential text embedding to sentiment index time-series $\{Y_t\}$

The final stage of the construction comprises mapping this time-series of distributions onto a scalar summary to create a sequence of summary statistics that will define the sentiment index time-series.

Using the embedding extracted from token occurrences, we construct additional time-series using properties of the empirical distribution of the embedded text. We acquire the density of the token proportions of Eq. (33):

$$g_N^{j,l}(\tilde{v}_N, \mathcal{D}^{j,l}) = \frac{\mathbb{I}^{j,l}(\tilde{v}_N) \tilde{\gamma}_N^{j,l}(\tilde{v}_N, \mathcal{D}^{j,l})}{\sum_{l=1}^{q_j} \mathbb{I}^{j,l}(\tilde{v}_N) \tilde{\gamma}_N^{j,l}(\tilde{v}_N, \mathcal{D}^{j,l})} \tag{34}$$

where, as before, \tilde{v}_N denotes the n -gram at time-step N , and the indicator function $\mathbb{I}^{j,l}(\tilde{v}_N)$ selects the n -gram terms:

$$\mathbb{I}^{j,l}(\tilde{v}_N) = \mathbb{I}(\tilde{v}_N, \mathcal{D}^{j,l}) = \begin{cases} 1, & \text{if } l \in \{l' : \mathcal{D}^{j,l'} \in \tilde{v}_N \text{ for some } w\} \\ 0, & \text{otherwise} \end{cases}, \tag{35}$$

and then we can effectively study the density itself, that changes per n -gram, or use a suitable summary of it.

We expect that the frequency with which words are used in the course of the text, as well as the richness of the dictionary, will reflect on the value of the entropy of the empirical distribution of proportions, which we use to construct our time-series. The entropy is a vector-valued process of dimension p , $\mathbf{H}_N = [H_N^{(1)}, \dots, H_N^{(p)}]$, whose marginal component that corresponds to the j th dictionary is given, for $j = 1, \dots, p$, by:

$$H_N^{(j)}(\tilde{v}_N) \left\{ g_N^{j,l}(\cdot) \right\}_{l=1 : q_j} = \begin{cases} - \sum_{l=1}^{q_j} \mathbb{I}^{j,l}(\tilde{v}_N) g_N^{j,l} \ln(g_N^{j,l}(\cdot)), & \exists l \text{ s.t. } g_N^{j,l}(\tilde{v}_N, \mathcal{D}^{j,l}) \neq 0 \\ 0, & \text{otherwise.} \end{cases} \tag{36}$$

Using this framework, we construct the daily sentiment index per news source for positive and negative polarities. We then provide the robust median of the sentiment per day in each polarity to produce a robust, polarity-based collection of daily

sentiment indices in which to study relationships between retail sentiment and price dynamics.

Formally, our daily Entropy Sentiment index is constructed as follows (Chalkiadakis et al., 2020):

$$\text{entropy_index}(\tau) = \text{median}(H_{i,\tau}, \dots, H_{i+k,\tau}), \quad (37)$$

where $H_{i,\tau}, \dots, H_{i+k,\tau}$ are the entropy values of the text segments $i, \dots, i+k$, coming from articles written on the same calendar day τ .

4.2.1 Reference dictionary

The previous description of our framework for the construction of a lexicon-based sentiment index makes evident the requirement of an expressive dictionary (lexicon) of English words that will be purpose-built for the crypto space, as well as a collection of crypto-specific words annotated with sentiment information for that space—positive, negative or neutral words. The lexicon will be the basis upon which all text tokens are related.

To construct the dictionary, often people collect the most frequent tokens present in the corpus of documents that is available for training and evaluation of their model. However, we argue that this approach significantly restricts the representational power of the dictionary. In contrast, we treated the construction of the dictionary as a separate task. We collected a general English dictionary, as well as a number of dictionaries covering different topics, including Engineering and Technology, Media, Business, Economics, Finance, Mathematics, and Computing, all of which are pertinent to the crypto space. The dictionaries were constructed by collecting words present in online dictionaries, mainly those of Oxford University. After obtaining the word lists via web scraping, we further curated them by cleaning the tokens from scraping artefacts. Finally, together with experts of the crypto community, we manually compiled a list of words that express positive, negative or neutral sentiment when used in the context of cryptocurrency markets.

4.3 Combining multiple news source sentiment time-series: volume-based weighting for crypto market sentiment

If we consider the different crypto assets (Bitcoin-BTC, Ethereum-ETH) that form the focus of this study, then the news articles written about each of these assets can be considered as “topics” in an NLP text processing context. We can then consider different options for combining the sentiment time-series from these different topics across different news sources.

Let $X_{\tau}^{(s,j,q)}$ denote the sentiment indices where the index s refers to sentiment polarity $s \in \{\text{positive, negative, absolute magnitude}\}$, the index j refers to asset $j \in \{\text{BTC, ETH}\}$, $q \in \{\text{Cryptodaily, Cryptoslate}\}$ refers to the news source of the

articles, τ is an n -gram “time” index, and $N_{s,j}$ denotes the total number of n -grams (or could also be sentences) of “topic” j , with sentiment s in all news sources. For calendar time units $t = 1, \dots, T$ we can partition $\left\{ X_{\tau}^{(s,j,q)} \right\}_{\tau=1}^{N_{s,j}}$ by grouping the observations that come from articles published on the same day in each news source: $\left\{ X_{\tau}^{(s,j,q)} \right\}_{\tau=1}^{n_t^{s,j,q}}$, for $t = 1, \dots, T$, $\sum_{t=1}^T \sum_q n_t^{s,j,q} = N_{s,j}$ and $n_t^{s,j,q} \geq 0$.

To capture a market wide sentiment for a given polarity $s \in \{\text{positive, negative, absolute magnitude}\}$ and asset j we use a volume-based weighting rule:

$$X_t^{(s,j)} = \sum_q w_t^{s,j,q} \tilde{X}_t^{(s,j,q)}, \quad w_t^{s,j,q} = \frac{n_t^{s,j,q}}{\sum_q n_t^{s,j,q}}, \tag{38}$$

where $\tilde{X}_t^{(s,j,q)} = m\left(\left\{ X_{\tau}^{(s,j,q)} \right\}_{\tau=1}^{n_t^{s,j,q}}\right)$, where $m(\cdot)$ denotes the mapping from each segment (set of n -grams) of topic j , news source q , and sentiment s that corresponds to time t . The weights are assigned according to the volume of n -grams per day for each topic, which ensures that article lengths have no effect on the weight.


In Fig. 3 we plot the smoothed volume-weighted positive and negative sentiment indices as well as the index of absolute sentiment strength, for articles referring to Bitcoin. The extracted crypto sentiments for Ethereum are available in the Supplementary Appendix, Section A.

5 Estimation of ARDL-MIDAS-Transformer long-memory regressions


In this section, we explain a simple five-stage estimation procedure to fit the infinite-lag Koyck ARDL-MIDAS-Transformer Gegenbauer Long Memory model. The procedure can be used to also fit intermediate models such as Multiple-MIDAS, the ARDL-MIDAS or other variants, such as the ARDL-MIDAS-NeuralNet models for different architectures; as we discuss this explicitly in the sentiment NLP context, we focus on Transformers for our study. The four stages proceed as follows:

- Stage 1: Crawl and scrape cryptocurrency news articles to create a corpus of crypto news for topics of relevance from each news feed identified. Munge the text articles according to a range of chosen pre-processing steps to address aspects of data cleaning and denoising, in terms of punctuation, numbering, letter casing, stemming, stopword removal, word compounds, and removal of low-frequency words. This stage produces an intra-daily time-series of n -grams (n tokens), time-stamped and ordered, which we denote by $\{S_t^{(m)}\}$. In addition, construct a time-series of article sentences, which will be the input to the BERT and VADER models.
- Stage 2: Construct the Entropy Sentiment time-series for each news source and combine them as described in Sect. 4.3 to make the response time-series $\{Y_t\}$. Then, construct the instrumental variable generative embeddings $\{\tilde{Y}_t\}$,

(a) Positive sentiment



(b) Negative sentiment



(c) Absolute sentiment




Fig. 3 Sentiment indices with 95% confidence intervals constructed from articles about Bitcoin published on Cryptodaily (<http://cryptodaily.co.uk>) and Cryptoslate (<http://cryptoslate.com>)

via the following steps. First, construct the BERT- and VADER-based sentiment time-series per news source and combine them in a single index across all news sources also according to the method of Sect. 4.3. Second, fit the regression models of Sect. 6.1.1 to generate a range of alternative possible IV $\{\tilde{Y}_t\}$.

- Stage 3: Fit the Koyck-transformed ARDL model and assess which of the different instrumental variables of the previous step at time scale t are appropriate. Perform statistical testing on the suitability of the Transformer/BERT and VADER as IV versus the crypto-specific entropy sentiment signal.
- Stage 4: Using the selected IVs at time scale t , fit the ARDL-MIDAS model to learn the optimal model structure for (p, K, m) , and estimate the MIDAS coefficient basis functions $B(k; \Psi)$ and the geometric decay rate for the infinite ARDL-Koyck transform γ .
- Stage 5: Fit the residuals from Stage 4 with a Range-Scale (R/S) estimation process for the Gegenbauer long memory to determine the Gegenbauer hyperbolic ACF decay parameter d and oscillation index u .

The advantage of this five-stage procedure versus a joint estimation of all components in one stage is that standard R and Python packages may be utilised to perform each stage of the estimation, which we have found to work adequately as outlined in the experimental results section (Sect. 6).

Below we add some further details on stages 3–5.

5.1 Stage 3: Estimation of ARDL(∞) regression

Consider time-series $\{Y_t\}$ given by the daily sentiment score based on our proposed sentiment time-series construction. We will then regress this sentiment scoring method against alternative sentiment extraction methods based on BERT (<https://huggingface.co/nlptown/>, model: bert-base-multilingual-uncased-sentiment) and VADER that we will transform into time-series covariates and denote them by $\{X_t^B\}$ and $\{X_t^V\}$ also constructed from daily measures of sentiment. Then we seek to fit the regression model:

$$Y_t = \beta_0 + \sum_{i=1}^p \gamma_i Y_{t-i} + \beta^B \sum_{j=0}^{+\infty} \phi_B^j X_{t-j}^B + \beta^V \sum_{j=0}^{+\infty} \phi_V^j X_{t-j}^V + \epsilon_t. \tag{39}$$

Since the estimation of the classical ARDL(∞) model via a Koyck geometric parametrisation is standard in the time-series literature, we defer the interested reader to a brief summary provided in the Supplementary Appendix, Section C. We present the analysis of this regression in the results Sect. 6.1.1.

5.2 Stage 4: Estimation of ARDL(∞)-MIDAS regression

Given the modified Koyck transform applied to the ARDL(∞)-MIDAS time-series regression structure as outlined in Proposition 2.1, one can perform estimation of

this model using standard MIDAS model estimation packages such as the R package `midasr` as detailed in Ghysels et al. (2016). This package works with MIDAS models generically specified in a vectorised form as follows:

$$\Phi_p(L)Y_t = \beta_k(L^{1/m})X_t^{(m)} + \epsilon_t. \tag{40}$$

Clearly, the modified Koyck Transformed ARDL-MIDAS model we proposed to develop can readily be represented in this form. Then, the package `midasr` performs a number of different model fitting structures including the U-MIDAS model which is an unrestricted variation of the MIDAS model formulation, in which a frequency alignment transformation and the estimation of the model is performed using Ordinary Least Squares (OLS), see further details in Foroni et al. (2015).

5.3 Stage 5: Estimation of long memory components

We focus on the class of non-oscillatory long-memory models based on the ARFIMA(0, d , 0) type of long memory, i.e. we set $u = 1$, and only have to estimate d , namely the long memory exponent in the model in Proposition 2.1. We adopted this setting as the empirical ACF was not oscillatory and so we simplified the generator of the long-memory fractional difference to the ARFIMA type.

We will estimate this long memory fractional difference parameter d on the residuals of the model from stage 3. One can then obtain an estimate for the strength of long memory d based on a Hurst exponent estimator, by first estimating the Hurst exponent H (Hurst, 1951) and then using the relationship $d = H - 0.5$.

In this work, we adopt the Rescaled Range R/S Hurst exponent estimator that measures the intensity of long-range dependence in a time-series and was originally developed by Hurst (1951). Given a time-series $Y_{t \in \{1,2,3,\dots,T\}}$, the sample mean and the standard deviation process are given by

$$\bar{Y}_T = \frac{1}{T} \sum_{j=1}^T Y_j \quad \text{and} \quad S_t = \sqrt{\frac{1}{t-1} \sum_{j=1}^t (X_j)^2}, \tag{41}$$

where $X_t = Y_t - \bar{Y}_T$ is the mean-adjusted series. Then a cumulative sum series is given by $Z_t = \sum_{j=1}^t X_j$ and the cumulative range based on these sums is

$$R_t = \text{Max} (0, Z_1, \dots, Z_t) - \text{Min} (0, Z_1, \dots, Z_t). \tag{42}$$

The following proposition describes the estimator of H as derived in Mandelbrot (1975).

