

APPLICATION OF MACHINE LEARNING
MODELS AND INTERPRETABILITY
TECHNIQUES TO IDENTIFY THE
DETERMINANTS OF THE PRICE
OF BITCOIN

2022

BANCO DE **ESPAÑA**
Eurosistema

Documentos de Trabajo
N.º 2215

José Manuel Carbó and Sergio Gorjón

**APPLICATION OF MACHINE LEARNING MODELS AND INTERPRETABILITY
TECHNIQUES TO IDENTIFY THE DETERMINANTS OF THE PRICE OF BITCOIN**

APPLICATION OF MACHINE LEARNING MODELS AND INTERPRETABILITY TECHNIQUES TO IDENTIFY THE DETERMINANTS OF THE PRICE OF BITCOIN

José Manuel Carbó (*)

BANCO DE ESPAÑA

Sergio Gorjón (*)

BANCO DE ESPAÑA

(*) The authors work in the Financial Innovation Division of Banco de España, and appreciate the comments received from José Manuel Marqués, Ana Fernández and Gabriel Pérez Quirós. The opinions and analyses expressed in this paper are the responsibility of the authors and, therefore, do not necessarily match with those of the Banco de España or the Eurosystem.

The Working Paper Series seeks to disseminate original research in economics and finance. All papers have been anonymously refereed. By publishing these papers, the Banco de España aims to contribute to economic analysis and, in particular, to knowledge of the Spanish economy and its international environment.

The opinions and analyses in the Working Paper Series are the responsibility of the authors and, therefore, do not necessarily coincide with those of the Banco de España or the Eurosystem.

The Banco de España disseminates its main reports and most of its publications via the Internet at the following website: <http://www.bde.es>.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

© BANCO DE ESPAÑA, Madrid, 2022

ISSN: 1579-8666 (on line)

Abstract

So-called cryptocurrencies are becoming more popular by the day, with a total market capitalization that exceeded \$3 trillion at its peak in 2021. Bitcoin has emerged as the most popular among them, with a total valuation that reached an all-time high of \$68,000 in November 2021. However, its price has historically been subject to large and abrupt fluctuations, as the sudden drop in the months that followed once again proved. Since bitcoin looks all set to continue growing while largely concentrating its activity in unregulated environments, concerns have been raised among authorities all over the world about its potential impact on financial stability, monetary policy, and the integrity of the financial system. As a result, building a sound and proper regulatory and supervisory framework to address these challenges hinges upon achieving a better understanding of both the critical underlying factors that influence the formation of bitcoin prices and the stability of such factors over time. In this article we analyse which variables determine the price at which bitcoin is traded on the most relevant exchanges. To this end, we use a flexible machine learning model, specifically a Long Short Term Memory (LSTM) neural network, to establish the price of bitcoin as a function of a number of economic, technological and investor attention variables. Our LSTM model replicates reasonably well the behaviour of the price of bitcoin over different periods of time. We then use an interpretability technique known as SHAP to understand which features most influence the LSTM outcome. We conclude that the importance of the different variables in bitcoin price formation changes substantially over the period analysed. Moreover, we find that not only does their influence vary, but also that new explanatory factors often seem to appear over time that, at least for the most part, were initially unknown.

Keywords: Bitcoin, machine learning, LSTM, interpretability techniques.

JEL classification: C40, C45, G12, G15.

Resumen

Las criptomonedas son cada día más populares. Sin ir más lejos, en 2021 su capitalización agregada llegó a superar los 3 billones de dólares, una cifra nunca antes registrada. Dentro de este amplio ecosistema, destaca el caso del *bitcoin*, cuyo precio alcanzó los 68.000 dólares en 2021, que marca un máximo histórico. Sin embargo, la evolución de esta cotización dista de ser consistente en el tiempo, pues se observan, con frecuencia, fluctuaciones considerables y abruptas, como las ocurridas en los meses que siguieron a los valores récord antes señalados. Ante el más que previsible crecimiento del *bitcoin* y la concentración de su actividad mayoritariamente en ambientes no regulados, crece la preocupación entre las autoridades financieras de todo el mundo acerca de su potencial impacto en la estabilidad financiera, en la política monetaria y en la integridad del sistema financiero. En consecuencia, apremia avanzar en la construcción de un marco regulatorio y supervisor sólido y consistente ante estos desafíos. A estos efectos, resulta necesario mejorar el grado de comprensión tanto de los factores subyacentes que influyen en la formación del precio del *bitcoin* como de su estabilidad a lo largo del tiempo. En este documento analizamos cuáles son las variables que determinan el precio al que se negocia el *bitcoin* en las plataformas de intercambio más relevantes. Para ello, utilizamos un modelo flexible de aprendizaje automático; concretamente, una red neuronal *Long Short Term Memory* (LSTM), para establecer el precio del *bitcoin* en función de una serie de variables que captan factores económicos, tecnológicos y de atención por parte de los inversores. Nuestro modelo LSTM replica razonablemente bien el comportamiento del precio del *bitcoin* en diferentes períodos. A continuación, empleamos una técnica de interpretabilidad —SHAP— para determinar las características que influyen más en los resultados del modelo LSTM. Conforme a lo anterior, concluimos que la importancia de las diferentes variables cambia sustancialmente a lo largo del período analizado. Además, encontramos que no solo varía su influencia, sino que, paulatinamente, aparecen nuevos factores explicativos que, al menos en su mayor parte, permanecen desconocidos.

Palabras clave: *bitcoin*, aprendizaje automático, redes neuronales LSTM, técnicas de interpretabilidad.

Códigos JEL: C40, C45, G12, G15.

1. Introduction

Crypto-asset markets have been gaining increased attention from both the private and the public sector ever since its early inception in 2009. Despite the fact that growth has been uneven for most its existence, in recent years, developments seem to point in a different direction, thus featuring a consistent upward trend which brought about an expansion unheard of in previous periods. As such, market capitalization rose from merely 15 billion US dollars, in early 2017, to around 300 billion in 2020, right before the pandemic outbreak. It then skyrocketed until a peak of around 3 trillion US dollars was finally reached in November 2022. At the time of writing the market has once again bounced back at around 1,7 trillion US dollars in what yet again appears to be a steady downward race. Taking into account that the volume of bitcoin exhibits a constant growth rate, these fluctuations in capitalization are derived from the large variations in the price of bitcoin.

Notwithstanding the above, crypto-asset markets have in parallel experienced profound transformations, giving rise -among other things- to greater institutional and retail involvement. This was mainly driven by both an increased role of traditional financial institutions in certain segments and the deployment of more sophisticated investment products such as ETFs, futures contracts and other collective investment vehicles. The market has further spread to encompass other applications like Non-Fungible Tokens (NFTs). In addition, it has also supported the emergence of so-called decentralized finance (DeFi): a highly speculative niche that offers significant returns against equally great risks. As result crypto-asset markets are progressively becoming more intertwined with the formal financial and monetary system, thus amplifying their potential to spill their inherent vulnerabilities over to the economy at large.

In this context, volatility arises as a critical subject of study on which a large number of contributions exist to this date. Yet, irrespective of the many different types of available crypto-assets, the bulk of the literature focuses undoubtedly on the particular case of bitcoin. Although its total market share has dropped from 75% in 2017 to 50% in 2021, bitcoin continues to play a critical role to explain overall market trends and to trigger the development of a wealth of other initiatives either as a role model or example of problems/shortcomings that may have to be addressed to promote greater take-up. Bitcoin further echoes the fact that crypto-assets remain largely unregulated, thus it helps identify potential courses of action for regulators and supervisory authorities across jurisdictions. As a result, ascertaining the determinants of bitcoin price formation and assessing their stability over time can shed light and help steer ongoing discussions on the best way to approach increased direct and indirect exposures of critical financial market participants to crypto-assets more broadly. This knowledge will allow to establish the actual materiality of the underlying risks and consistently guide the decision on the proportionality of applicable requirements.

The objective of this article is, therefore, to analyse which variables influence the price at which bitcoin is traded on exchanges and how they behave across time. With this in mind, we were particularly mindful of two distinctive issues. First of all, the fact that unlike other financial instruments bitcoin lacks intrinsic value nor is it backed by a pool of assets like the so-called stablecoins. And secondly, there is no agreed theoretical model that explains ex-ante the determinants of the price of bitcoin. For this reason, we decided to use a flexible machine learning (ML) model, specifically, a long and short term memory (LSTM) neural network, in order to anticipate the price of bitcoin based on a series of potential explanatory variables. This model is consistent with our goals and the underlying circumstances surrounding bitcoin in that it allows for a flexible approach which does not impose ex-ante restrictions on the relationship between the various features¹ and the price of bitcoin. Furthermore, the model can also accommodate multiple features in non-linear and non-stationary time series. To choose the features to be used in the model, we drew from the existing body of papers on

¹ In machine learning, “features” is the term used for the individual independent variables that are taken as an input to make predictions over a target variable. Throughout the article we will use the term “features” together with “factors”, “drivers” and “variables” interchangeably.

the possible determinants of the price of bitcoin and picked a set of economic, technological and investor-related variables that seemed to enjoy the widest support.

Our aim is neither to build a perfect predictive model nor to develop a sound investment tool. In fact, we are not interested in analyzing how the price of bitcoin reacts to past prices, or to strongly endogenous variables (such as bitcoin's market capitalization). On the contrary, as stated before the purpose of this article is to understand if there is a set of factors that confidently explain the fluctuations in the price of bitcoin, and to analyze if these potential relationships are stable over time, all of this with a view to providing authorities with additional input in their open reflections on how to best deal with the emerging reality of crypto-assets.

With this in mind, our analysis takes place over three different periods. First, we address the time interval that falls between 2015 and mid-2017. This is the time frame where price growth seems to be the less pronounced and more stable. We then focus on the stretch comprised between 2017 and 2019, known to feature the first big price bubble in the history of bitcoin. Finally, we pay attention to the period between 2020 and 2021 with its large fluctuations and all-time price records. Our findings show that the LSTM model is able to replicate reasonably well the behaviour of the bitcoin price in all the above time segments. We then apply an ML interpretability technique called SHAP, to deepen our understanding of the features that are most important to the LSTM predictions. In this way we can identify the main drivers of the price of bitcoin at different points in time.

As a result of the above exercise, we learned that technological variables (such as hash difficulty, block size, number of transactions or unique addresses) played a decisive role in the first two periods, yet they became irrelevant in 2021. On the other hand, attention-related variables, such as searches on Google Trends, showed an increasing relevance over the years. In fact, they turned out to be the most important category in 2021. What's more, in stark contrast to part of the literature, we could not amass evidence that the SP500 and gold were ever among the main drivers of bitcoin price. We, therefore, conclude that the influence of specific variables in the price of bitcoin is largely unstable and seems to change substantially -and in ways hard to anticipate- across the different time periods observed. Not only does its influence vary, but our research suggests that new explanatory factors might also appear which oftentimes remain opaque, at least in its early stages. We believe these findings to be of relevance for that work that both local and global financial authorities are carrying out in order to help inform the design of prospective public policy actions.

We acknowledge the existence of both opportunities and challenges in choosing ML models over more traditional econometric techniques. On the one hand, as mentioned before, ML allows for a high degree of flexibility and better out-of-sample performance than traditional econometric techniques. On the other hand, the use of ML has a cost in terms of interpretability: i.e. traditional econometric alternatives are inherently interpretable, while ML models must resort to additional tools for this very purpose. Despite this trade-off, we believe that post hoc interpretability techniques, particularly SHAP, work reasonably well (Alonso and Carbo, 2022; Molnar, 2022), while also becoming more widely used in the context of ML (Albanesi and Vamossy, 2019; Chen et al. 2021). Therefore, on account of the particular circumstances of our study, we consider that ML advantages outweigh its potential limitations as it works better in predicting and replicating the data and interpretability shortcomings can be addressed in a practical way.

The paper is divided as follows. Section 2 explains the main features and recent evolution of the crypto-asset market, highlighting key public policy issues. Section 3 presents a literature review and underlines our contribution. Section 4 sheds light on the data that were used for this study. Section 5 is devoted to the empirical analysis, which is split into (i) predictions by the LSTM model (sections 5.1 and 5.2) and (ii) interpretation of those predictions and the determinants of the price of bitcoin (sections 5.3 and 5.4). Section 6 concludes.

2. Crypto-asset markets: notable characteristics, latest developments and key public policy issues

Crypto-assets, more frequently and imprecisely known as crypto- or virtual currencies, are a new type of private asset that depends primarily on cryptography and DLT (or an equivalent technology) as part of its perceived or inherent value (FSB, 2018). As such, crypto-assets embrace a diverse set of representations whose main common trait is, precisely, the fact that they usually combine several distinctive elements from a wide range of financial instruments. As a result, crypto-assets typically emerge as hybrid products that rise both significant conceptual and practical challenges in trying to accommodate them into any pre-existing legal framework² (Foz, 2021)

Despite the potentially manifold interpretations that may stem from the above constraints -both in relation to their nature and effective scope-, from a technical standpoint crypto-assets seem to exhibit a number of prominent and prevailing attributes at all times: (i) the deployment of a decentralized ledger that enables new ways of issuance, registration and exchange of underlying assets, (ii) the application of protocols which set out the rules by which the execution of transactions is to take place, and (iii) the emergence of a complex ecosystem featuring a wealth of participants that take up different roles as regards the distribution, validation, trading, transfer and storing of digital assets.

It is now commonplace to consider that the inception of the crypto-asset markets was a by-product of bitcoin's outbreak. However, its theoretical foundations can be traced back to the works of earlier authors like Wei Dai (1998) and Nicholas Szabo (2005) among many others. These researchers were keen on highlighting the critical part played by trusted intermediaries (e.g. central banks and/or credit institutions) in securing the proper functioning of traditional payment circuits and yet, how this very circumstance was prone to triggering other types of negative consequences for the society as a whole³.

In an attempt to address these challenges, both authors advocated for the establishment of an alternative monetary system based on distributed networks and cryptographic proofs. They further supported the deployment of tools that could help ensure a more predictable growth of the underlying settlement asset. However, it was indeed the actual launch of bitcoin that first provided a material example of a practical type of money holding the promise of avoiding an indigenous loss of its value over the long term (Nakamoto, 2008).

Thus, bitcoin's ultimate purpose was that of becoming a credible alternative to fiat money for which it leverages a novel electronic payment system that supports peer-to-peer exchanges of value in an allegedly safe manner. According to its proponents, such particular circumstance further facilitates the deployment of affordable and irreversible payments since transaction costs are lowered on account of forgoing the role traditionally assigned to intermediaries in all regular payment channels but cash (Conesa, 2019).

In order to achieve this goal, bitcoin hinges on a network of users -nodes- which rely on the Internet to communicate with one another on the basis of a common protocol by masking their respective real-world identities under public keys. Said protocol is widely accessible to any interested parties in its condition as open software.

² Regardless of the growing number of initiatives that intend to provide a specific regulatory framework for crypto-assets, oftentimes where not yet in place other approaches apply. Thus, some jurisdictions are trying to address this regulatory gap by extending to certain aspects of crypto-assets' operations currently applicable requirements on account of the similarities with other types of regulated activities; e.g. payment services or financial instruments.

³ According to these authors, salient examples of potential negative implications of intermediation include hyperinflation as well as the potential for authorities to tax or regulate social and economic activity through the threat of force.

In the interest of avoiding the so-called double spending problem, i.e. the possibility that the same units of crypto-assets are sent to more than one recipient at once, bitcoin marshals an inspired solution: each and every network participant enjoys full visibility over all the transactions that have taken place across them. Moreover, they all are entitled to act as potential validators. Hence, every prospective operation ends up being published in a decentralized repository and awaits final validation with the concurrence of any number of parties that decide to join the respective process. The various nodes will also serve as a common database, keeping a complete or partial copy of the historic record of all executed transactions at any given point in time.

With a view to providing a consistent and unified picture of said repository and so, ensuring that it reflects the legitimate consent, bitcoin has chosen to rely on a computationally costly verification procedure. Nodes agreeing to voluntarily perform the role of validators (miners) will, therefore, have an option to be compensated for their efforts both with newly created units of crypto-assets⁴ and by cashing in user fees. Yet, rewards will only be then granted to the node that first demonstrates to have rightfully solved the cryptographic puzzle posed by the algorithm.

Completing this task successfully demands a great amount of resources as it is basically a trial and error process that aims at brute-forcing a given result. Thus, in principle, there is some randomness over which node will finally achieve this accomplishment⁵. Transactions pending validation are collected in a block which is then trimmed down to form a 256-bit block hash value. Once a block has been mined correctly it will reference the previous one, forming an unbroken cryptographic chain back to the first bitcoin block (Brühl, 2017). This setup helps warrant that ultimately everyone agrees on the transaction record. It furthermore makes it highly unlikely that anybody could tamper with blocks in the chain since re-mining all the following blocks would be computationally unfeasible.

Bitcoin's positive reception paved the way for an expansion of the crypto-asset's ecosystem which led to the emergence of a wide range of initiatives, each one with different goals and purposes. Hence, in addition to providing new means of exchange, some novel crypto-assets have proven to be highly useful in raising funds for innovative and risky projects. Likewise, others have been shaped in such a manner that they resemble an equity in the stock of a company, further furnishing their holders with voting, subscription or appraisal rights as well as with dividends and other entitlements. Eventually, certain token's main functionality is to grant their users access to products and/or services developed and distributed by its issuer over its own technical platform (utility tokens).

From the point of view of their respective users, crypto-assets are oftentimes leveraged to cater for payments, hedging or speculative purposes in a manner that is somewhat disconnect from its original intent. Along the same lines, some tokens are mainly conceived as a symbolic gesture or statement such as Dodgecoin or Jesus Coins (Kim et al., 2018). In addition, so-called non-fungible tokens (NFT) have been gaining momentum lately in that they have proven particularly useful to attest the scarcity and provenance of rare assets –e.g. original artwork, pictures, collectibles, or trademarks, among many others-. This circumstance is creating new and significant opportunities to trade unique digital goods⁶.

⁴ This amount is cut in half at specific points in time. It is the result of a scheduled event, known as halving, whose goal is to ensure that the maximum supply of bitcoins does not exceed a fixed referenced volume (21 million) in order to mimic the finite quantity of physical gold. By design, the number of bitcoins minted per block is reduced by 50% after every 210.000 blocks which is equivalent to about once every four years. This helps keep the pace of growth predictable. Halving seems to have some interesting knock-on effects on its market price and the behaviour of other markets (El Madhy, 2021; Meynkhard, 2019).

⁵ Yet, miners with most efficient hardware may potentially benefit from better hash rates which is likely to increase their probability of successfully mining a block and so, capture the reward.

⁶ NFTs are tradable rights to digital assets, where ownership is recorded in smart contracts on a blockchain. Due to its lack of fungibility they are mainly intended as pure assets (Dowling, 2022).

The compounded effect of the above developments alongside the consolidation of its underlying technology proposition as well as the expansion of decentralized finance (DeFi) have all caused a visible surge of crypto-asset markets. The additional impact of the recent lockdown due to the COVID-19 outbreak has significantly amplified this trend. As a result, the market has furthered its size and reached levels previously unheard of. In fact, market capitalization of crypto-assets exceeded USD3 trillion in 2021, a 400% growth in less than 12 months that can be largely explained by bitcoin and ether⁷. Moreover, both these crypto-assets ranked among the world's top 20 traded assets (Iyer, 2022).