Proposition 5.1 Consider a time-series $Y_t \in \mathbb{R}$ and define S_t and R_t in Eqs. (41) and (42) respectively, then $\exists C \in \mathbb{R}$ such that the following asymptotic property of the Rescaled Range R/S holds

$$[R/S](T) = \frac{1}{T} \sum_{t=1}^T R_t/S_t \sim CT^H, \quad \text{as } T \rightarrow \infty.$$

In addition, for small sample size T , the Rescaled Range R/S can be adjusted with the following formula Annis and Lloyd (1976):

$$\mathbb{E}[R/S(T)] = \begin{cases} \frac{T-1/2}{T} \frac{\Gamma((T-1)/2)}{\sqrt{\pi(T/2)}} \sum_{j=1}^{T-1} \sqrt{\frac{T-j}{j}}, & \text{for } T \leq 340 \\ \frac{T-1/2}{T} \frac{1}{\sqrt{T\pi/2}} \sum_{j=1}^{T-1} \sqrt{\frac{T-j}{j}}, & \text{for } T > 340 \end{cases},$$

where the $\frac{T-1/2}{T}$ term was added by Peters (1994). The estimate of H can then be obtained by a simple linear regression

$$\log \left((R/S(T) - \mathbb{E}[R/S(T)]) \right) = \log C + H \log T.$$

Hence, we have the following definition for \hat{H} , the estimator of H , based on the unadjusted Rescaled Range R/S analysis and given by:

$$\hat{H}_{R/S} = \frac{T \left(\sum_{t=1}^T \log R/S(t) \log t \right) - \left(\sum_{t=1}^T \log R/S(t) \right) \left(\sum_{t=1}^T \log t \right)}{T \left(\sum_{t=1}^T (\log t)^2 \right) - \left(\sum_{t=1}^T \log t \right)^2}. \quad (43)$$

The empirical confidence interval of \hat{H} given in Eq. (43) with sample size $T = 2^N$ (Weron, 2002) is

$$(0.5 - \exp(-7.33 \log(\log N) + 4.21), \exp(-7.20 \log(\log N) + 4.04) + 0.5).$$

Note that no asymptotic distribution theory has been derived for the estimated Hurst parameter H for R/S analysis, however, we can apply bootstrap methods to find related properties to test for statistical significance of the estimates to detect the long memory properties.

6 Results and discussion

In this section, we will present two real data case studies. The first is to illustrate that the sentiment time-series index constructed in this work is distinct in its information content compared to those extracted either from deep neural network solutions obtained via pre-trained non fine-tuned applications of BERT, or from rule-based systems like VADER. We demonstrate that both of these systems can be combined together to produce a viable time-series index for the design of an effective Instrumental Variable to act as a proxy for the proposed entropy sentiment signal when fitting ARDL(∞)-MIDAS models via OLS, which would otherwise produce biased estimators.

Note that fine-tuning of Transformer-based models and re-construction of domain-specific rule-based systems for sentiment extraction is particularly difficult in this study context which is a “small data” problem relative to typical sentiment studies. The reason for this is that the number of crypto articles available is relatively small compared to the size of corpora used to train deep neural network architectures or inform decisions about potential rule sets that would generalise well. This is one of the key motivations we see for our proposed method, in that it is directly interpretable and applicable in relatively small data contexts such as the case study in question. Therefore, in the first case study we demonstrate that our proposed crypto sentiment index contains significantly different information, compared to the typical approach of just applying BERT and VADER as a black-box package without tailoring or fine-tuning. We demonstrate the value of our sentiment time-series index through an ARDL regression example to show that the covariates of the competing sentiment methods are not strongly expressive of the variation in our daily sentiment signal.

In the second case study, we treat our daily sentiment index as the target response time-series and we seek to explore changes in daily sentiment for cryptocurrency markets in terms of intra-daily crypto price fluctuations and technology factor variations. This will be meaningful for both interpretation of sentiment and price discovery as well as forecasting sentiment at the end of the day, given observations of current intra-daily price and technology network factors. We fit the sequence of infinite lag Koyck-transformed ARDL-MIDAS and ARDL-MIDAS-Transformer Gegenbauer long-memory models to undertake this second case study, exploring along the way each component of the model.

6.1 Case study I: ARDL structure of Y_t and explanatory power of BERT and VADER sentiment methods for instrumental variable construction of \tilde{Y}_t

Let $\{s_{1,\tau_1}, s_{2,\tau_2}, \dots, s_{N,\tau_N}\}$ be the collection of sentences from all articles, ordered according to article publication date τ_i and order of appearance in the article.

Valence Aware Dictionary for sEntiment Reasoning (VADER, Hutto et al., 2014) is a rule-based sentiment model derived from human annotation of online texts. In VADER, first a gold-standard sentiment dictionary is extracted, then validated using qualitative methods (based on human annotation) and lexical features are extracted together with five rules that incorporate grammatical or syntactical conventions that people use to express sentiment intensity.

In our setting, to construct the sentiment index from VADER we use a daily median filter:

$$\text{VADER_index}(\tau) = \text{median}\left(\text{VADER}(s_{i,\tau}), \dots, \text{VADER}(s_{i+k,\tau})\right) \in [-1, 1], \quad (44)$$

where $\{s_{i,\tau}, \dots, s_{i+k,\tau}\}$ are sentences from articles written on the same calendar day τ and $\text{VADER}(\cdot)$ returns the output of the VADER sentiment model.

Naturally, a challenge with this method arises in the context we consider as there is domain-specific knowledge and terminology in the cryptocurrency context that is

not adequately captured by the standard formulation of VADER. Nevertheless, one may find examples of the use of the VADER sentiment model in the crypto market, e.g. Abraham et al. (2018), Kraaijeveld et al. (2020) and Kim et al. (2016). These studies significantly differ from our study in three main aspects: first, the type of sentiment model utilised; second, the type and quality of data used to produce the sentiment model; and third, the way in which sentiment is analysed or utilised. We specifically have not concentrated on social media sentiment, which VADER claims to extract, because of data quality challenges stemming from the short, mainly informal nature of social media text. We have instead focused on public news articles from community accepted reliable websites, which undergo editorial processing before publication, and have developed specific sentiment indices able to capture the particular nature of the vocabulary used in the domain due to our purpose-built crypto dictionary.

Furthermore, alternative methods to rule-based approaches include the word embedding-based models. Word embeddings are real-valued high dimensional vectors that correspond to specific words, and are obtained via a complex non-linear optimisation process. This approach has been prevalent in the neural network-based NLP paradigm, and the optimisation process that obtains the embeddings aims to either learn a decomposition of the document-term matrix of a corpus of documents [e.g. GloVe (Pennington et al., 2014)], or minimise an entropy measure (‘perplexity’) for a model that predicts the word that follows a given word sequence (‘language modelling’), e.g. the Transformer-based BERT (Devlin et al., 2019) that we also utilise in this work).

For BERT, we used a pre-trained model based on an implementation from Hugging Face (<https://huggingface.co/nlptown/>, model: bert-base-multilingual-uncased-sentiment), a group well-known in the NLP community for code quality. The model has been pretrained on a corpus of product reviews, yet we did not fine-tune, i.e. further train, the model with data from our domain for two reasons: first, to our knowledge there are no datasets of crypto-related public articles that have been annotated with sentiment labels, and second, we did not want to undertake this task as it would require manually annotating more than 3000 articles, which places such a process out of our research scope. On the contrary, we want to illustrate: (i) that our approach needs a lot less annotating, i.e. only for the domain dictionary construction, (ii) contrary to computationally expensive neural models which may capture an unclear concept of sentiment, our method is efficient and offers interpretable and informative results, and (iii) the problem of domain mismatch and lack of generalisation of such models in specialised areas, despite the fact that generalisation is one of the main arguments in their favour.

The selected BERT model returns a categorical sentiment score of five levels (0–4), corresponding to the “star rating” of a review: 0 for very negative and 4 for very positive. We again used a daily median filter to construct the sentiment index based on BERT:

$$\text{BERT_index}(\tau) = \text{median}(\text{BERT}(s_{i,\tau}), \dots, \text{BERT}(s_{i+k,\tau})) \in \{0, 1, 2, 3, 4\}, \quad (45)$$

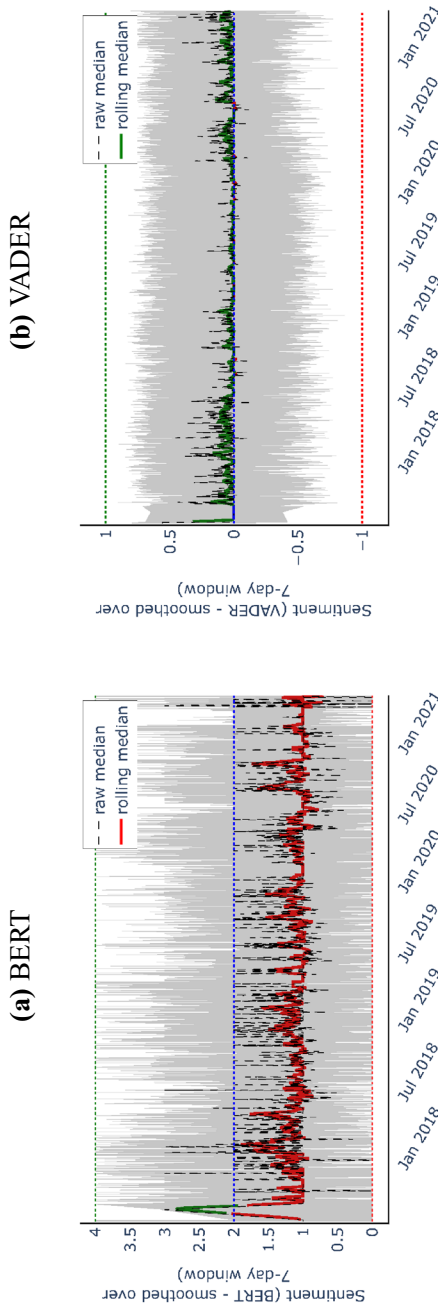


Fig. 4 Sentiment indices based on BERT and VADER with 95% confidence intervals constructed from articles about Bitcoin published on Cryptodaily (<http://cryptodaily.co.uk>) and Cryptoslate (<http://cryptoslate.com>)

where $\{s_{i,\tau}, \dots, s_{i+k,\tau}\}$ are sentences from articles written on the same calendar day τ and $BERT(\cdot)$ is the output of the BERT model on a sentence.

As discussed, in the current case study, in addition to the custom entropy sentiment time series, we also will be employing sentiment time series constructed using the BERT and VADER models. The time series for BERT and VADER will be constructed following the same procedure of Sect. 4.3, where instead of n -grams, we will be assembling sentences from articles on the same day. We plot the sentiment indices constructed based on BERT and VADER in Fig. 4 for BTC, where we indicate polarity with a different choice of color, namely green for positive and red for negative. The corresponding plots for ETH are in the Supplementary Appendix, Section B.

6.1.1 Analysis of entropy based sentiment time-series versus instrumental variable construction from BERT and VADER indices

In this section, we investigate whether the sentiment indices constructed from the BERT and VADER models can be used as explanatory variables for our entropy sentiment index. For this purpose, we fit distributed lag time-series regression models where the BERT and VADER indices are in the set of regressors, whereas the dependent variable is each of the following Entropy Sentiment indices: absolute sentiment strength, negative sentiment and positive sentiment. Subsequently, the significance of the model parameters relevant to the BERT and VADER covariates is assessed. The model structure we adopt is the following:

$$Y_t = \beta_0 + \sum_{i=1}^p \gamma_i Y_{t-i} + \sum_{j=0}^{+\infty} \vec{\beta}_j^T \vec{X}_{t-j} + \epsilon_t, \tag{46}$$

where $\epsilon_t \sim N(0, \sigma^2)$, $\vec{\beta}_j = [\beta_j^B \ \beta_j^V]^T$, $\vec{X}_t = [X_t^B \ X_t^V]^T$ and the superscripts B, V stand for BERT and VADER respectively.

We draw attention to the fact that our goal with this model structure is not to develop a predictive model for our sentiment index Y_t but rather to investigate if the covariates from BERT and VADER have any explanatory power for Y_t , given our novel Entropy Sentiment time-series model. We add an autoregressive component in the covariates to account for potential serial dependence in the dependent variable. If we do not account for that, we risk being misled by the output of the regression: if Y_t has serial dependence then the BERT and VADER covariates can appear to be significant for the dependence structure but that would not mean that they are explanatory for Y_t —it would be an artefact of the chosen model structure.