In terms of price, bitcoin stood out as a particularly volatile crypto-asset in the period 2020/2021. At its peak, it experienced an almost 200% increase in comparison to its point of departure in January 2021. This was also about five times greater than its previous record. Other crypto-assets such as ether or binance coin followed suit, yet their rate of growth was far less pronounced. In parallel, stablecoins also saw a significant expansion, being the most traded asset class within the crypto-asset ecosystem. Their current market share stands at around 7%.

The latest episode in the expansion of crypto-asset markets shows a number of distinctive characteristics that help tell it apart from any other previous phases. Despite the substantial role played by retail investors in driving up both volumes and values, institutional investors –including hedge funds- have been particularly active in this field⁸. Consequently, both the average size of trades and the maturity of associated portfolios has increased accordingly. Similarly, the demand for crypto-assets has soared among large IT corporations that now consider them as a useful tool for the treasury management strategies. Despite still being of a relatively low importance in terms of size, these movements seem to have had a noteworthy signaling effect on other potential investors, eventually giving rise to a bubble (Shu et al., 2021).

This period has also seen a rise in the number of crypto-asset related custody and trading services that are being offered by traditional financial institutions while also featuring their mounting integration in well-established and widely-used private payment networks such as those commonly associated with cards. Again, this latter aspect is of great relevance for the future prospects of crypto-assets in that they may find it easier to enter and consolidate in the broader financial system at a faster pace.

The transformation of the crypto-asset markets is further augmented by the deployment of more sophisticated and complex investment vehicles such as, e.g. derivatives, futures or ETFs, which can help stimulate the appetite of a wealth of investors that still lay off boundary due to the uncertainties and limitations surrounding these markets.

From the point of view of financial authorities, all these changes signal the true potential of crypto-assets to become a critical element of both the financial and economic blood circuit of the society at large. They also highlight the sheer size of challenges that they need to face speedily in order to safeguard the orderly functioning of both the system as a whole and of its underlying parts. A larger footprint of crypto-assets is, thus, seen as a potential source of distortions for both financial stability and monetary policy (ECB, 2019) while also giving rise to other relevant concerns such as the effectiveness of consumer protection mechanisms (Australian Parliament, 2021), threats to financial integrity (FATF, 2014) as well as their foreseeable negative environmental impact (Moshin et al., 2020).

In this context, the distinctively high volatility of crypto-assets alongside an increased reliance on leverage trading strategies (IMF, 2021) emerge as the two chief elements which largely fuel international discussions.

⁷ Together, bitcoin and ether account for about 2/3 of the market's capitalization value.

⁸ PwC & Elwood estimate that, between 2019 and 2020, the size of crypto-asset related AuM held by hedge funds doubled, reaching a scale of USD 2.000 million. Other research seems to confirm the growing relevance of institutional investors for the bitcoin market as well: e.g. according to data by Chainanalysis, in December 2021, bitcoin-related transactions below USD 10.000 fall by 22% whereas those with a nominal amount above USD 1 million increased by 32% in the same period.

Aspects such as their ultimate influence on market, credit, liquidity risks and how well are these being managed are, therefore, core issues for authorities to carefully look into⁹. As a result, they are growingly expected to design and deploy the most appropriate safeguards in defense of the public interest while still succeeding in keeping alive most of the promised benefits that the development of these markets could bring about.

For this purpose, authorities are presently engaged in a comprehensive and globally coordinated exercise to review thoroughly applicable regulatory and supervisory frameworks so as to decide when and how to adapt current rules and standards (e.g. Basel III) and when to complement those with novel ones (e.g. MiCA). This approach implies combining traditional measures -including capital, liquidity and leverage requirement- with other types of actions such as an enhanced monitoring of market behavior or promoting improved transparency, yet their proper calibration demands to deepen their understanding over which elements and forces most likely shape crypto-asset markets' behaviour.

One such fundamental aspect is the determination of triggers of price as well as the assessment of their stability over time. This knowledge will be crucial to gauge the actual materiality of the underlying risks and consistently decide on the proportionality of applicable requirements. In the next sections, we try to address some of these questions by applying a set of novel tools.

3. Literature review

Bitcoin has been of interest to both the general public and the academia since its inception, when Nakamoto wrote his seminal white paper (2008) and launched the cryptocurrency in January 2009. Due to this growing attention, there are several branches of academic literature dealing with topics that concern bitcoin. On the one hand, there is a debate on whether bitcoin could be considered as a safe haven asset, a hedge, a diversifier, or just a speculative asset. For more on this, see Bouoiyour and Selmi (2015), Bouri et al (2017), Urquhart and Zhang (2019) and Bariviera and Merediz (2021). On the other hand, there is a wealth of papers that focus on examining what variables could influence the price and return of bitcoin. This particular topic is seeing a substantial increase in the number of actual contributions which, oftentimes, exhibit very heterogeneous results. Our work is framed within this broad universe.

Bitcoin has been active since 2009 which means that existing articles address different time periods and employ very different methodologies. While some of them rely on traditional time series like ARIMA or Vector Error Corrections, others choose machine learning methods instead, like boosting or neural networks. In order to help better understand our contribution to the literature, we resolved to structure and cluster the former into three distinctive types of debates: (i) the role of economic/financial variables such as exchange rates or economic indicators, (ii) the relevance of technology and its many blockchain-specific underlying features, and (iii) the role of gold and the degree of attention/interest raised by crypto-assets among investors and the general public.

The first line of thought explores whether macroeconomic indicators or the most relevant exchange rates could explain the price of bitcoin. Using ARDL and Bayesian quantile regression, Bouoiyour and Selmi (2015; 2017) found that the Shanghai market and macroeconomics events -like Brexit or Venezuela currency demonetization- indeed exercised visible influence. Cermak (2017) leveraged GARCH models to conclude that macroeconomic conditions in China, USA, and the EU could affect bitcoin's price volatility. Moreover,

⁹ Notwithstanding other equally relevant aspects such as those related to increased operational risks due to the maturity of the technology as well as associated legal and reputational risks.

using Vector Error Correction, Zhu *et al.* (2017) established that the US Dollar Index, the Dow Jones, and the Federal funds had a long term impact on the price of bitcoin.

Over and above, the critical role of the SP500 is further emphasized in three additional papers. Sovbetov (2018) resorted to an ARDL in order to infer that the SP500 had a weak and positive effect in the long run. Along these lines, Kjarland *et al.* (2018), discovered that the SP500 was among the most important variables determining the price of bitcoin. In addition, Kapar and Olmo (2020) used VECM to understand the dynamics of bitcoin returns and concurred as well with the fact that the SP500 had positive correlation with bitcoin returns. The potential relevance of other economic factors, like return to investment, was explored by Alessandretti *et al.* (2021) via Boosting and LSTM who found it to be the most crucial factor determining the price of bitcoin. Finally, Chen *et al.* (2021) combined LSTM and random forest and found that both the SP500 and the NASDAQ, have a positive influence on bitcoin price.

Conversely, within this literature, there are also papers that question the importance of some of the aforementioned factors. For instance, leveraging ARDL and VECM, Ciaian *et al.* (2016) and Cian *et al.* (2018) posited that macroeconomic and financial variables do not have a long-term impact on the price of bitcoin. Similarly, Baur (2018) ascertained that the price of bitcoin is uncorrelated with traditional asset classes such as stocks, bonds and commodities both in normal times and in periods of financial turmoil. Finally, Pyo and Lee (2019), found that price changes in bitcoin after macroeconomic announcement are insignificant.

On the point of technology, the role of the hash rate and the difficulty of the algorithm (see appendix for definitions) seem to be overly present in the literature. Some papers find evidence that the competition in the network of producers, the rate of unit production, and the difficulty of the algorithm used to “mine” the “cryptocurrency” could be influencing the price of bitcoin -see, e.g. Kristoufek (2015), Ciaian *et al.* (2016), Bouoiyour and Selmi (2017) or Hayes (2018)-. Other research considers that the role of hash rate or hash difficulty is negligible for determining said price. For example, Fantazzini and Kolodin (2020) found that there is neither evidence of Granger-causality or cointegration between hash rate and the price of bitcoin. Kjarland *et al.* (2018) established that the hash rate is irrelevant for modeling bitcoin price dynamics.

Finally, there seems to be wide consensus on the positive impact of gold, investor attention and online searches in the evolution of the price/return of bitcoin. Bouoiyour and Selmi, (2017), Panagiotidis, (2018) as well as Kapar and Olmo, (2020) all highlighted the importance of gold. Kristoufek (2013), Kaminski (2014), Kristoufek (2015), Ciaian (2016), Kim *et al.* (2016), Sovbetov (2018), Panagiotidis (2018), Lyócsa *et al.* (2020), Chen *et al.* (2020) and Zhu *et al.* (2021) drew similar conclusions about the role on investor attention. For this purpose, they measured it either as the number of tweets, the number of online searches on Google, or the number of enquires in Wikipedia.

Some of our findings are in line with this literature. As we show in section 4.4, investor attention, proxied by the volume of both searches in Google Trends and mentions in Twitter, is one of the most important variables influencing the price of bitcoin. Regarding the hash rate, algorithm difficulty and other technology-related variables inherent to bitcoin, we gather that they exhibited a strong influence on the price of bitcoin both in 2017 and 2018, but not any longer in 2021. Finally, concerning gold and macroeconomic indicators like the SP500, we conclude that they never rank among the top factors that explain the evolution of the price of bitcoin.

Our paper is mostly related to Panagiotidis *et al.* (2018), and Chen *et al.* (2021). The former is entirely focused on the drivers of bitcoin rather than in predicting, while the later also uses LSTM and interpretability techniques to uncover the determinants of the price of bitcoin. Our contribution is twofold. We first identify which are the potential explanatory elements in the evolution of the price of bitcoin at different points of its lifecycle, including the most recent spike of 2021. In this way, we are able to provide evidence of whether these different variables remain the same or change and are replaced by others. We further distinguish whether other, still unknown factors at the time of writing, emerge which are not captured by our model.

In addition, contrary to the approach followed by Chen et al. (2021), we address the issue of interpretability of the results in a different way. Instead of leveraging feature importance to uncover the relevance of each of the underlying variables to the LSTM method, we rather rely on SHAP, a global interpretability technique, that might be better suited when using correlated features and when trying to interpret neural networks (Alonso and Carbó, 2022). Moreover, unlike Chen et al. (2021), we try to avoid highly endogenous variables among the potential determinants, like lagged values of bitcoin price or market capitalization.

4. Data

In accordance with most the existing literature, we consider three different types of potential explanatory factors linked to the price of bitcoin: (i) the specific technology features of bitcoin, (ii) the evolution of the economy, and (iii) the level of attention/interest it arises among the public at large. Starting with the technology dimension we took into account the following elements: difficulty in finding the hash, unique addresses, commissions to miners (fees), hash rate, sum of blocks, average block size, sum of transfers, and average transfer size¹⁰.

Regarding the economic variables, we chose to include the following ones: the price of gold and oil (separately), the SP500, the FTSE, the DOW30, the NASDAQ, and the exchange rate of several international currencies (i.e. the Euro, the British Pound, the Yuan, the Yen and the Swiss Franc) and the US dollar. Finally, as a proxy for the level of public attention we placed our focus on (i) how the search term “bitcoin” was captured in Google Trends, and (ii) the number of Tweets per day that were published with “bitcoin” as the distinctive hashtag¹¹¹². We further considered using specific term searches in Wikipedia but we discarded this approach soon as it did not enhance the information already provided by Google Trends and Twitter.

While agreeing that a larger selection of Google trends’ keywords / different Twitter hashtags may help better capture the public's degree of attention or interest in bitcoin, we feel that our approach already partly addresses this concern. In fact, when searching for “bitcoin” on Google trends, we performed this search “as a topic”. This implies that Google’s search engine seeks result worldwide, regardless of underlying language. Furthermore, it takes into account other searches that are considered to be related to the “topic” of bitcoin according to Google trends (Carbó and García, 2021). In any case, to more holistically reflect the actual interest of investors on bitcoin, more specialized searches could be performed the choice of the option “related topics or related searches” of Google Trends or, alternatively, using different keywords (cryptocurrency, bitcoin-usd, etc.) as in Aslanidis et al (2022) and Urquhart (2018).

All this data was collected from Coinmetrics (technological factors), Yahoo Finance (economic factors), Google (Google Trends), and Bitinfocharts (Number of Tweets). In total, we relied on 25 features to determine the price of bitcoin.

We worked with daily frequencies in our set of data (business days only) from January 2015 to July 2021. We further divided the sample into three periods of time. We called the first slot launch period (from 1 April 2015 to 1 April 2017) It distinctively shows a steady growth in the price of bitcoin. The second slot was referred to as the expansion period (1 April 2016 to 1 April 2018). It features the first spike in the price of bitcoin, namely in December 2017, when it topped 20.000 US dollars. The third and last slot goes from 15 June 2019 to 15

¹⁰ We include in the appendix a section in which we explain in detail each of these possible determinants

¹¹ Due to the way in which we have obtained the data, we do not know if the intensity of searches on Google and the tweets have been generated by a specialized audience (investors) or a more general one. That is why from now on we will refer to these variables broadly as public attention variables.

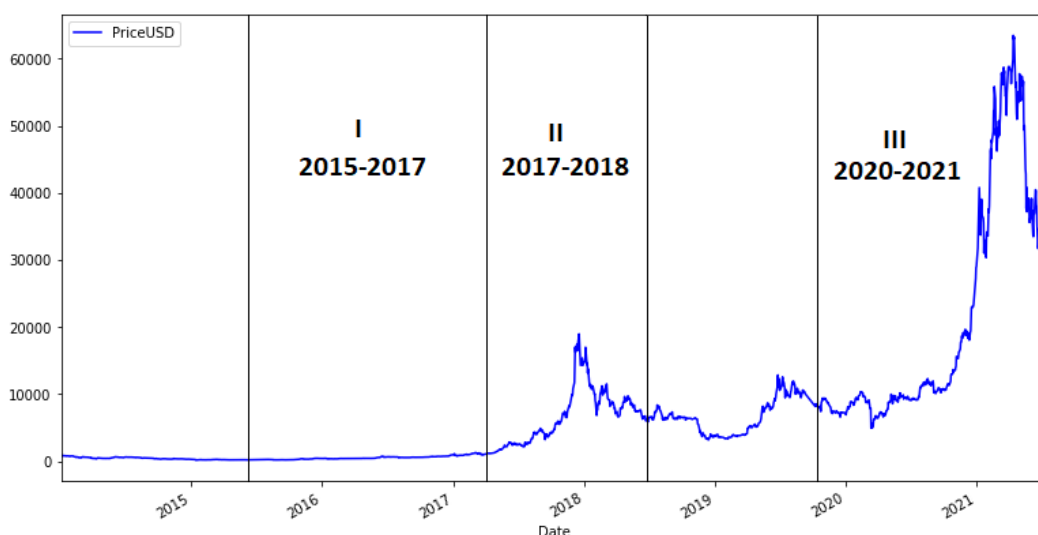
¹² The information from Twitter corresponds to any tweet that has the hashtag “bitcoin”, so it can be from anywhere in the world in any language. The information from Google Trends is also collected globally and in any language

June 2021. This was branded as the *consolidation period* and runs from early post-pandemic days until reaching the heights of the price of bitcoin in March and April 2021 (i.e. 60.000 US dollars).

Figure 1 and table 1 depict the three periods mentioned above and further present information about their precise duration. In our view, each of them is to be associated with key moments in the lifecycle of bitcoin. For each one of them, we trained, validated and tested our LSTM model, determining the target variable (i.e. bitcoin exchange price as expressed in US dollars) as a function of the 25 features considered. We then complemented this exercise by analyzing in detail which features were the most important ones to explain price movements that took place during the various individual time slots.

For each period, the exact start and end dates were chosen at our discretion so as to ensure that the test samples had a minimum duration of five months and reflected the key events that interested us: i.e. the constant growth at the end of 2016, the first peak in the price of bitcoin at the end of 2017, and its all-time highs during the first months of 2021. In addition, when choosing the dates, we made sure to count at least 12 months of training, and between one and two months of validation prior to the test samples. This way, in all the three periods we secured a 70% partition of the data to train, 5% to validate, and 25% to test. We also checked the robustness of our results by selecting different start and end dates in the interval contained four weeks before and four weeks after the original start and end dates.

Figure 1: Bitcoin price evolution and the three stages considered



Source: Author's own calculations (2022).

5. Empirical exercise

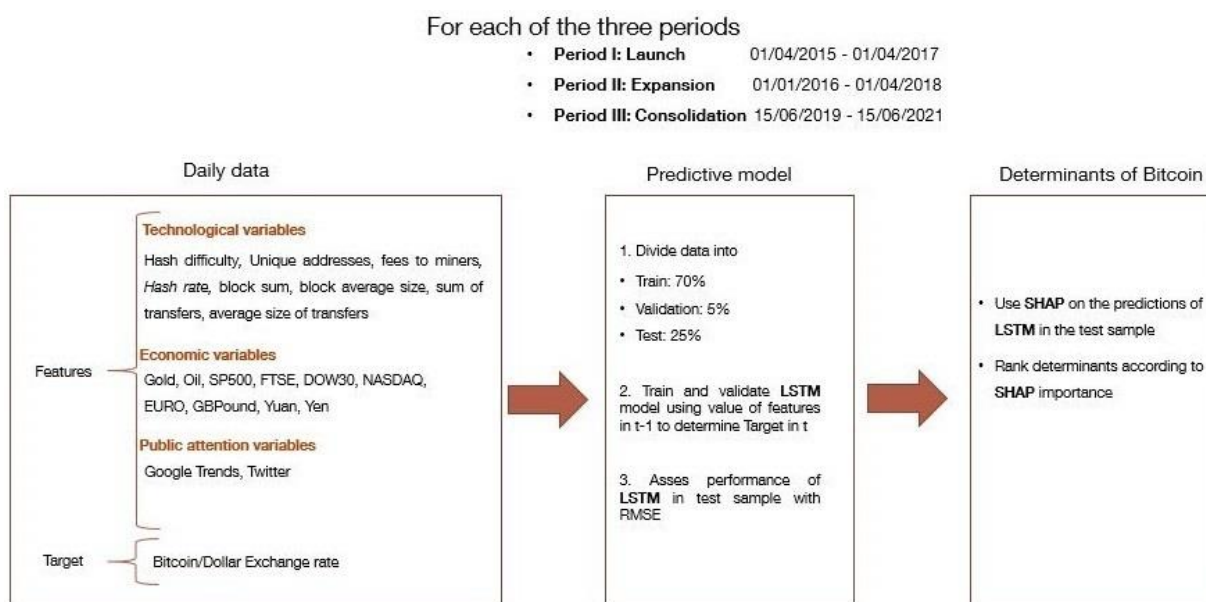
The empirical exercise proposed in this paper is summarized in **figure 2**. For each of the different time periods we chose to analyze, data on all the potentially explanatory factors -as mentioned in the previous section- was collected and further classified into three distinctive groups. The next step was to use the values of those features in $t-1$ to train, validate and, finally, test our LSTM in order to see whether we could determine the price of bitcoin in t . We then checked the resulting prediction error by calculating the root mean squared error metric. This allowed us to corroborate that the LSTM model was capable of satisfactorily replicating the price of bitcoin.