To obtain a parsimonious representation of the model and thus reduce the number of parameters we have to estimate, we need to find a suitable functional expression for the coefficients $\vec{\beta}$. As discussed in more detail in Sect. 2.1, $\vec{\beta}_j$ must form an L_2 sequence so that the sum of the corresponding term is square summable (provided X_t have finite moments) and the process converges, and one option for the functional

expression of $\vec{\beta}_j$ is to make them a geometric sequence Hill et al. (2001). This is convenient as we can use the generator form of such sequences which has only two parameters to estimate:

$$\begin{aligned} \beta_j^B &= \beta^B \phi_B^j, & 0 < \phi_B < 1 \\ \beta_j^V &= \beta^V \phi_V^j, & 0 < \phi_V < 1 \end{aligned} \tag{47}$$

for each covariate correspondingly. The final form of our model then becomes:

$$\begin{aligned} Y_t &= \alpha + \sum_{i=1}^p \gamma_i Y_{t-i} + \sum_{j=0}^{+\infty} \beta_j^B X_{t-j}^B + \sum_{j=0}^{+\infty} \beta_j^V X_{t-j}^V + \epsilon_t \\ &= \alpha + \sum_{i=1}^p \gamma_i Y_{t-i} + \beta^B \sum_{j=0}^{+\infty} \phi_B^j X_{t-j}^B + \beta^V \sum_{j=0}^{+\infty} \phi_V^j X_{t-j}^V + \epsilon_t. \end{aligned} \tag{48}$$

If our analysis using this model shows that Y_t is conditionally independent, or only weakly dependent on the covariates of BERT and VADER, then our sentiment index captures different information to these alternative models. In that case, it may be valuable, for example, to consider all indices together as components in a multi-modal sentiment index model, where modality would correspond to the source, namely underlying model, of each component that captures a different sentiment aspect. Or, as we demonstrate in Case Study II, one can use the BERT and VADER sentiment signals combined as an Instrumental Variable for estimation of a model using our Entropy Sentiment signal that has infinite lag structure transformed by a Koyck transform method. We fit the model in Eq. (48) in rolling windows with a length of three months with an one-month overlap, where the regressors are the average-smoothed, z-scaled BERT and VADER indices.

To begin with, we regress the Entropy Sentiment Y_t against its $1, \dots, p$ lags with an AR(p) model and obtain the raw residuals which we denote as \tilde{E}^y in what follows next.

Model I: BERT covariate for Y_t , IV regression with BERT

We construct the IV as follows:

$$\tilde{E}^y_{t-1} = \tilde{\mu}_0 + \tilde{\mu}_1 X_{t-1}^B + \epsilon_{t-1}. \tag{49}$$

We perform a t -test on the regression parameters and compute the IV \tilde{E}^y_{t-1} using only those that are statistically significant. Then we conduct the following OLS regression for E_t^y :

$$E_t^y = \alpha(1 - \phi) + \phi \tilde{E}^y_{t-1} + \beta^B X_t^B + \epsilon_t - \phi \epsilon_{t-1}. \tag{50}$$

At this stage, we want first to evaluate the quality of the IV and then the quality of the fit with respect to X^B . To evaluate the instrumental variable we test for autocorrelation in the error terms; if they are not autocorrelated then we have successfully constructed an IV that is not correlated with the errors and the use of OLS was appropriate. We test for error autocorrelation using the Breusch–Godfrey test (LM

test for autocorrelation, Breusch, 1978), for which the null hypothesis is the lack of serial correlation of any order up to p . If we have evidence to accept the null then we proceed to test the parameters for statistical significance, otherwise we consider this fit invalid—we would need to construct a different IV, for example by adding more structure to the corresponding model. If they are significant, we next assess the model by means of the AIC.

The procedure just described for Model I is followed for all subsequent models, therefore we will next present only the different model structures we used.

Model II: VADER covariate for Y_t , IV regression with VADER

IV regression:

$$\tilde{E}^y_{t-1} = \tilde{\mu}_0 + \tilde{\mu}_1 X^V_{t-1} + \epsilon_{t-1}. \tag{51}$$

E^y_t regression:

$$E^y_t = \alpha(1 - \phi) + \phi \tilde{E}^y_{t-1} + \beta^V X^V_t + \epsilon_t - \phi \epsilon_{t-1}. \tag{52}$$

Model III: BERT and VADER covariates for Y_t , IV regression with BERT

IV regression:

$$\tilde{E}^y_{t-1} = \tilde{\mu}_0 + \tilde{\mu}_1 X^B_{t-1} + \epsilon_{t-1}. \tag{53}$$

E^y_t regression:

$$E^y_t = \alpha(1 - \phi) + \phi \tilde{E}^y_{t-1} + \beta^B X^B_t + \beta^V X^V_t + \epsilon_t - \phi \epsilon_{t-1}. \tag{54}$$

Model IV: BERT and VADER covariates for Y_t , IV regression with VADER

IV regression:

$$\tilde{E}^y_{t-1} = \tilde{\mu}_0 + \tilde{\mu}_1 X^V_{t-1} + \epsilon_{t-1}. \tag{55}$$

E^y_t regression:

$$E^y_t = \alpha(1 - \phi) + \phi \tilde{E}^y_{t-1} + \beta^B X^B_t + \beta^V X^V_t + \epsilon_t - \phi \epsilon_{t-1}. \tag{56}$$

Model V: BERT and VADER covariates for Y_t , IV regression with BERT and VADER

IV regression:

$$\tilde{E}^y_{t-1} = \tilde{\mu}_0 + \tilde{\mu}_1 X^B_{t-1} + \tilde{\mu}_2 X^V_{t-1} + \epsilon_{t-1}. \tag{57}$$

E^y_t regression:

$$E^y_t = \alpha(1 - \phi) + \phi \tilde{E}^y_{t-1} + \beta^B X^B_t + \beta^V X^V_t + \epsilon_t - \phi \epsilon_{t-1}. \tag{58}$$

Finally, we compare the AIC results for the appropriate models and choose the best among those.

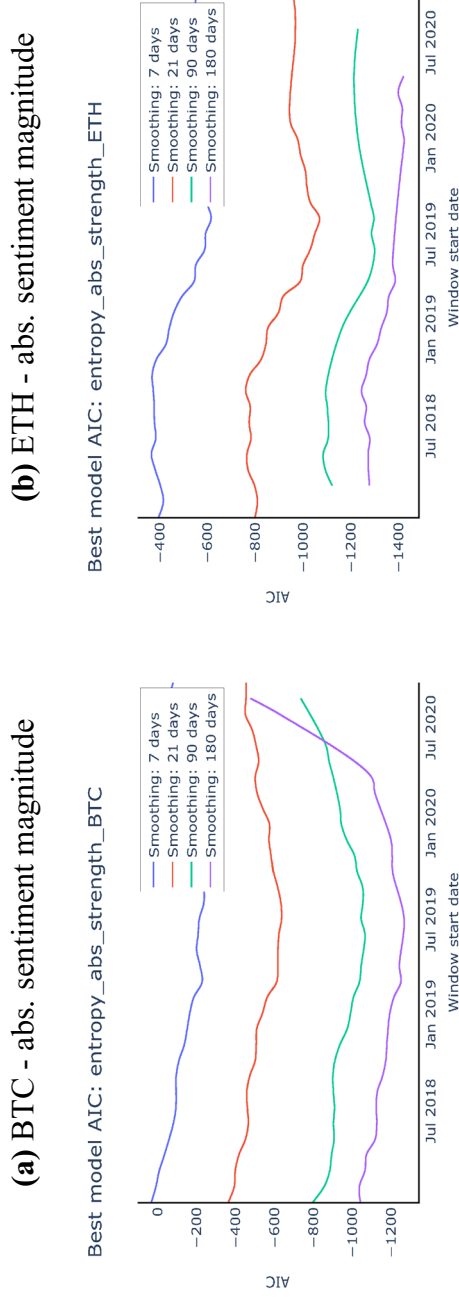


Fig. 5 Stage I rolling window fits: best AR(p) model AIC score

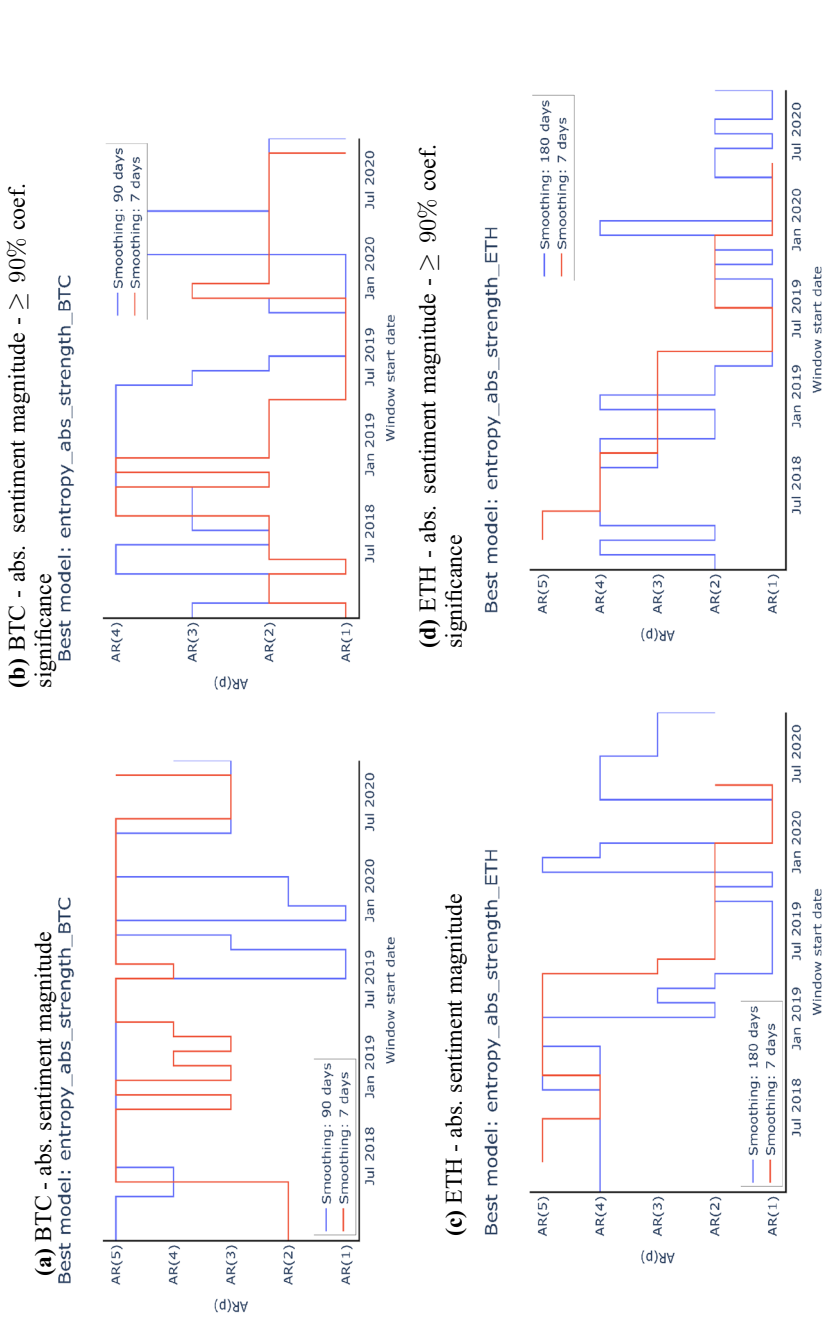


Fig. 6 Stage I rolling window fits: best AR lag structure

6.1.2 Stage 2: Model selection and calibration

We begin with an analysis of the fitting process for Stage 2 of Sect. 5. For fitting the Stage 2 model we performed a sequence of fits in rolling windows of seven months, overlapping by one month. Before each sequence of fits, we considered a range of different settings for the response sentiment signal and the model structure:

1. we considered different smoothing options for the sentiment response; we applied a rolling median filter with window sizes of 1 week, 3 weeks, 3 months, and 6 months;
2. we considered a range of lag structures for the autoregressive covariates, with lag order $p = 1, \dots, 5$.

We demonstrate here the results for the total sentiment—the remaining equivalent results for the polarity of positive and negative sentiment are available in Supplementary Appendix, Section C.1.