As explained before, our goal was not to build a model that could be used to flawlessly predict the future price of bitcoin. That is, by nature, an almost impossible task. What we pursued instead, was to develop a

model that, without imposing or assuming any specific ex-ante relationship between the various features and the target variable, could roughly replicate the movements of the price of bitcoin. Finally, using SHAP, a ML interpretability technique, we checked the importance of the individual features in order to understand which ones might be the most important to determine the price of bitcoin in each period.

Below you will find all the details about our empirical study. In section 5.1 we explain the basics of LSTM model and why we used it. In section 5.2 we show the performance of LSTM based on its out-of-sample price prediction. In section 5.3 we introduce SHAP, the technique that we used to understand the outcome of the LSTM model. In section 5.4 we apply SHAP to uncover the determinants of the price of bitcoin in each of the analyzed periods. Finally, in the appendix we show how the performance of LSTM could be improved if we included the lag of the price of bitcoin as an additional variable. But this would complicate the analysis of the determinants of bitcoin, since interpretability techniques would attribute great importance to the lag in the price of bitcoin, and less to the rest of the variables. That is why we prefer not to include the lag of the price of bitcoin as an explanatory variable, obtaining a slightly worse prediction but with an interpretability of the results that allows us to better understand the price determinants.

Figure 2: Empirical exercise proposed in this study



Source: Author's own calculations (2022).

5.1 LSTM model

The problem at hand, establishing the price of bitcoin as a function of the values adopted in the past by a set of explanatory variables, is a challenging task. Most of the data that were collected for this project are highly nonlinear and non-stationary. Therefore, the use of traditional time series will require imposing additional assumptions on the data. Since there's no generally-accepted theory that could be leveraged for this purpose, on account of the fact that our ultimate aim was to understand the extent to which certain factor movements were compatible with price movements, we chose to address this problem by relying on a Long Short Term Memory neural network (LSTM)¹³. The advantage of this approach is the flexibility that the model provides in

¹³ The choice of this ML model was the result of a selection process to find the best performing one. In fact, priori to settling on the LSMT, we first tested different tree-based models, such as regression trees, random forest and XGBoost. In each of the periods, we

that it does not impose any specific ex-ante relationship between the target variable (bitcoin price) and the features. Thus, it can be used in non-linear and non-stationary time series scenarios (Bala and Singh 2019, Abbasimehr et al 2020 and Wang et al 2021). The model can work well in the presence of multicollinear variables because it has a nonlinear structure (via activation functions) and also uses regularization techniques. Regularization allows to stabilize the coefficients assigned to the features, thus reducing the problem in the case that two features are multicollinear.

LSTM is a variation of feedforward neural networks which are capable of learning the time dimension of the data. We show a graphical representation of a simple fed forward neural network in **figure 3**, left. A feedforward neural network consists of an input layer, which corresponds to the input data, one or more hidden layers, and an output layer. The input layer has as many nodes as features in the data (four in the case of **figure 3**). The hidden layer is composed by nodes that represent linear combinations of weights of the nodes from the input layer (in this case, three nodes). These combinations are processed by activations functions (tangent, sigmoid, etc.) that result from the output of each hidden layer node. If there is more than one hidden layer (in **figure 3** there is only one hidden layer), the nodes of any hidden layer are based on the linear combinations of the output of the nodes of the previous hidden layer. Finally, there is an output layer based on a linear combination of nodes of the last hidden layer, to which we apply a nonlinear activation function.

The drawback of this architecture is that it does not take advantage of the information available from past values of the features or decisions of the net. Consequently, it might not be an ideal tool for performing time series analysis. However, this is partially solved by using Recurrent Neural Networks (RNN), which are distinguished from regular neural networks in that they can deal with sequential data and can be trained to hold the knowledge about the past. This is achieved by applying a feedback loop that is connected to their past decisions, thus ingesting their own outputs moment after moment as new input (**figure 3**, right). We explain in more detail how this feedback loop works for RNN in **figure 4** left.

Let's focus on the RNN cell at time t . The previous hidden state h_{t-1} and the input x_t are combined in a vector that will go through a tahn activation function, resulting in the output of the cell h_t , which represents the memory for the next time step. The tahn activation function ensures that the values of the vector are always between -1 and 1. This architecture allows the RNN to make use of sequential information, in a similar fashion to a Markov model. But long term information might not be used, since RNN are exposed to the vanishing gradient problem. Neural networks, including RNNs, update parameters based on an optimization algorithm called gradient descent, in which models learn via gradient values. If the gradient values are extremely small, the parameter updates become negligible, and as a result, the model stops learning or takes too long to learn.

LSTM are a special version of RNN that can solve this problem. They were first proposed in 1997 by Hochreiter and Schmidhuber. The basic architecture of a LSTM cell is shown in **figure 4** right. The main difference from RNN is the way the model controls what information should be removed from memory, and what information should remain. This is done using information outside the normal flow of the network in a closed cell or memory, called cell state. This cell state can be seen as the horizontal line in the top of the cells, from h_{t-1} to h_t and to h_{t+1} . The LSTM decides whether to store or delete information from the cell state based on the importance it places on the information. There are three gates that can transform the cell state: forget gate, input gate, and output gate. These doors are neural networks that determine whether past information is

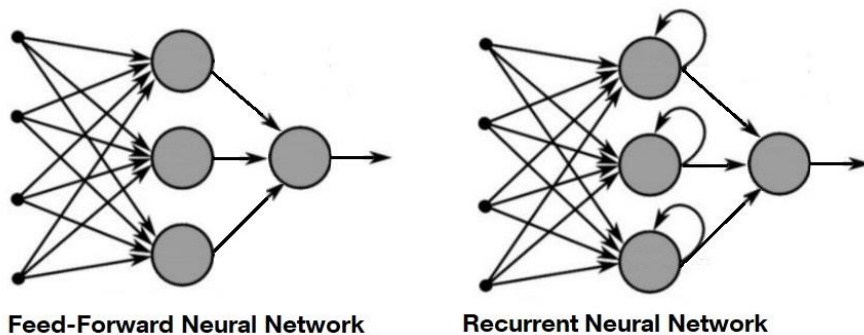
found the prediction error of these models to be considerably higher than for the LSTM, with a root mean square error (RMSE) as high as twice the error borne by the LSTM. We tested further traditional deep learning models (without the characteristic time loop of recurrent neural networks or LSTMs), and although their performance proved to be better than with tree-based models, the prediction error was still at least 10% higher than compared with the LSTM in all the three periods considered.

eliminated (forget gate), whether or not new information is allowed (input gate), and what information is moved to the next state (output gate).

The forget gate combines the previous hidden state h_{t-1} , and the current input x_t through a sigmoid function, delivering a value between zero and one, where zero means non important and one means important. This value is then applied through pointwise multiplication to the cell state, in this way the cell state can forget unimportant information. The input gate combines as well the previous hidden state and the current input through a sigmoid function and a tahn function, and multiplies both values to regulate which information will be added through pointwise addition to the cell state. This updates the cell state according to the values that the input gate considers relevant. Finally, the output gate also combines the previous hidden state and the current input through a sigmoid function, which multiplies the tahn output of the modified cell state. This results in the new hidden state for the next time step.


While the architecture and composition of the basic cell of a LSTM is more complicated than the one for RNN, the LSTM allows us to remember inputs for a long period of time, and it solves the vanishing gradient problem. This is the case since the gate structure of the LSTM ensures that the gradients do not converge to zero, keeping a relatively short training time.

Figure 3: Feed forward neural network and Recurrent neural networks (RNN) model



Source: Own authors (2022)

Figure 4: Recurrent neural networks (RNN) and LSTM model



Source: Olah, 2015.

As with any other ML model, we must avoid the problem of overfitting to produce a model that generalizes well to new and unknown data. If we train the LSTM model on the entire data sample, it will end up learning

the optimal weights of each feature and fit the model's predictions perfectly for the in-sample data, but it will not be able to predict well out-of-sample. That is why it is common practice in ML to split data into three groups, train, validation, and test. We fit the model in the training sample. In our case, the target variable was the price of bitcoin in t , while as features we used the 25 listed elements in section 3 with their values in $t-1$. In the validation set we evaluated the performance of the fitted model to select the hyper parameters and architecture of the model. Finally, we checked the accuracy of our model by making predictions in the test sample. In this way, we were able to assess the performance of the model against data sets which had not been directly involved in the training process. We followed this approach for each of the three periods of observation in which we chose to break up the lifecycle of bitcoin, dividing the sample as follows: 70% was used to train, 5% was used to validate (so that we had at least 30 days of validation) and 25% was used to test. In table 1 we show the exact days used for the training, validation and test samples. In **table 2** we summarize the parameters we selected through the validation process.

Table 1. Time intervals for the three periods

Period	Training sample	Validation sample	Test sample
Launch	01/01/2015 to 04/07/2016	04/07/2016 to 15/08/2016	15/08/2016 to 01/04/2017
Expansion	01/01/2016 to 25/07/2017	25/07/2017 to 22/09/2017	22/09/2017 to 01/04/2018
Consolidation	15/06/2019 to 22/11/2020	22/11/2020 to 23/12/2020	23/12/2020 to 15/06/2021

Source: Own authors (2022)

Table 2. Parameters

Parameters	Values
Input lags	3-5 days of memory
Activation function	Linear
Learning rate	0,001
Nodes in hidden layer	256, 128, 64 (three hidden layers)
Dropout value	0,1
Longest training duration	60 Seconds
Loss function	MSE

Source: Own authors (2022)

Table 3. Performance of the model in different periods

Period	RMSE	Test sample
Launch	5.7%	01/10/2016 to 01/04/2017
Consolidation	13.2%	01/10/2017 to 01/04/2018
Expansion	21.2%	29/12/2021 to 15/06/2021

Source: Own authors (2022)

5.2 LSTM results

We measure the effectiveness of the LSTM model through the metric RMSE (Root Mean Squared Error), which is based on the mean of the square root of the squared differences between the predicted values and the observed values. The formula for RMSE is as follows:

$$\text{RMSE} = \sqrt{\frac{(\sum_{t=1}^T \hat{y}_t - y_t)^2}{T}}$$

Where T is the number of days, \hat{y}_t is the predicted price of bitcoin on day t , and y_t is the actual price of bitcoin on day t .

Our first finding is that the LSTM model performs reasonably well in all three periods considered. This is a particularly positive outcome in order to cement the results of our exercise since we didn't use lagged values of the price of bitcoin as additional features. Moreover, the RMSE that we obtained was between 5% and 20% of the price (**table 3**).

Another interesting observation is that the model offers its best outcome in the so-called *launch period* (5.7% of RMSE), followed by the *expansion period* (13.2% of RMSE). The error in the prediction is higher in the *consolidation period*, 21.2%, particularly during March and April 2021.

These results could be different if we changed the exact start and end days of the period. As we discussed in section 4, the start and end days of the period were chosen so that the test samples were associated with key moments in the lifecycle of bitcoin. As a robustness exercise, we tested the performance of the LSTM model by slightly altering the exact dates of each of the three periods. Specifically, we carried out our exercise on different start and end dates, choosing dates between four weeks before and four weeks after the original start and end dates.

For both the *launch period* and the *expansion period*, we observed that results do not change significantly, with the RMSE varying between 4% and 6% for the *launch period*, and between 12% and 18% for the *expansion period*. Along these lines, the RMSE of the LSTM model ranged between 20% and 26% during

the *expansion* period when advancing two weeks and moving four weeks backwards both the start date and the end date. Interestingly, when pushing forward the start and end dates more than two weeks, the LSTM performed worse, with an RMSE that could reach 38%. It is worth mentioning that this, could change our conclusions around interpretability. This inferior prediction performance is explained by the fact the LSTM model cannot predict the big price increase that took place in 2021 when the training sample leaves out a critical amount of data points reflecting the price increase trend of bitcoin (late 2020). We recognize that this is a limitation of the study, and we leave for future research to delve into possible solutions to this problem.

The actual performance results of the LSTM model for each time slot are depicted in **figures 5, 6 and 7**, where the blue line represents the actual price of bitcoin, and the orange line represents the prediction made by the LSTM model.

Going further into the details, as regards the *launch period* (**figure 5**), we trained the model from 1 January 2015 to 4 August 2016, we further validated it from 4 August 2016 to 15 September 2016 and, finally, we had it tested for the following six months. As a result, we noticed that the predicted price of LSTM followed very closely the actual price of bitcoin. The same holds true for the *expansion period*, **figure 6**, in which we trained the model from 1 January 2016 to 25 August 2017, validated it from 25 August 2017 to 22 October 2017, and had it, again, tested for the following six months. While the difference between the predicted and the actual price is higher in this latter case, given the steep price during the test sample with respect to the training and validations sample, we believe it is fair to say that the performance of LSTM remains reasonably good.

Figure 7 shows the resulting prediction values for the *consolidation period*. Here, the model was trained from 1 June 2019 to 19 November 2020. The period used for validation ranged from 19 November 2020 to 29 December 2020, while testing took place until 15 June 2021. In this period, we did appreciate a considerable gap between predicted and actual price of bitcoin during March and April of 2021. On account of the above, we can first conclude that, based on the same features, the LSTM model's performance is generally good in terms of RMSE for the first two periods considered but it worsens considerably during the third one. These results could, however, be improved if the lag price value of bitcoin is included as an extra feature.

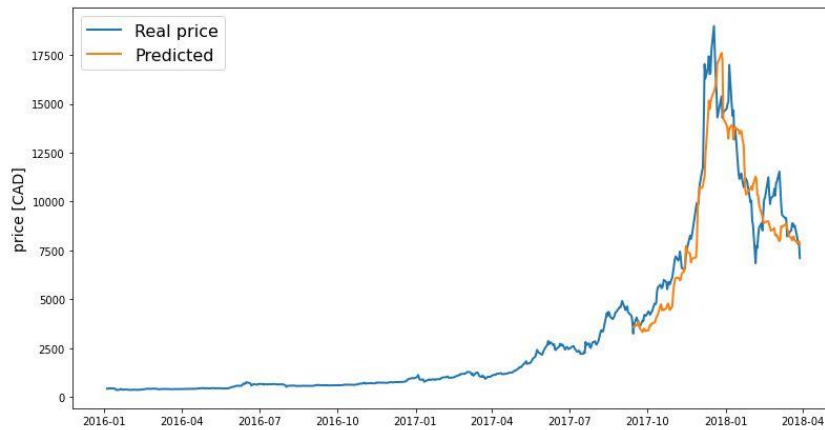
In the appendix we perform such an exercise which proves to render a considerably lower RMSE in the *consolidation period*. In other words, when included as a complementary feature, the lag of the price of bitcoin helps the LSTM model make a better prediction of future prices. However, this is achieved at the expense of interpretability since SHAP places a lot of weight on the lag of the price of bitcoin. Obviously, for the purposes of our exercise, these results are neither be very informative nor useful. Therefore, we refrained from taking this path for the rest of the exercise.

Figure 5: Predicted and actual bitcoin price. Launch period




Source: Author's own calculations (2022).

Figure 6: Predicted and actual bitcoin price. Consolidation period



Source: Author's own calculations (2022).

Figure 7: Predicted and actual bitcoin price. Expansion period



Source: Author's own calculations (2022).

5.3 SHAP

The increase in the use of ML models has awoken an interest in how to explain their outcome. There are different techniques that can accomplish this (see Molnar 2020 for a detailed review and a comprehensive list of methods). Some of the most popular techniques are the so-called model agnostic or post hoc interpretability techniques, that can be applied to any model. They can either be used to explain which features matter for a single prediction (local interpretability), or to explain which features matter more in the whole dataset (global interpretability). In this paper we focus on latter aspect, since our goal is to establish which factors determine the price of bitcoin over long periods, rather than on a single day.

The two main global interpretability techniques are SHAP (Lundberg and Lee, 2017, Lundberg et al 2020)¹⁴ and Permutation Feature Importance (introduced by Breiman 2001 for Random forest, and model agonist version provided by Fisher, et al. 2018). Both SHAP and Permutation Feature importance are global

¹⁴ The Shapley values can be used as local interpretability technique.

interpretability techniques that can yield a comprehensive ranking of the importance of features to the ML model in question. We have decided to use SHAP instead of Permutation Feature Importance as the main analysis for the following reasons.

First, SHAP is growingly gaining traction in the context of Deep Learning models (Albanesi and Vamossy 2019). In addition, it seems to have an advantage when features are correlated (Molnar 2020, Alonso and Carbó 2022). While it's true that highly correlated features will imply unrealistic permutations in both the above methods, these will be particularly significant in the case of Permutation Feature Importance. And third, while Permutation Feature Importance informs about which features matter more for a prediction's error, SHAP delivers directly the importance of a given feature in terms of its impact on the prediction itself. This allows us to compute the percentage impact of each feature on the price of bitcoin. In any case, none of these techniques is absent criticism (see, e.g., Rudin, 2019).

How does SHAP determine the importance of each feature? SHAP is a technique that measures the contribution of a variable to the predicted outcome, on a particular day, compared to the average prediction. These contributions are called the Shapley values. Once we have the Shapley values, for each variable and for each day, these can be added to obtain the final importance of the variable. Therefore, SHAP can be used as a local and global interpretability technique.

The approach to compute Shapley values can be explained from a game theory perspective. The game would be to reproduce the result of the model (in our case, the price of bitcoin). The players would be all possible coalitions of variables. Finally, the reward would be the contribution of each coalition towards the final outcome of the model. For illustrative purposes, here is an example. Suppose we decide to use the following variables: "Gold," "SP500," and "Hash rate". Next, consider that we want to know the extent to which the "SP500" is important in order to establish the price of bitcoin on a given day t . Based on the above, these are the four possible coalitions of variables without "SP500":

- Without variables
- Gold
- Hash rate
- Gold and Hash rate

For the four coalitions, we calculate the price in t with and without the "SP500". The Shapley value of the "SP500" in relation to the price of bitcoin on day t is the weighted average of those marginal contributions. To obtain the global importance of the "SP500" for the price of bitcoin in the test sample, we repeat the process for all days within the sample and compute the average of the Shapley values. In the appendix we provide further details as to how to find the solution analytically.

5.4 Determinants of the price of bitcoin according to SHAP

Since we considered three different periods, we ended up with three different sets of predictions. In this section, we take stock of SHAP (see section 4.3) to analyze the extent to which each of the 25 features were important for those predictions. The results can be found in **figures 8, 9 and 10**.

Period #1 – Launch.

Figure 8 depicts the top 15 features which, according to SHAP, influenced the price of bitcoin in the *launch period*. A hallmark of this initial phase was how technology-related variables inherent to bitcoin showed clear signs of playing a critical role. In fact, Hash difficulty emerged as the most important variable, with an effect

on the price of bitcoin of 13% over the mean prediction in the test sample. Other proprietary technological factors like, e.g., unique addresses, number of transactions, and fees to miners ranked among the top seven features as well. All of them helped explain price behavior with an effect that is calculated to fall above 2%. Regarding those variables which reflect public attention, Google Trends appeared to be the third most relevant driver, with an effect of around 4% over the mean prediction. As we will see later on, both Google Trends and Twitter, gain greater importance in subsequent periods. Finally, among the variables which reflect the status of the economy results were heterogeneous. In particular, the Yuan appeared as the second most relevant feature, but gold and the SP500 were nowhere near the top 10 spots.

Period #2 - Expansion.