For each configuration, we stored the successfully fitted models and ranked them according to the AIC score. In Fig. 5, we see the AIC of the best fitting model per window for all smoothing parametrisations. Informed by these traces, we focus on the best-performing model to continue our analysis. The next question we addressed was whether we would need to have a different lag structure per window to capture the sentiment signal's variability. For this purpose, we plot the lag order per fitting window for the best-fitting model (Fig. 6, left panels), and in addition, we plot the

Table 1 Percentage of windows that fail to reject the null hypothesis of the performed tests

Test	Entropy Abs. strength-BTC	Entropy negative- BTC	Entropy positive- BTC	Entropy Abs. strength-ETH	Entropy negative- ETH	Entropy positive- ETH
1-Week smoothed signals						
KS	0.00	0.00	0.00	0.00	0.00	0.00
VS	100.00	96.67	100.00	93.33	97.06	94.12
KPSS	100.00	100.00	100.00	100.00	100.00	100.00
3-Week smoothed signals						
KS	0	0	0	0	0	0
VS	91.18	97.06	97.06	93.75	90.91	97.06
KPSS	100	100	100	100	100	100
3-Month smoothed signals						
KS	0	0	0	0	0	0
VS	96.55	92.86	93.1	100	91.3	100
KPSS	100	100	100	100	100	100
6-Month smoothed signals						
KS	0	0	0	0	0	0
VS	95.83	92.31	100	77.27	100	86.21
KPSS	100	100	100	100	100	100

lag order per fitting window for the best fitting model after dropping the coefficients that were not statistically significant to any level up to 90% (Fig. 6 right panels). We also plot for comparison (in red trace) the worst performing model according to the AIC. We observe that the differences in model selection are significant when we ignore the coefficient significance level, and therefore we were correct in utilising a different lag structure per fitting window. When we also account for the significance level, we observe that for the absolute sentiment magnitude we still benefit from adapting the lag structure per window.

6.1.3 Stage 3: Evaluation of model fit and quality of instrumental variable design

Having selected the best Stage 2 model per signal window, we then evaluate the quality of the fitting by performing a series of statistical tests on the residuals. Furthermore, with this analysis we aim to verify our hypothesis of normality of the error distribution. We employ the widely used Kwiatkowski–Phillips–Schmidt–Shin (KPSS, Kwiatkowski et al., 1992), the Kolmogorov–Smirnov (KS, Dimitrova et al., 2020), and the Vasicek–Song (VS, Lequesne & Regnault, 2020; Song, 2002; Vasicek, 1976) tests to investigate trend stationarity and normality of the error distribution. Table 1 shows the percentage of fitted windows that did not produce evidence to reject the null hypothesis of the tests. The null hypothesis of the KPSS test is that the signal is trend stationary, whereas the null hypothesis of the KS, Vasicek and Song tests is that the tested signal follows a normal distribution. In our experiments, we approximated the distribution of the null hypothesis of the Vasicek test via Monte Carlo sampling, which is the basic formulation of this test.

From Table 1 we first note that the windowed signal is trend stationary. Furthermore, with regards to the error distribution, we observe that the KS test always rejects the null hypothesis of normality, but the Vasicek entropy test shows that for the majority of the fits the error normality assumption holds. This ambiguity is resolved if we consider that the KS test places equal weight on the median tendency of the error distribution and its tails. Therefore, if there is even a little skewness or kurtosis in the distribution, it may reject the null of normal distribution whilst this may not be true. This is the effect we see here, where the Vasicek test does not reject the null hypothesis as it places higher importance on the median tendency of the distribution.

To illustrate this point further and justify our choice of normal errors, we first construct the Q – Q plots of the residuals in a window of the absolute sentiment magnitude for ETH, which we see in Table 1 that exhibits the lowest compliance with the normal error assumption. Figure 7 shows the Q – Q plots obtained for normal and t -Student distributions. We observe that there is evidence of heavier tails in the residuals than a normal distribution, and therefore we fit t -Student distributions of differing degrees of freedom to reflect different strengths of kurtosis that may be present; we illustrate this for three settings 3, 10, and 15 degrees of freedom. In so doing, we note that the best fit is achieved for t distributions with many degrees of freedom ($t = 15$), which is close to a normal, and therefore the assumption of normal errors is adequate for our purposes. Next, in Fig. 8 we plot the density estimate of the residuals in the same window, and in Fig. 9 we plot the residuals against




Fig. 7 Stage I residual distribution example: 6-month window starting on 2020-04-27 for ETH absolute sentiment strength

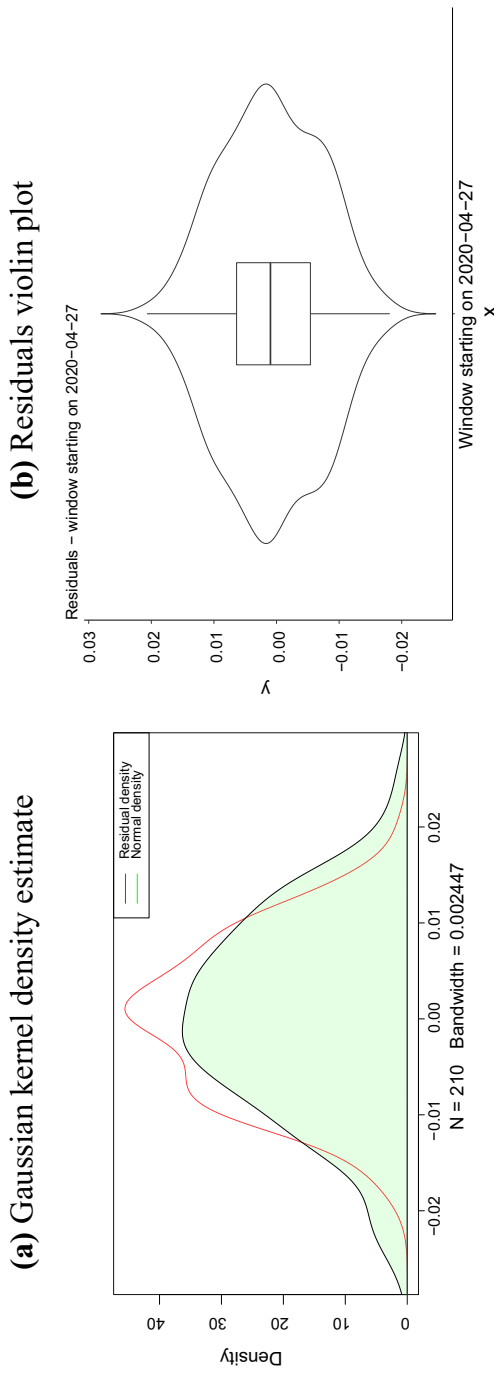
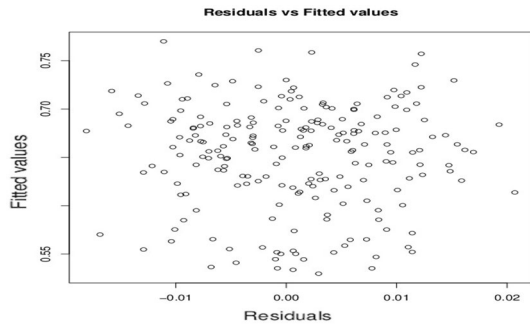


Fig. 8 Stage I residual distribution density example: 6-month window starting on 2020-04-27 for ETH absolute sentiment strength. A normal density is illustrated in green and the residual density is plotted in red (colour figure online)

Fig. 9 Stage I residuals vs fitted response values: 6-month window starting on 2020-04-27 for ETH absolute sentiment strength



the fitted response to verify the absence of any uncaptured trend structure. Both the shape of the density estimates and the residual against fitted values scatterplot verify our assumption.

6.1.4 Stage 3: Quantitative evaluation of the explanatory power of the BERT and VADER indices for instrumental variable design

Our goal is to quantify the explanatory power of the BERT and VADER covariates for the sentiment response, after we have obtained an appropriate instrumental variable. Table 2 shows the results of the rolling window regressions for a selection of windows in which we have obtained an instrumental variable with one of the explored model structures. The quality of the instrumental variable was measured with the Breusch–Godfrey test (Breusch, 1978), and only windows that showed decorrelated errors (95% significance levels of the coefficients in the IV regression) were kept for the regression against sentiment.

First, we observe that there were very few periods over the course of almost 4 years that the BERT and VADER models could explain part of our entropy sentiment index, for all polarities universally. Secondly, we note that the model structures most suitable for the IV regression were M3 and M5, where we used both covariates for the sentiment response, and either BERT only (M3) or BERT and VADER (M5) for the IV regression. Thirdly, it is interesting to see that BERT and VADER are more explanatory for BTC news sentiment rather than ETH, which may be attributed to Bitcoin's popularity leading to simpler, less technical and, therefore, easier to understand language in the relevant articles. Fourthly, we observe that there are specific periods where the covariate significance was high ($\geq 99\%$): those starting March 3, 2018, April 8 2018, and 27 February 2020. March–October 2018 was a period of relatively low volatility in the Bitcoin price which culminated in Winter 2018 and was sparking a lot of concern among retail investors at the time for fear of a price plunge. In Fig. 4, we see that BERT is negative all of the time during that period, and VADER shows some of its few negative sentiment indications at that time. Therefore, we would expect that both indices would help explain more of the negative sentiment content of our index at that time, as we see that is the case. This also applies to the 7-month window starting 8 April 2018 and extending to November of the same year, when the lack of volatility was more pronounced. The same,

Table 2 Statistical significance levels of coefficients of BERT/VADER sentiment covariates

Sentiment type	Window start date	Best IV model	BERT coefficient with standard errors	BERT significance level	VADER coefficient with standard errors	VADER significance level
1-Week smoothed signals						
Entropy negative-BTC	2018-03-09	M5	-1.13E-01 ± 7.41E-03	**	3.1E-01 ± 2.1E-02	*
Entropy positive-BTC	2019-02-02	M3	4.91E-02 ± 2.34E-04	**	6.42E-02 ± 5.80E-04	.
Entropy positive-ETH	2018-04-08	M1	6.90E-02 ± 1.83E-03	*	-	-
Entropy positive-ETH	2018-05-08	M1	6.35E-02 ± 1.24E-03	*	-	-
3-Week smoothed signals						
Entropy negative-BTC	2018-03-09	M5	-1.39E-01 ± 1.0E-04	***	4.03E-01 ± 1.91E-04	***
Entropy negative-BTC	2018-04-08	M5	-1.2E-01 ± 2.0E-04	**	3.13E-01 ± 4.38E-04	**
Entropy positive-BTC	2019-02-02	M1	3.96E-02 ± 2.07E-05	*	-	.
Entropy positive-ETH	2019-09-30	M5	6.04E-02 ± 3.34E-04	*	-2.28E-01 ± 1.28E-03	*
3-Month smoothed signals						
Entropy Abs. strength-BTC	2018-04-08	M2	-	.	2.0E-01 ± 4.71E-08	*
Entropy Abs. strength-BTC	2018-05-08	M4	1.0E-02 ± 5.34E-06	*	3.0E-02 ± 1.11E-04	.
Entropy Abs. strength-BTC	2020-05-27	M1	7.80E-02 ± 2.86E-07	*	-	.
Entropy Abs. strength-BTC	2020-08-25	M3	5.98E-02 ± 1.31E-06	.	3.50E-02 ± 5.65E-07	*
Entropy negative-BTC	2020-02-27	M3	9.08E-02 ± 1.35E-07	**	3.24E-02 ± 4.68E-08	**
Entropy positive-BTC	2018-04-08	M4	1.23E-02 ± 8.1E-09	.	2.11E-01 ± 1.03E-07	*
Entropy positive-BTC	2018-05-08	M4	1.18E-02 ± 6.18E-07	*	7.75E-02 ± 1.28E-05	.
6-Month smoothed signals						
Entropy Abs. strength-BTC	2019-10-30	M2	-	.	1.13E-01 ± 7.27E-07	*
Entropy Abs. strength-BTC	2019-11-29	M2	-	.	1.43E-01 ± 2.28E-07	**

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 '.' 1




Fig. 10 Cumulative sentiment content of the BTC and ETH sentiment indices constructed from articles from Cryptodaily (<http://cryptodaily.co.uk>) and Cryptoslate (<http://cryptoslate.com>)

but less pronounced and at a higher price level, phenomenon is also observed for the window starting 27 February 2020 and extending to September 2020.

Both of these observations explain why the BERT and VADER covariates appear to have some explanatory power for our negative sentiment index for BTC. However, it is still evident that our index captures different sentiment content than the state-of-the-art alternative approaches, which we will also demonstrate qualitatively next, before proceeding to show how to leverage our sentiment index in the crypto space.

6.1.5 Stage 3: Qualitative evidence for the difference between the entropy index and the BERT and VADER indices

In the previous section, we statistically proved and quantified the difference in the content of the entropy-based sentiment index versus the indices based on BERT and VADER. In this section we provide visual evidence that illustrates the difference between the indices. In Fig. 10 we show the cumulative sentiment content of the different indices for BTC and ETH, each relativised with respect to the beginning of the time period we study.