In accordance with SHAP, the top 15 features that explain the price of bitcoin during this phase are shown in **figure 9**. In 2018, technology-related variables inherent to bitcoin continued to be among the most important determinants, although they were no longer as important as was the case in the *launch period*. Yet, Hash difficulty still was the most critical price driver, with an even more pronounced effect over the mean prediction: around 20%. In contrast, other variables capturing the technology dimension like, e.g., unique addresses, number of transactions, mean block size, etc. were much less relevant during the *expansion period* in comparison to the role they played during the *launch period*.


On the other hand, in the *expansion period* the variables related to the level of public attention/interest became much more critical. In particular, Google Trends ranked as the second most important explanatory factor, with an effect of 13% over the mean prediction. Likewise, Tweets appeared among the top 5 ones. Regarding economic variables, again heterogeneity was the rule. Opposite to what was the case in the *launch period*, the variable Yuan lost all of its importance. Moreover, gold remained unimportant, and the Dow, Nasdaq and the SP500 despite showing up in the top 10 spots, exhibit a rather low impact (3.5%, 2.5% and 2.4% respectively) in comparison with the top variables.

Period #3 - Consolidation.

As above, **figure 10** displays the top 15 features that help explain the price of bitcoin as derived from SHAP in the *consolidation period*. It is worth highlighting how technology-related variables associated with bitcoin lost visibly importance: while fees to miners were now the second most important aspect, with an effect of 15% over the mean prediction, other variables falling into this category did not appear among the top 15. It is especially remarkable as well how Hash difficulty, the most critical variable in both the *launch* and the *expansion period* now suddenly became completely irrelevant. On the other hand, the variables related to the level of public attention/interest definitely took over: Google Trends replaced previous candidates as the main explanatory factor with great impact: i.e. over 23% over the mean prediction. Again, the results were very different across the variables which reflect the state the economy is in. Overall they seem to have an intermediate importance, never at the top, but also not irrelevant at all. Among those, the most important ones for this period seemed to be GBPound exchange rate, the Nasdaq and the Yuan, all of them with an impact of 5% over the mean production.

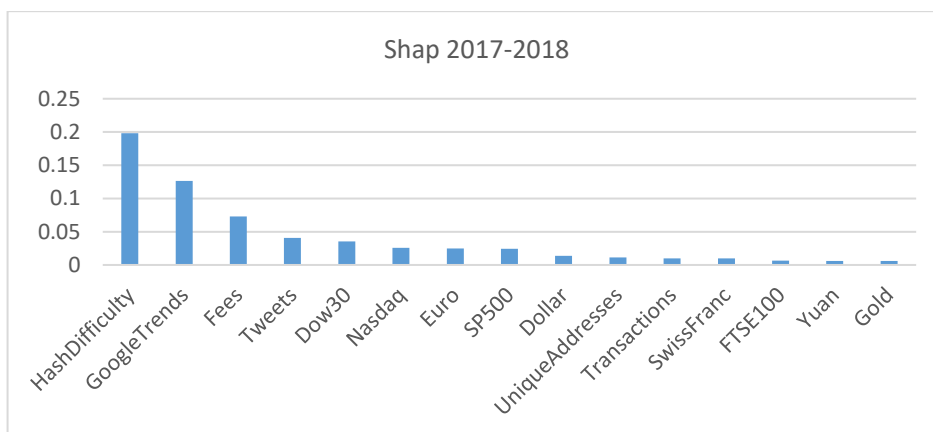
Finally, in order to obtain a general overview of how each category evolved to become more or less relevant, we aggregated all the features within each distinctive group: i.e. technological variables, economic variables, and public attention variables. For each period, we combined the SHAP values of all features, and we computed which percentage belonged to each category. These results are summarized in **figures 11, 12 and 13**. Thus, in the first two periods, technological variables were clearly the most important ones. Interestingly in the last period, they lost relevance quite visibly: i.e. from above 60% of all the impact in the first period, to less than 21% in 2021. Variables related to sentiment gained importance as the years went by. Their overall effect started at around 9% in the *launch period*, and climbed to 34% in the *consolidation period*. Economic variables did not present a clear trend.

Figure 8: Difference with respect to the mean prediction of including or not the variable, in percentage terms. **Launch period**



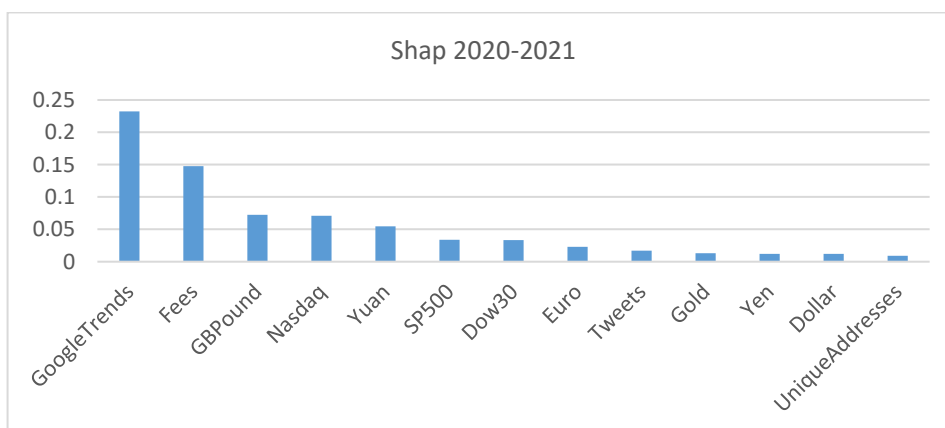
Source: Own authors (2022)

Figure 9: Difference with respect to the mean prediction of including or not the variable, in percentage terms. **Expansion period**



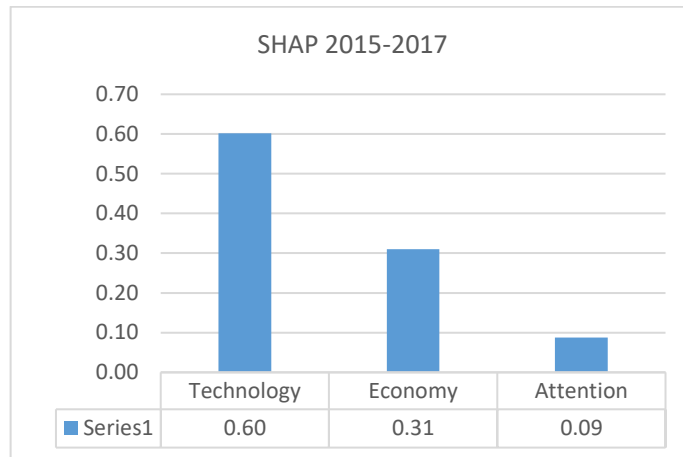
Source: Own authors (2022)

Figure 10: Difference with respect to the mean prediction of including or not the variable, in percentage terms. **Consolidation period**



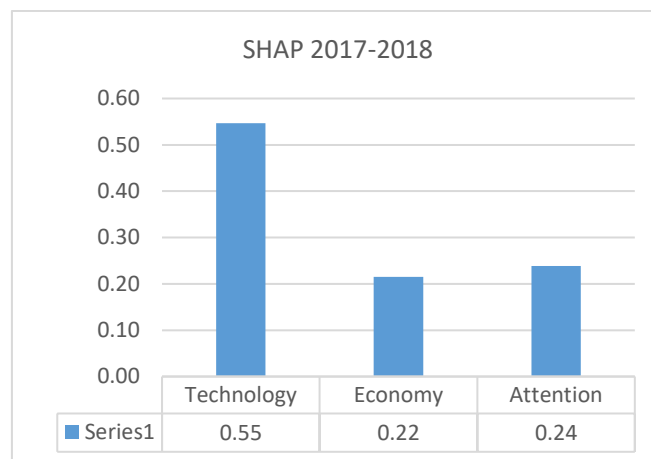
Source: Own authors (2022)

Figure 11: Aggregation by category. *Launch period*



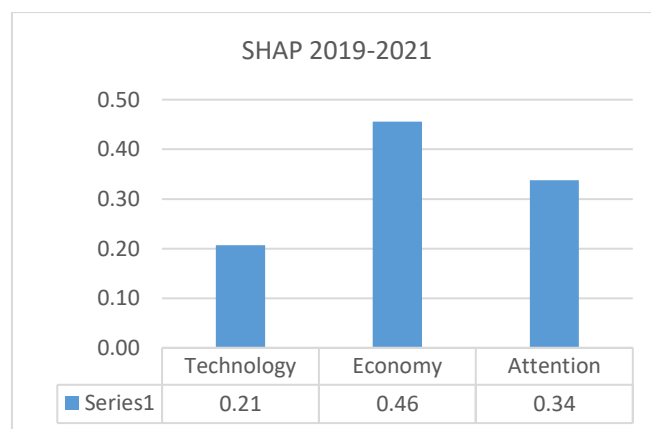
Source: Own authors (2022)

Figure 12: Aggregation by category. *Expansion period*



Source: Own authors (2022)

Figure 13: Aggregation by category. *Expansion period*



Source: Own authors (2022)

6. Conclusions

In this article we have explored which elements are most likely to explain the behavior of bitcoin prices at a given point in time: i.e. in between (i) 2015 and 2017, (ii) 2017 and 2018, and (iii) 2020 and 2021. We have also analyzed the extent to which these features may be longstanding or change over the years. The aim of this exercise is to help shed some light on the degree of maturity achieved by such digital asset in comparison to other financial instruments and alternative crypto-related value propositions such as those which are backed by underlying assets (i.e. stablecoins). This way we expect to potentially contribute to the current broader public policy debate on future actions to be adopted by relevant macro and microprudential authorities for the safeguarding of financial stability.

As a point of departure, we have first taken stock of the existing literature on both the underlying nature of bitcoin (e.g. safe haven, hedge, diversifier or speculative asset) and its price/return determinants. As a result, we have identified three main set of potential drivers -(i) macroeconomic factors, technology-related ones as well as those that reflect the degree of attention that said crypto-asset may be drawing- that we have leveraged on for our purposes.