The rates of change of the traces reveal whether the aggregation of sentiment is continuous or intermittent over time, and whether it is affected by sentiment polarity or not. We observe the linear growth of cumulative sentiment, and almost constant rate of change for the most period of time, which means that there is a smooth variation in the way information is captured: a growing amount of different information is accumulated over time, as opposed to certain abrupt rare events being the main drivers of the index. The fact that the growth is linear for both assets is in part due to the fact that the information is rich, i.e. there are articles on Cryptodaily and Cryptoslate regarding BTC and ETH almost on a daily basis during the time period we focus on. Note that the linear pattern is not something dictated by any of the sentiment index models and may change if we study different news sources or assets.

From Fig. 10 we understand that BERT and VADER accumulate sentiment at different rates. VADER is slower at gathering information, and in addition gathers sentiment information that is different to any of the information captured by our entropy indices. This is because its rate of change is significantly lower than the rate of the entropy indices and almost constant for the case of BTC (Fig. 10, left panel). For ETH (Fig. 10, right panel), we note that at the beginning of the period, VADER is able to capture information at a higher rate, which however saturates around the middle of 2018, meaning it was unable to capture any significant information, and seems to slightly pick up again at the beginning of 2021. These findings, respectively, are in agreement with the almost flat line we observe in the smoothed index of Fig. 4 (right, VADER) after mid 2018, and the also almost flat line for Ethereum (see Supplementary Appendix, Section B, Figure 2b) in early 2019. We can attribute this to the fact that VADER is trained on general online media content and therefore lacks the specificity required for the crypto domain. For the change of rate we observe in the case of Ethereum around early 2019, we remark that this may be indicative of a change in the language of the articles, namely before that period

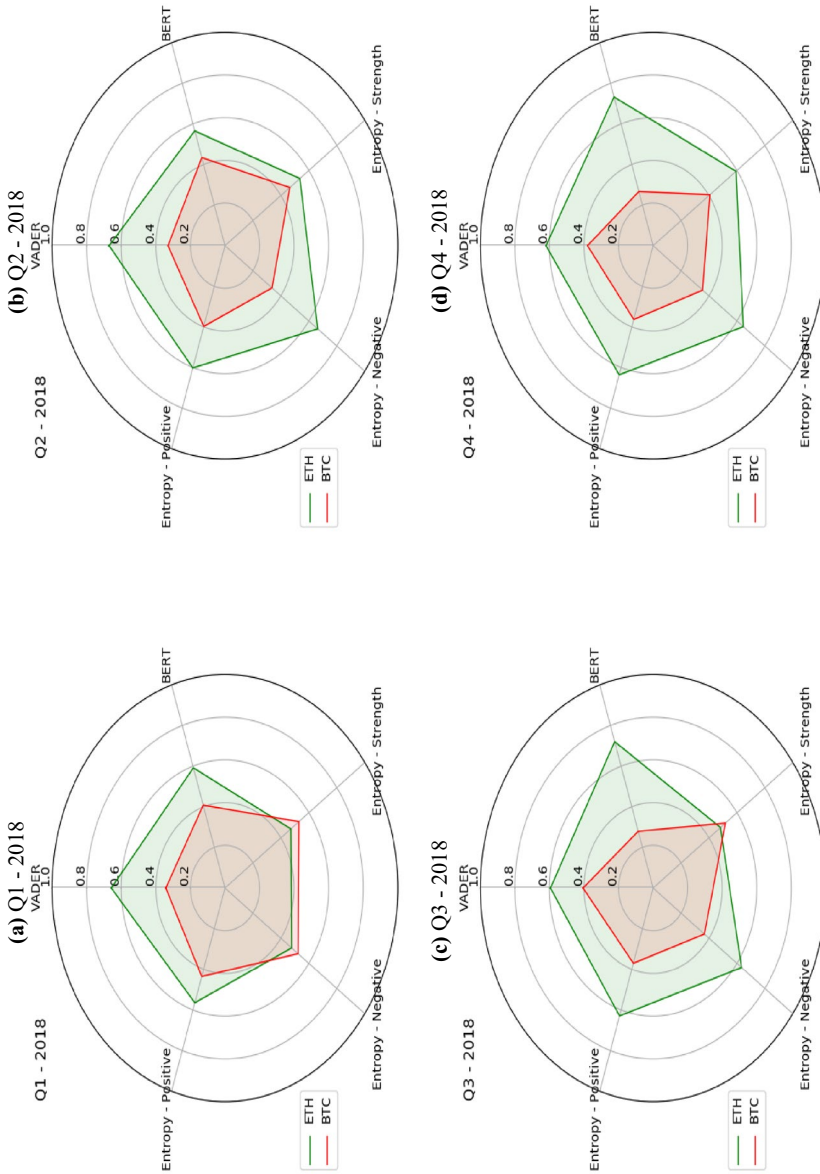


Fig. 11 Sentiment index IQR per quarter for each of the sentiment indices

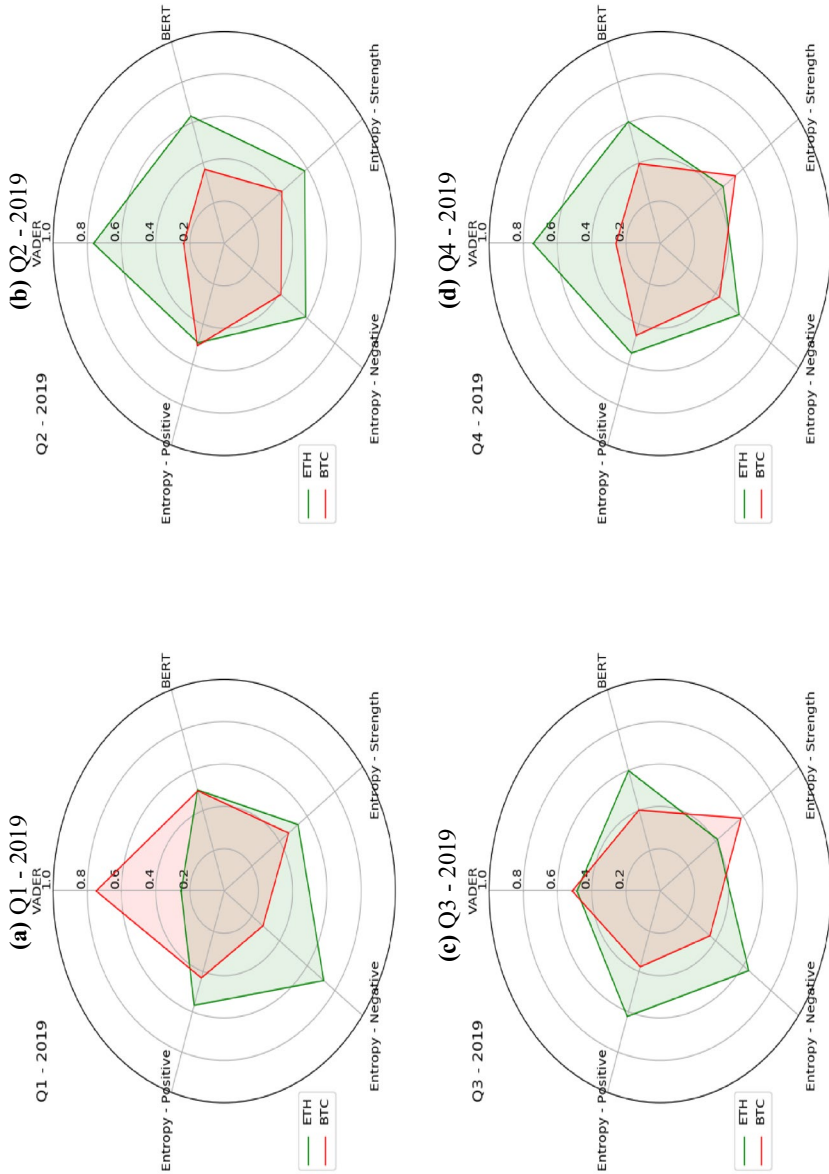


Fig. 12 Sentiment index IQR per quarter for each of the sentiment indices



Fig. 13 Sentiment index IQR per quarter for each of the sentiment indices

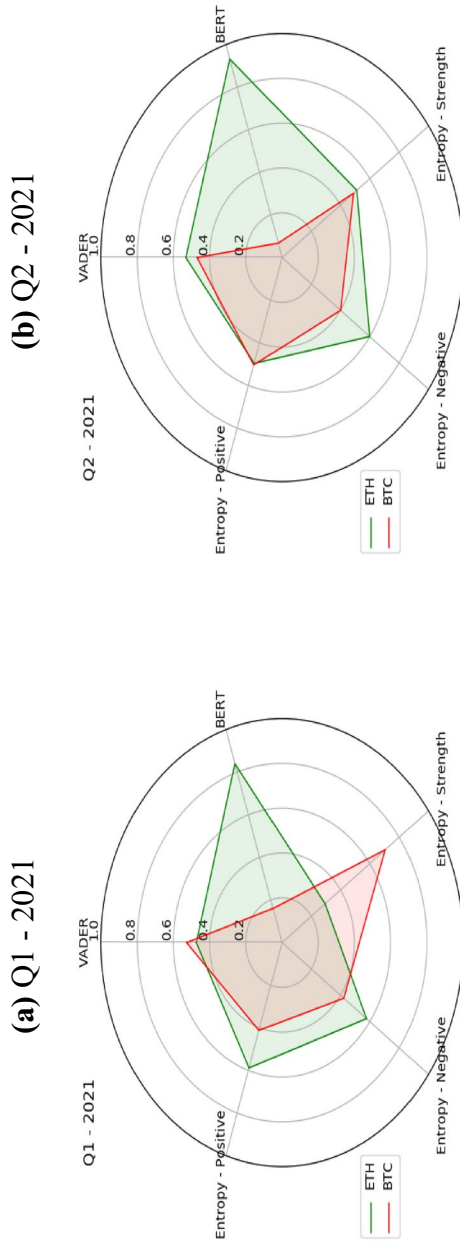


Fig. 14 Sentiment index IQR per quarter for each of the sentiment indices

authors wrote in a less technical and more similar way to the average social media user, which we know is the type of language VADER was trained on. On the contrary, we see that BERT's growth rate is higher than the rest of the indices in the case of BTC and almost parallel to our positive polarity entropy sentiment index for ETH. For the latter, given that BERT has identified predominantly negative sentiment as we saw in Supplementary Appendix, Section B, Figure 3a we interpret this observation as a difference in the understanding of sentiment polarity between our positive entropy index and BERT, given that both are based on the same texts: what the entropy index perceives as positive sentiment, BERT sees as mostly negative.

In addition, we see that for BTC (Fig. 10, left), the negative index is almost identical to the absolute sentiment strength, which means that most of the time the negative index captures the same global sentiment tendency as the absolute strength index. However, for ETH (Fig. 10, right), even if the two indices appear to grow almost in parallel in terms of information, they intersect and deviate in mid 2021. This is consistent with the price boom that ETH experienced at the beginning of the second quarter of 2021, hence the negative sentiment would no longer dominate in the absolute sentiment strength.

Finally, in both plots of Fig. 10, and more evidently in the case of ETH (right), we can see that the rate of change is not constant all of the time—it is higher at the start of the period. This is more distinctly shown in the positive entropy index. This period follows the period of the ETH price peak of January 2018 and we can see that our positive sentiment index reacts more strongly to the resulting sentiment signal.

To further illustrate the difference between the signals captured by the examined sentiment constructions we construct the plots of Figs. 11, 12, 13 and 14 where we have summarised the sentiment content of the indices in yearly quarters. These figures are obtained by using the interquartile range (IQR) to summarise the sentiment content over each quarter, which would capture the volatility of the signal at each period.

In addition to demonstrating that the sentiment content of the indices is very different, these plots may also reveal significant information about the behaviour of the sentiment indices. We also provide comparative quarterly results for sentiment analysis of the median sentiment and the sentiment volatility under each polarity of sentiment signal. The median sentiment analysis is provided in Supplementary Appendix, Section C.2. In Fig. 11, we can see that the IQR of the VADER and BERT indices did not significantly change in Q3 and Q4 of 2018. Given that during those particular quarters Bitcoin was going through a very low volatility period followed by a price crash, hence the sentiment signal, due to mixed reporting about price speculation, was highly varying, it is evident that the BERT and VADER models were unable to capture this variation in sentiment. On the contrary, we observe that the proposed sentiment signals we developed based on Entropy indices are very reactive to the sentiment signal variation. In Q1–Q2 2019, when Bitcoin price started rising again, news reporting started becoming more positive capturing the investors regaining optimism.

Similarly, observing Figs. 13 and 14, we see that the positive Entropy index clearly shows an increase in volatility during Q3–Q4 2020 and Q1 2021, when Bitcoin price was trending upwards, while in the second quarter of 2021, when there

Table 3 The options that we explored for the lag structure in the regression covariates

Regression component	Lag structure
Response autoregressive component	1–5 days, i.e. 1–5 daily lags
BTC high-frequency component	1–5 days, i.e. 24–120 hourly lags
Hash Rate low-frequency component	1–6 months, i.e. 30–180 daily lags

was a price reversal, we remark that the variability of the negative sentiment index starts increasing to reflect the high volatility of the price at that time. Similarly, we observe that BERT reacts significantly to the sentiment signal change between end of 2020 and beginning of 2021, even though we know from Fig. 4 (left - BERT) that at that time BERT mostly identified negative sentiment. VADER on the other hand seemed unable to adapt to any change in sentiment after Q4 2020, showing only little change in IQR.

6.2 Case study II: MIDAS-Koyck transform calibration results

In the second case study, we are interested in exploring a regression relation between daily sentiment, as captured by our daily Entropy index, versus a technology-based time-series covariate given by daily hash rate, and the price signal given by intra-daily asset price on an hourly time frame.

We focus on the intra-daily close price of Bitcoin, and extract the price and hash rate data from CoinMarketCap (<https://coinmarketcap.com/>) for the period of September 2017–May 2021.

6.2.1 Model selection for finite-lag ARDL-MIDAS-Transformer time-series regressions

Before proceeding with the analysis, we investigate the fitting of a wide range of model parametrisations with respect to the autoregressive lags, the covariates and the parameters of the Almon weight functions, which, for this study, was the Exponential Almon function.

We adopt in this section the framework of infinite-lags in the ARDL structure, where we apply the Koyck-MIDAS transform to obtain a reparametrised model family as described in Eq. (16) using the instrumental variable of Eq. (30) constructed from BERT and VADER sentiment signals.

Informed by previous research (Chalkiadakis et al., 2020), we explored the options for the lag structure that are included in Table 3. In addition, we used all three different types of sentiment polarity (absolute sentiment strength, positive, negative), and explored a number of optimisation routines from the R package `optim`. We assessed the models that fit successfully for the whole range of data using the AIC and Mean Squared Error (MSE) criterion and provide the top-4 fitting models in Table 4 according to these two model selection criteria.

Table 4 Best fitting models for different configurations of lag structures and training set-ups

Model	Sentiment type	Autoregressive lags (days)	Hash rate lags (months)	BTC hourly close price lag (days)	Smoothed covariates (1-week median filter)	Hash rate Almon lag number	BTC close price Almon lag number	AIC	MSE
M1	Abs. sent. strength	5	6	5	True	1	1	-3.9969E+02	0.04112
M2	Abs. sent. strength	5	6	5	True	1	2	-3.9805E+02	0.04110
M3	Abs. sent. strength	5	6	5	False	2	1	-3.9007E+02	0.04146
M4	Abs. sent. strength	5	6	5	False	1	1	-3.8732E+02	0.04163

Table 5 Percentage of rolling windows with statistically significant covariates in the regression. The model configuration is in terms of: (autoregressive daily lags, BTC/h daily lags, HR monthly lags, HR Almon param., BTC Almon param., Box–Cox applied, Box–Cox Lambda if applied, window step size, window size)

Model configuration	% Significant covariates	% Significant BTC/h Almon parameters	% Significant HR Almon parameters	% Significant at 99% level	% Significant at 99.9 % level
5-5-6-1-1-TRUE_-0.5-30-1260	61.364	37.5	50	61.364	61.364
5-5-6-1-1-TRUE_-0.5-30-1080	60	30	50	60	60
5-5-6-1-1-TRUE_0.30-1260	59.091	37.5	50	59.091	59.091
5-5-6-1-1-FALSE_0.30-1260	59.091	37.5	50	59.091	59.091
5-5-6-1-1-TRUE_-1-30-1260	59.091	37.5	50	59.091	59.091
5-5-6-1-1-TRUE_0.30-1080	58.182	25	50	58.182	58.182
5-5-6-1-1-TRUE_-1-30-900	56.97	43.333	36.667	56.97	56.97
5-5-6-1-1-TRUE_0.30-900	56.818	40.625	40.625	56.818	56.818
5-5-6-1-1-TRUE_-0.5-30-900	56.818	37.5	40.625	56.818	56.818
5-5-6-1-1-FALSE_0.30-1080	56.364	20	50	56.364	56.364
5-5-6-1-1-FALSE_0.30-900	54.545	33.333	43.333	54.545	54.545
5-5-6-1-2-FALSE_0.30-1260	51.515	16.667	16.667	51.515	51.515
5-5-6-1-2-TRUE_-1-30-1080	50.505	27.778	50	50.505	50.505
5-5-6-1-2-TRUE_-0.5-30-1080	50.505	27.778	44.444	50.505	50.505
5-5-6-1-2-TRUE_-0.5-30-1260	50	12.5	37.5	50	50
5-5-6-1-1-TRUE_-1-30-720	47.934	27.273	29.545	47.934	47.934
5-5-6-1-2-TRUE_-1-30-1260	47.727	12.5	25	47.727	47.727
5-5-6-1-1-TRUE_0.30-720	47.521	27.273	31.818	47.521	47.521
5-5-6-1-2-TRUE_0.30-1080	47.273	15	40	47.273	47.273
5-5-6-1-1-TRUE_-0.5-30-720	47.107	25	29.545	47.107	47.107
5-5-6-1-2-TRUE_-1-30-900	46.853	38.462	30.769	46.853	46.853
5-5-6-1-2-TRUE_-0.5-30-900	46.753	32.143	35.714	46.753	46.753
5-5-6-1-1-FALSE_0.30-720	45.868	31.818	27.273	45.868	45.868

Table 5 (continued)

Model configuration	% Significant-covariates	% Significant BTC/h Almon parameters	% Significant HR Almon parameters	% Significant at 99% level	% Significant at 99.9 % level
5-5-6-1-2-TRUE_0-30-900	45.455	41.667	29.167	45.455	45.455
5-5-6-1-2-TRUE_0-30-1260	45.455	12.5	25	45.455	45.455
5-5-6-1-2-FALSE_0-30-900	44.056	26.923	34.615	44.056	44.056
5-5-6-1-2-FALSE_0-30-1080	43.939	8.333	41.667	43.939	43.939
5-5-6-1-2-TRUE_-1-30-720	36.364	12.5	34.375	36.364	36.364
5-5-6-1-2-TRUE_0-30-720	35.758	26.667	30	35.758	35.758
5-5-6-1-2-FALSE_0-30-720	35.227	18.75	28.125	35.227	35.227
5-5-6-1-2-TRUE_-0-5-30-720	34.091	15.625	25	34.091	34.091

The rows in bold indicate the best-performing models considering all columns

6.2.2 Assessment of time-series ARDL-MIDAS-Transformer regression model over time

Based on the model search of the previous section, we proceed to study the calibration of the best fitting model in a rolling window fashion, to study how the regression relationship between the response sentiment and the covariates evolves over time. We explore four window sizes, i.e. 720-, 900-, 1080-, and 1260-day rolling windows with a step size of 30 days, and conduct our analysis on the best fitting model according to the AIC (M1) and the MSE (M2).

In addition to the parametrisation of the models as presented in Table 4, we also perform the fitting after applying the Box–Cox transform on the sentiment response, as the initial fittings revealed heteroscedasticity in the residuals. The Box–Cox transform (Box & Cox, 1964) aims to reduce that by transforming the response to resemble a normal variable. The transform, in the case of a positive variable, which is our sentiment response, is given as follows:

$$y(\lambda) = \begin{cases} \frac{y^\lambda - 1}{\lambda}, & \text{if } \lambda \neq 0; \\ \log(y), & \text{otherwise,} \end{cases} \quad (59)$$

where $-5 \leq \lambda \leq 5$. In practice, we found a change in the regression performance for a small range of negative λ values close to zero, and we therefore investigated only the values of $\lambda = -1, -0.5, 0$.

The results of the rolling window fits are presented in Table 5 in terms of the percentage of rolling windows where the covariates of the regression were statistically significant. The list is ordered with first key the total number of windows with statistically significant covariates, and second and third keys the total number of windows where the hourly close price (BTC/h) and Hash Rate (HR) covariates, respectively, were significant. Based on the results of this analysis, we select the best window to employ in the following studies, which is a 1080-day window for the best fitting model according to MSE (M2 in Table 4), and a 1260-day window for the best-fitting model according to AIC (M1 in Table 4).

Next, we provide details about the study of the Almon polynomial coefficients over time for the low- and high-frequency covariates. The majority of the details for the model studies showing dynamics of the Exponential-Almon lag structures for each covariate are provided in the accompanying Supplementary Appendix in Section D, and here we present the results for Model M1 (best according to AIC) with the coefficients for the BTC hourly close price and the daily hash rate in Figs. 15 and 16, when fitting the model without applying the Box–Cox transform and with Box–Cox applied with $\lambda = 0$.

6.2.3 Statistical significance of interactions between data at mixed frequencies

In this section, we continue to explore the results of model M1 by looking at the statistical significance of the coefficients of the autoregressive response covariate, as

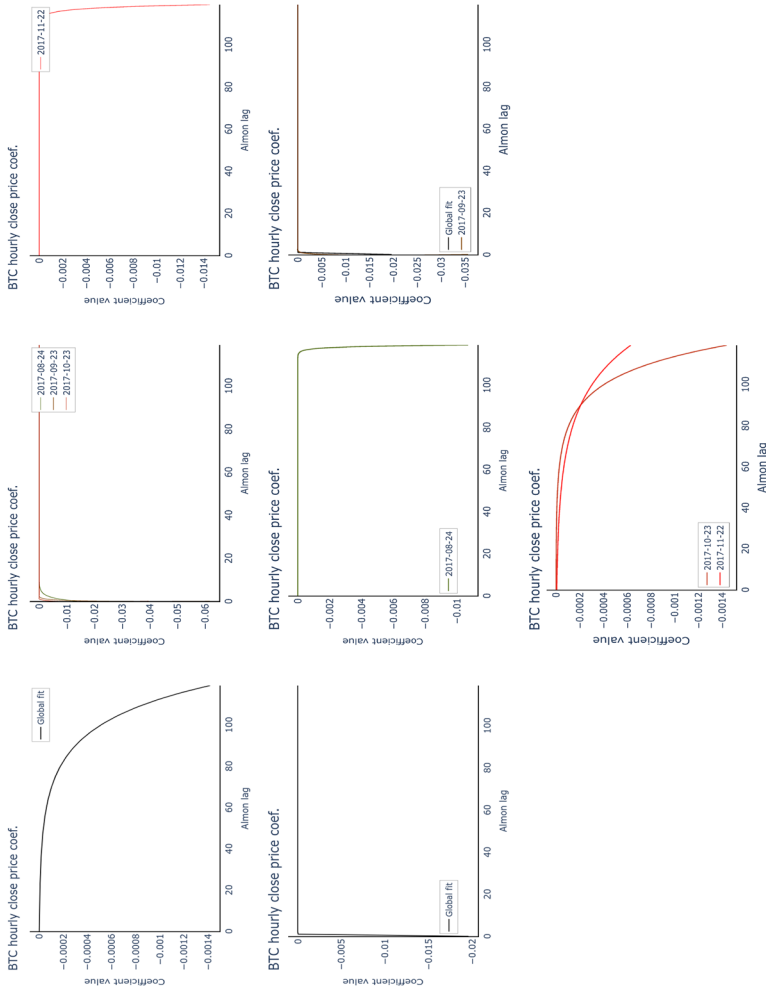


Fig. 15 M1: Exponential Almon coefficients structure of BTC hourly close price over time for the model fit on the complete dataset. Top row plots: No Box–Cox Transform applied. Bottom plots: Box–Cox Transform with $\lambda = 0$ applied

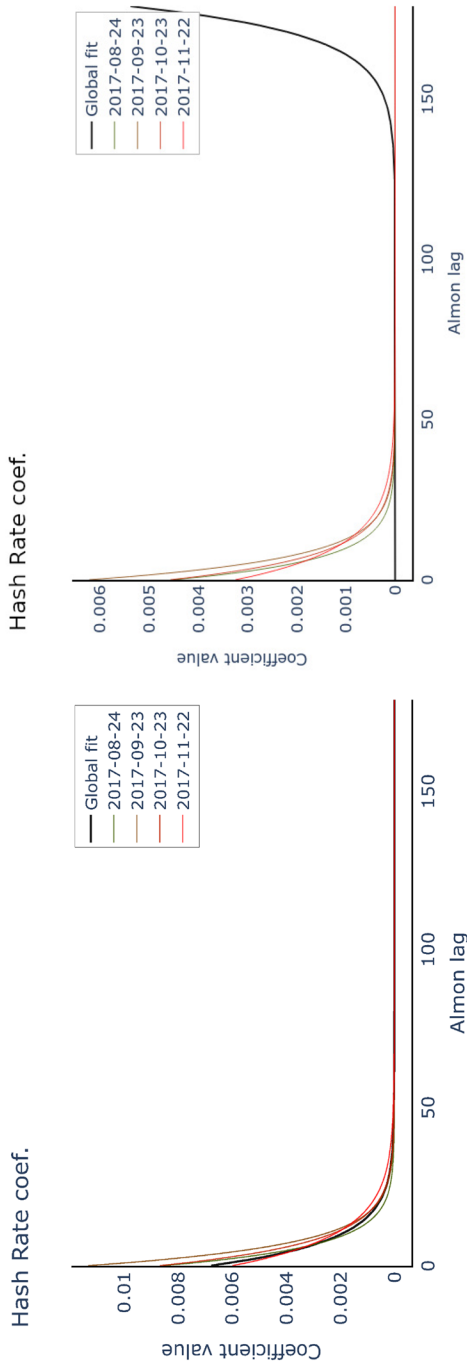


Fig. 16 M1: Exponential Almon coefficients structure of Hash Rate over time for the model fit on the complete dataset. Left plots: No Box-Cox Transform applied. Right plots: Box-Cox Transform with $\lambda = 0$ applied

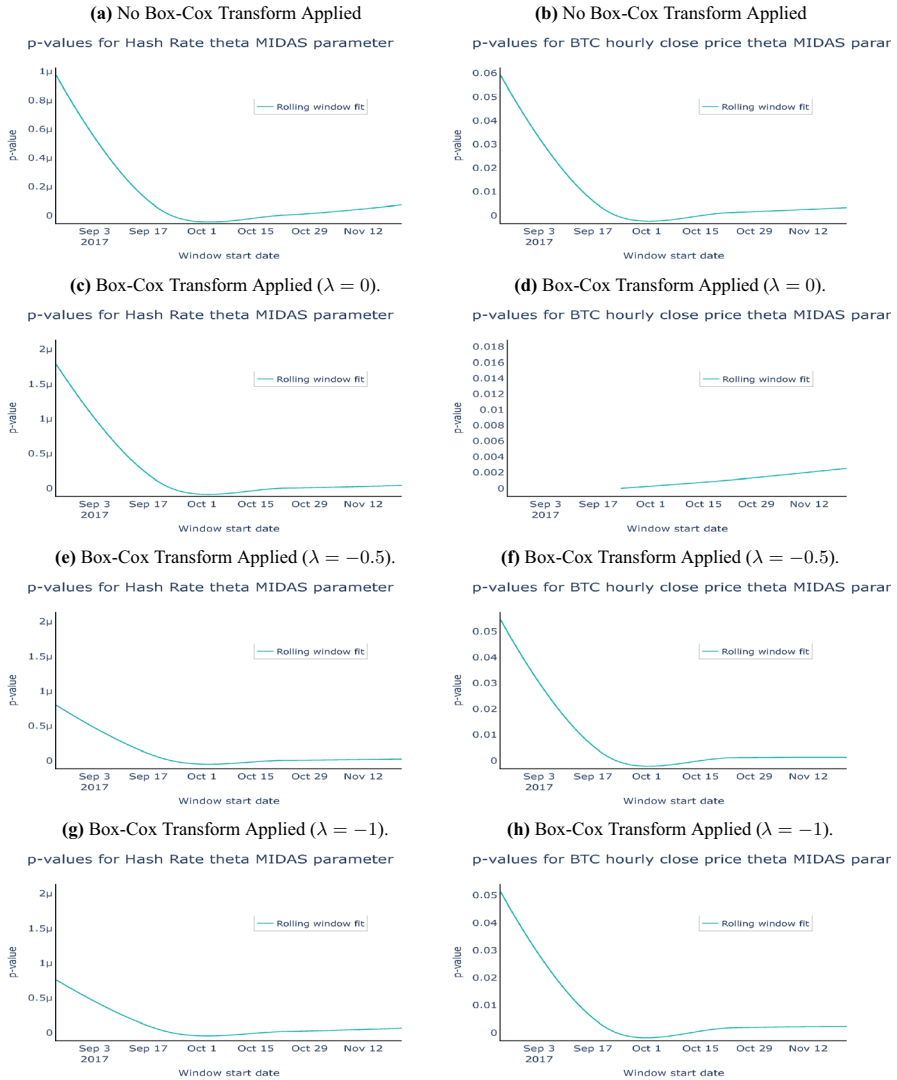
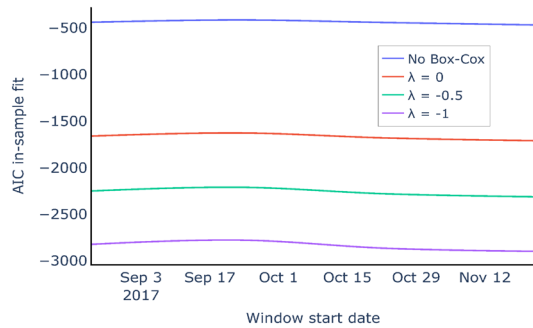


Fig. 17 M1: Left subplots: Hash Rate low-freq. covariate. Right subplots: BTC hourly close price high-freq. covariate. p -values for the statistically significant Exponential Almon parameters (β_1, ψ)

well as the low- and high-frequency covariates. In Fig. 17 we plot the p -values for the fitted M1 models in which the Distributed Lag MIDAS covariates were found to be statistically significant. For the corresponding plots for model M2, please refer to the Supplementary Appendix, Section D. In terms of performance, we observe in Fig. 18 that the models fit with a Box–Cox transform exhibit an improved performance compared to the models without.

Fig. 18 M1: AIC for each rolling window fit



Furthermore, in the Supplementary Appendix, Section D.1, we demonstrate the p -values on the significance of the AR lags of the MIDAS-Koyck transformed coefficients in Eq. (16) for the cases of no Box–Cox transformation and a few cases of Box–Cox transformed fitted models, for models M1 and M2.

6.2.4 Mixed-data long memory structures

We now seek to explore how the weight function that we employ for the MIDAS coefficients affects the persistence properties of the high- and low-frequency covariates. To investigate, we estimate the Hurst exponent from the covariates, the response, and the residuals per fitting window and explore their relationship. We plot the estimated Hurst exponents of the response and the residuals in Figs. 19, 20, 21, 22 and 23. Note that the Hurst exponent is a positive number upper-bounded by 1—the small bias that is observed in some estimates is consistent with the behaviour of the estimator in large-scale synthetic studies that we performed when testing the behaviour of this estimator, before applying it in this real data case study.

We observe that the long memory strength in the response, as captured via the residuals, is attenuated compared to the long memory in the covariates as a result of the MIDAS structure in the model. Therefore, the long memory structure of the covariates is not trivially transferred to the response in such settings, which has to be considered if including the long memory is an important desired feature for the model.

Note that should we wish to use Gegenbauer polynomial MIDAS weights, we could estimate the d and u parameters from the residuals. The long memory d is defined as $d = H - 0.5$, whilst, in this instance, we can estimate the cyclic frequency u by observing the autocorrelation plots of the residuals per window, as we observed that the ACF exhibits almost no oscillation around the x axis, which means that $u = 0$, i.e. the long memory is coming from an ARFIMA-type process. We have provided a detailed analysis of the autocorrelation profiles of the regression for model M1 and model M2 with and without the Box–Cox transform in the Supplementary Appendix, Section D.2.




Fig. 19 M1: BTC/h (left) and Hash rate (right) vs Residuals and Sentiment response Hurst exponents. No Box-Cox transform was applied. Red, diamond-shaped markers correspond to the response-derived Hurst, whilst the blue, circle-shaped markers to the residual-derived Hurst (colour figure online)




Fig. 20 M1: BTC/h (left) and Hash rate (right) vs residuals and Sentiment response Hurst exponents. The Box-Cox transform with $\lambda = 0$ was applied. Red, diamond-shaped markers correspond to the response-derived Hurst, whilst the blue, circle-shaped markers to the residual-derived Hurst (colour figure online)




Fig. 21 M1: BTC/h (left) and Hash rate (right) vs Residuals and Sentiment response Hurst exponents. The Box–Cox transform with $\lambda = -0.5$ was applied. Red, diamond-shaped markers correspond to the response-derived Hurst, whilst the blue, circle-shaped markers to the residual-derived Hurst (colour figure online)




Fig. 22 M1: BTC/h (left) and Hash rate (right) vs residuals and sentiment response Hurst exponents. The Box-Cox transform with $\lambda = -1$ was applied. Red, diamond-shaped markers correspond to the response-derived Hurst, whilst the blue, circle-shaped markers to the residual-derived Hurst (colour figure online)





Fig. 23 M2: BTC/h (left) and Hash rate (right) vs Residuals and Sentiment response Hurst exponents. No Box–Cox transform was applied. Red, diamond-shaped markers correspond to the response-derived Hurst, whilst the blue, circle-shaped markers to the residual-derived Hurst (colour figure online)

(a) M1 - best model according to AIC



(b) M2 - best model according to MSE




Fig. 24 Forecasting with rolling-windows fitted models. The true sentiment signal has been detrended in the following plots as we did for the fitting. The red trace denotes the cubic spline-smoothed signal, and the grey band denotes the 95% confidence interval. (color figure online)

Table 6 Forecasting with rolling-window fitted models

(a) Performance of model M1 (best according to AIC)

30-Day forecast start	M1
2021-02-05	1.15E-01
2021-03-07	7.47E-02
2021-04-06	6.36E-02

(b) Performance of model M2 (best according to MSE)

30-Day forecast start	M2
2020-09-08	9.58E-08
2020-11-07	3.45E-02
2021-02-05	1.21E-02
2021-04-06	5.78E-02

6.2.5 Mixed-data infinite-lag Koyck regressions: estimation of γ

In this section, we study the in-sample fitting performance of the ARDL-MIDAS model with the Koyck transform to obtain tractability with an infinite number of lags. We estimate the model with and without (see Sects. 6.2.1–6.2.4) the adjustments dictated by the Koyck transform, and then, having obtained the regression coefficients for the autoregressive lags of Eq. (16) before and after the transformation, we can form a system of equations to solve for decay parameter γ of the transform under the constraint $0 < \gamma < 1$.

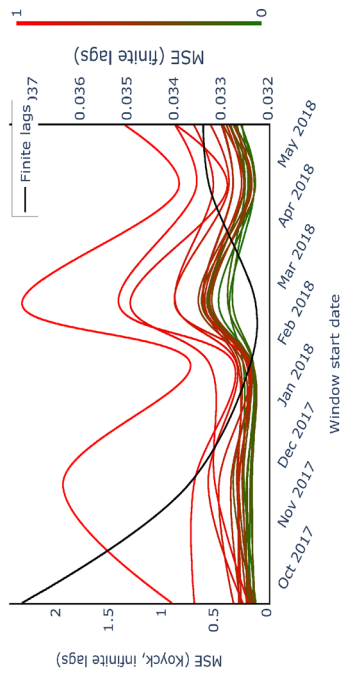
We inform our selection for the autoregressive covariate in Eq. (16) from our analysis in Case Study I, and use as instrumental variable a linear combination of the Entropy Sentiment and the BERT and VADER sentiment covariates, which decorrelate the sentiment response from the errors in this formulation as studied in Case Study I. Specifically, the IV in the current study is the difference between the Entropy Sentiment index, median-smoothed weekly in a rolling window fashion, and the average of the also weekly median-smoothed BERT and VADER sentiment signals.

6.2.6 Trend forecasting and the effect of the infinite-lag Koyck transform

We finalise the study with an analysis of how well the proposed infinite lag ARDL-MIDAS-Transformer long memory time-series regression models perform on out-of-sample forecasting of daily crypto sentiment. To assess the forecasting performance of the fitted models, we use each model fitted on a rolling window to forecast one month ahead of the window. First, for the models without the infinite lag Koyck adjustment, we perform the fitting without applying the Box–Cox transform, and present the results in Fig. 24 and Table 6(a) and (b).

Second, we extend the estimated models with the Koyck adjustment of Eq. (16) for a range of values of γ on a grid: 0.01, 0.05, 0.1, ..., 0.95, where the step size after

(b) M2 - best model according to MSE



(a) M1 - best model according to AIC

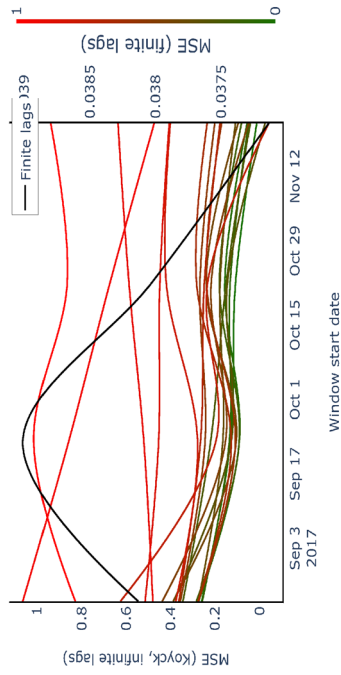


Fig. 25 MSE forecasting performance with the rolling-window fitted models of Sect. 6.2.1, after having added the infinite-lag Koyck transform for a range of values for γ . The latter are sampled from a grid on (0, 1), which is shown with the color graduation: the redder the trace, the higher the value of γ . The reference model without the Koyck transform is plotted in black. (color figure online)

0.1 is 0.05. We evaluate the forecast performance in terms of MSE and illustrate the results in Fig. 25. We observe that the performance after the Koyck adjustment (left y-axis) is significantly diminished, which means that the geometric decay we adopted for the coefficients in the Koyck transform is not optimal in this setting.

7 Conclusion

In this work, we have proposed a novel class of time-series regression models that incorporate several relevant interpretable features: infinite-lag ARDL regressions, Mixed Data Sampling (MIDAS) multiple time resolution regression structures, Deep Neural Network architectures for Instrumental Variable design for reducing the estimation bias in the ARDL-MIDAS-Transformer class of models, and, finally, fractional integration in the form of Gegenbauer long memory polynomials for the MIDAS configuration. Each of these model components is carefully explained and detailed and then a thorough real data statistical analysis is undertaken on cryptocurrency market sentiment constructed daily from news articles. The daily sentiment is then regressed against intra-daily hourly closing price dynamics and money supply, as captured by mining Hash Rate. Overall, in our study and application we see the advantage of the infinite lag ARDL-MIDAS-Transformer time-series regression models in the rigorous incorporation of sophisticated long-memory signal characteristics, interpretation capabilities for the interplay between multi-resolution covariates, as well as fitting and forecasting performance. These advances will make IV regressions with expressive neural network instruments more accessible to researchers with time-series, econometrics familiarity, and will also assist with MIDAS model development and selection. The results of the real data studies demonstrated that end-of-day sentiment can be accurately forecast at any time point intra-daily, given the appropriate time-series model structure, from the hybrid ARDL-MIDAS-Transformer class of models proposed in this manuscript. We are confident that the proposed model class will find interesting applications not only in cryptocurrencies but also in traditional financial markets, commodities markets, text and sentiment signals analysis. We leave the investigation of case studies in these research areas, as well as further informative extensions to the MIDAS polynomial structure to future work.

8 Software and technical appendix

The reader is also referred to the supplementary Technical Appendix for additional results and analyses.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s42521-023-00079-9>.

Acknowledgements This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability Code and data for reproducibility purposes are available at <https://github.com/ichalkiad/ardlmidasdnn>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abraham, J., Higdon, D., Nelson, J., & Ibarra, J. (2018). Cryptocurrency price prediction using tweet volumes and sentiment analysis. *SMU Data Science Review*, 1(3), 1. <https://scholar.smu.edu/datasciencereview/vol1/iss3/1>
- Andreou, E., Ghysels, E., & Kourtellos, A. (2011). Forecasting with mixed-frequency data. In Michael P. Clements & David F. Hendry (Eds.), *The Oxford Handbook of Economic Forecasting* (online ed., pp. 225–246). Oxford Academic. <https://doi.org/10.1093/oxfordhb/9780195398649.013.0009>
- Andreou, E., Ghysels, E., & Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business & Economic Statistics*, 31(2), 240–251.
- Annis, A., & Lloyd, E. (1976). The expected value of the adjusted rescaled Hurst range of independent normal summands. *Biometrika*, 63(1), 111–116. <https://doi.org/10.1093/biomet/63.1.111>
- Beran, J. (1994). *Statistics for long-memory processes* (Vol. 61). CRC Press. <https://doi.org/10.1201/9780203738481>
- Box, G. E. P., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2), 211–252. <http://www.jstor.org/stable/2984418>
- Breusch, T. S. (1978). Testing for autocorrelation in dynamic linear models. *Australian Economic Papers*, 17(31), 334–355. <https://doi.org/10.1111/j.1467-8454.1978.tb00635.x>
- Chalkiadakis, I., Yan, H., Peters, G. W., & Shevchenko, P. V. (2021). Infection rate models for COVID-19: Model risk and public health news sentiment exposure adjustments. *PLoS ONE*, 16(6), 1–39. <https://doi.org/10.1371/journal.pone.0253381>
- Chalkiadakis, I., Zaremba, A., Peters, G. W., & Chantler, M. J. (2020). On-chain analytics for sentiment-driven statistical causality in cryptocurrencies. *Blockchain: Research and Applications*, 3(2), 100063. <https://doi.org/10.1016/j.bcra.2022.100063>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the Association for Computational Linguistics: Human language technologies (Long and short papers)* (Vol. 1, pp. 4171–4186). Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1423>
- Dhrymes, P. J., Klein, L. R., & Steiglitz, K. (1970). Estimation of distributed lags. *International Economic Review*, 11(2), 235–250.
- Dimitrova, D. S., Kaishev, V. K., & Tan, S. (2020). Computing the Kolmogorov–Smirnov distribution when the underlying CDF is purely discrete, mixed, or continuous. *Journal of Statistical Software*, 95(10), 1–42. <https://doi.org/10.18637/jss.v095.i10>
- Faroni, C., Marcellino, M., & Schumacher, C. (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 178(1), 57–82.
- Ghysels, E., Kvedaras, V., & Zemlyš, V. (2016). Mixed frequency data sampling regression models: The R package midasr. *Journal of Statistical Software*, 72(1), 1–35.
- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2004). *The MIDAS touch: Mixed data sampling regressions, manuscript*. University of North Carolina and UCLA.

- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2005). There is a risk-return trade-off after all. *Journal of Financial Economics*, 76(3), 509–548.
- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131(1–2), 59–95.
- Ghysels, E., Sinko, A., & Valkanov, R. (2007). Midas regressions: Further results and new directions. *Econometric Reviews*, 26(1), 53–90.
- Hannan, E. J. (1965). The Estimation of Relationships Involving Distributed Lags. *Econometrica: Journal of the Econometric Society*, 3(1), 206–224. <https://doi.org/10.2307/1911896>
- Harris, Z. (1954). Distributional structure. *Word*, 10(23), 146–162. <https://doi.org/10.1080/00437956.1954.11659520>
- Hartford, J., Lewis, G., Leyton-Brown, K., & Taddy, M. (2017). Deep IV: A flexible approach for counterfactual prediction. In: Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17, JMLR.org, pp. 1414–1423.
- Hassani, H., Beneki, C., Unger, S., Mazinani, M. T., & Yeganegi, M. R. (2020). Text mining in big data analytics. *Big Data and Cognitive Computing*, 4(1), 2504–2289. <https://doi.org/10.3390/bdcc4010001>
- Hill, R. C., Griffiths, W. E., & Judge, G. G. (2001). *Undergraduate econometrics* (2nd ed.). Wiley.
- Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116(1), 770–799. <https://doi.org/10.1061/TACEAT.0006518>
- Hutto, C. J., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the 8th international conference on weblogs and social media (ICWSM), 1–4 June 2014* (pp. 216–225). The AAAI Press. <https://www.scinapse.io/papers/2099813784>
- Jiao, X., Yin, Y., Shang, L., Jiang, X., Chen, X., Li, L., Wang, F., & Liu, Q. (2020). TinyBERT: Distilling BERT for natural language understanding. In: *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online (pp. 4163–4174). <https://doi.org/10.18653/v1/2020.findings-emnlp.372>
- Kim, Y. B., Kim, J. G., Kim, W., Im, J. H., Kim, T. H., Kang, S. J., & Kim, C. H. (2016). Predicting fluctuations in cryptocurrency transactions based on user comments and replies. *PloS One*, 11(8), e0161197. <https://doi.org/10.1371/journal.pone.0161197>
- Klein, L. R. (1958). The Estimation of Distributed Lags. *Econometrica: Journal of the Econometric Society*, 26(4), 553–565. <https://doi.org/10.2307/1907516>
- Koyck, L. M. (1954). *Distributed lags and investment analysis*. Contributions to economic analysis. North-Holland Pub. Co.
- Kraaijeveld, O., De Smedt, J., et al. (2020). The predictive power of public twitter sentiment for forecasting cryptocurrency prices. *Journal of International Financial Markets, Institutions and Money*. <https://doi.org/10.1016/j.intfin.2020.101188>
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54(1), 159–178. [https://doi.org/10.1016/0304-4076\(92\)90104-Y](https://doi.org/10.1016/0304-4076(92)90104-Y)
- Lequesne, J., & Regnault, P. (2020). vsgoftest: An R package for goodness-of-fit testing based on Kullback–Leibler divergence. *Journal of Statistical Software, Code Snippets*, 96(1), 1–26. <https://doi.org/10.18637/jss.v096.c01>
- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1), 1–167. <https://doi.org/10.2200/S00416ED1V01Y201204HLT016>
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65. <https://doi.org/10.1111/j.1540-6261.2010.01625.x>
- Mandelbrot, B. B. (1975). Limit theorems on the self-normalized range for weakly and strongly dependent processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 31(4), 271–285. <https://doi.org/10.1007/BF00534968>
- Pang, B., Lee, L., et al. (2008). Foundations and trends® in information retrieval. *Foundations and Trends® in Information Retrieval*, 2(1–2), 1–135.
- Pennebaker, J., Chung, C., Ireland, M., Gonzales, A., & Booth, R. (2007). *The development and psychometric properties of LIWC2007*. Austin, TX: University of Texas at Austin.
- Pennington, J., Socher, R., & Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532–1543). Association for Computational Linguistics. <https://doi.org/10.3115/v1/D14-1162>

- Peters, E. E. (1994). *Fractal market analysis: Applying chaos theory to investment and economics*, Wiley finance editions. Wiley.
- Singh, A., Hosanagar, K., & Gandhi, A. (2020). Machine learning instrument variables for causal inference. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, Association for Computing Machinery, New York, NY, USA, 2020 (pp. 835–836).
- Song, K.-S. (2002). Goodness-of-fit tests based on Kullback–Leibler discrimination information. *IEEE Transactions on Information Theory*, 48(5), 1103–1117. <https://doi.org/10.1109/18.995548>
- Stock, J. H., & Trebbi, F. (2003). Retrospectives: Who invented instrumental variable regression? *Journal of Economic Perspectives*, 17(3), 177–194.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Proceedings of the 27th international conference on neural information processing systems-Volume 2, NIPS'14* (pp. 3104–3112). MIT Press.
- Tenney, I., Das, D., Pavlick, E. (2019a). BERT Rediscovered the Classical NLP Pipeline. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R. T., Kim, N., Van Durme, B., Bowman, S. R., Das, D., et al. (2019b). What do you learn from context? Probing for sentence structure in contextualized word representations. In: *7th International Conference on Learning Representations, ICLR 2019*, New Orleans, LA, USA 2019.
- Vasicek, O. (1976). A test for normality based on sample entropy. *Journal of the Royal Statistical Society: Series B (Methodological)*, 38(1), 54–59. <https://doi.org/10.1111/j.2517-6161.1976.tb01566.x>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, U., & Polosukhin, I. (2017). Attention is all you need. In *Proceedings of the 31st international conference on neural information processing systems, NIPS'17* (pp. 6000–6010). Curran Associates Inc.
- Weron, R. (2002). Estimating long-range dependence: Finite sample properties and confidence intervals. *Physica A: Statistical Mechanics and its Applications*, 312(1), 285–299. [https://doi.org/10.1016/S0378-4371\(02\)00961-5](https://doi.org/10.1016/S0378-4371(02)00961-5)
- Xu, L., Chen, Y., Srinivasan, S., de Freitas, N., Doucet, A., & Gretton, A. (2021). Learning Deep Features in Instrumental Variable Regression. In: *International Conference on Learning Representations, ICLR 2021 (online conference)*
- Zhang, L., & Liu, B. (2017). *Sentiment analysis and opinion mining* (pp. 1152–1161). Springer US. https://doi.org/10.1007/978-1-4899-7687-1_907

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Ioannis Chalkiadakis^{1,4}  · Gareth W. Peters² · Matthew Ames³

✉ Ioannis Chalkiadakis
ioannis.chalkiadakis@iscpif.fr

Gareth W. Peters
garethpeters@ucsb.edu

Matthew Ames
matt.ames@resilientml.com

¹ School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, Scotland, UK

² Department of Statistics and Applied Probability, University of California Santa Barbara, Santa Barbara, California, USA

³ ResilientML, Melbourne, Australia

⁴ Institut des Systèmes Complexes de Paris Île-de-France, CNRS/UAR 3611, Paris, France