Against this light, we found that technological variables emerged as the more relevant ones for the determination of bitcoin prices during the first two periods of our sample. However, they lost all its significance as we entered the last observation period. More precisely, variables such as hash difficulty, block size, number of transactions or unique addresses rendered virtually irrelevant to elucidate the evolution of bitcoin prices as we neared 2021. Conversely, variables pointing at the degree of public attention enjoyed by bitcoin -like Google Trends- grew progressively in importance as we came closer to the present day. In fact, in stages defined by high price volatility (2018 and 2021), the interest of the public takes on a very notable role.

It is also worth noting that variables highlighting the role of macroeconomic and financial development never took up a leading place in the establishment of bitcoin prices. Throughout the entire observation period, they seem to exercise a limited, yet constant influence while also showing a great degree of heterogeneity among each of its individual components. This appears to be in contradiction with the findings of part of the literature which oftentimes underscores the critical positive impact of the SP500 and gold on the determination of price.

On account of the above, our research leads us to conclude that the formation of the price of bitcoin is still a highly complex phenomenon whose underlying causes are difficult to anticipate with an acceptable degree of uncertainty. While most of the determinants highlighted in the literature seem to clearly play a role in the evolution of bitcoin prices over time, we prove that their influence does change substantially at short notice. Moreover, possibly due to the immaturity of crypto-assets markets, oftentimes new explanatory factors emerge unexpectedly which, furthermore, may remain undetected and opaque to both investors and authorities for long periods of time. This may be one of the reason why our predictive model performs notably worse in 2021.

For the above reasons, compared to other well-known and well-established asset classes, bitcoin - and, by extension, its namesakes - seems to continue to exhibit a difficult-to-predict behavior, thus making it a high-risk investment in the current landscape. It is, therefore, advisable for financial authorities to be fully aware of this fact upon deciding, at least, on the prudential treatment to be assigned to the potential exposures of banks to unbacked crypto-assets, in particular as regards market and liquidity risk, as well as in relation to the adoption of other relevant conduct-related measures in defense of investors and consumers at large.

This may, for example, call for deeper reflections by authorities on the implied model risk and further vindicate the amount of public warnings on crypto-assets that both national competent authorities and regional regulators have been issuing over time. In addition, such circumstance is supportive of more recent measures

aimed at supervising the way these offerings are advertised in order to better cope with the existing asymmetries in end-users' knowledge and understanding of the actual risks these digital assets entail.

More broadly, our findings could also be of interest to macroprudential authorities in their assessment of the materiality of the potential risks crypto-assets place on global financial stability and the need and timeliness of the deployment of effective regulatory and supervisory actions. Against this light, financial authorities may further want to consider maintaining conservative positions regarding their regulation so as to avoid the transmission of potentially systemic risks to the financial system as a whole.

7. References

- Abbasimehr, H., M. Shabani and M. Yousefi (2020). "An optimized model using LSTM network for demand forecasting", *Computers & Industrial Engineering*, 143, 106435.
- Albanesi, S., and D. F. Vamossy (2019). *Predicting consumer default: A deep learning approach*, No. w26165, National Bureau of Economic Research.
- Alessandretti, L., A. El Bahrawy, L. M. Aiello and A. Baronchelli (2018). *Anticipating cryptocurrency prices using machine learning*, Complexity.
- Alonso, A., and J. M. Carbó (2022). *Accuracy and stability of explanations of machine learning models for credit default prediction*, forthcoming.
- Aslanidis, N., A. F. Bariviera and Ó. G. López (2022). "The link between cryptocurrencies and Google Trends attention", *Finance Research Letters*, 102654.
- Bala, R., and R. P. Singh (2019). "Financial and non-stationary time series forecasting using LSTM recurrent neural network for short and long horizon", in 2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT) (pp. 1-7), IEEE.
- Bariviera, A. F., and I. Merediz-Solà (2021). "Where do we stand in cryptocurrencies economic research? A survey based on hybrid analysis", *Journal of Economic Surveys*, 35(2), pp. 377-407.
- Baur, D. G., K. Hong and A. D. Lee (2018). "Bitcoin: Medium of exchange or speculative assets?", *Journal of International Financial Markets, Institutions and Money*, 54, pp. 177-189.
- Bouri, E., P. Molnár, G. Azzi, D. Roubaud and L. I. Hagfors (2017). "On the hedge and safe haven properties of Bitcoin: Is it really more than a diversifier?", *Finance Research Letters*, 20, pp. 192-198.
- Bouoiyour, J., and R. Selmi (2015). "What does Bitcoin look like?", *Annals of Economics & Finance*, 16(2).
- Bouoiyour, J., and R. Selmi (2017). *The Bitcoin price formation: Beyond the fundamental sources*, arXiv preprint arXiv:1707.01284.
- Breiman, L. (2001). "Random forests", *Machine Learning*, 45(1), pp. 5-32.
- Carbó, J. M., and E. Diez García (2021). *El interés por la innovación financiera en España: un análisis con Google Trends*, Documentos Ocasionales, No. 2112, Banco de España.
- Cermak, V. (2017). *Can bitcoin become a viable alternative to fiat currencies? An empirical analysis of bitcoin's volatility based on a GARCH model* (May 2).
- Ciaian, P., M. Rajcaniova and D. A. Kancs (2016). "The economics of BitCoin price formation", *Applied Economics*, 48(19), pp. 1799-1815.
- Ciaian, P., D. A. Kancs and M. Rajcaniova (2018). *The price of BitCoin: GARCH evidence from high frequency data*, arXiv preprint arXiv:1812.09452.
- Chen, W., H. Xu, L. Jia and Y. Gao (2021). "Machine learning model for Bitcoin exchange rate prediction using economic and technology determinants", *International Journal of Forecasting*, 37(1), pp. 28-43.

- Fisher, A., C. Rudin and F. Dominici (2019). "All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously", *Journal of Machine Learning Research*, 20(177), pp. 1-81.
- Hayes, A. S. (2017). "Cryptocurrency value formation: An empirical study leading to a cost of production model for valuing bitcoin", *Telematics and Informatics*, 34(7), pp. 1308-1321.
- Hochreiter, S., and J. Schmidhuber (1997). "Long short-term memory", *Neural Computation*, 9(8), pp. 1735-1780.
- Kaminski, J. (2014). *Nowcasting the bitcoin market with twitter signals*, arXiv preprint arXiv:1406.7577.
- Kapar, B., and J. Olmo (2021). "Analysis of Bitcoin prices using market and sentiment variables", *The World Economy*, 44(1), pp. 45-63.
- Kim, Y. B., J. G. Kim, W. Kim, J. H. Im, T. H. Kim, S. J. Kang and C. H. Kim (2016). "Predicting fluctuations in cryptocurrency transactions based on user comments and replies", *PLOS ONE*, 11(8), e0161197.
- Kjærland, F., A. Khazal, E. A. Krogstad, F. B. Nordstrøm and A. Oust (2018). "An analysis of bitcoin's Price dynamics", *Journal of Risk and Financial Management*, 11(4), p. 63.
- Kristoufek, L. (2015). "What are the main drivers of the Bitcoin price? Evidence from wavelet coherence analysis", *PLOS ONE*, 10(4), e0123923.
- Lundberg, S. M., and S. I. Lee (2017). "A unified approach to interpreting model predictions", *Advances in neural information processing systems*, pp. 4765-4774.
- Lundberg, S. M., G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, and S. I. Lee (2020). "From local explanations to global understanding with explainable AI for trees", *Nature Machine Intelligence*, 2(1), pp. 56-67.
- Lyócsa, Š., P. Molnár, T. Plíhal and M. Širaňová (2020). "Impact of macroeconomic news, regulation and hacking exchange markets on the volatility of bitcoin", *Journal of Economic Dynamics and Control*, 119, 103980.
- Molnar, C. (2020). *Interpretable machine learning*, Lulu.com.
- Nakamoto, S. (2008). "Bitcoin: A peer-to-peer electronic cash system", *Decentralized Business Review*, 21260.
- Olah, C. (2015). *Understanding LSTM Networks*, retrieved from <https://colah.github.io/posts/2015-08-Understanding-LSTMs/>
- Panagiotidis, T., T. Stengos and O. Vravosinos (2018). "On the determinants of bitcoin returns: A LASSO approach", *Finance Research Letters*, 27, pp. 235-240.
- Pyo, S., and J. Lee (2020). "Do FOMC and macroeconomic announcements affect Bitcoin prices?", *Finance Research Letters*, 37, 101386.
- Rudin, C. (2019). "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead", *Nature Machine Intelligence*, 1(5), pp. 206-215.

Sovbetov, Y. (2018). "Factors influencing cryptocurrency prices: Evidence from bitcoin, ethereum, dash, litcoin, and monero", *Journal of Economics and Financial Analysis*, 2(2), pp. 1-27.

Urquhart, A., and H. Zhang (2019). "Is Bitcoin a hedge or safe haven for currencies? An intraday analysis", *International Review of Financial Analysis*, 63, pp. 49-57.

Urquhart, A. (2018). "What causes the attention of Bitcoin?", *Economics Letters*, 166, pp. 40-44.

Wang, T., S. Zhao, G. Zhu and H. Zheng (2021). "A machine learning-based early warning system for systemic banking crises", *Applied Economics*, 53(26), pp. 2974-2992.

Zhu, Y., D. Dickinson and J. Li (2017). "Analysis on the influence factors of Bitcoin's price based on VEC model", *Financial Innovation*, 3(1), pp. 1-13.

Zhu, P., X. Zhang, Y. Wu, H. Zheng and Y. Zhang (2021). "Investor attention and cryptocurrency: Evidence from the Bitcoin market", *PLOS ONE*, 16(2), e0246331.

8. Appendix

8.1 Definitions

As taken from Coinmetrics:

Difficulty finding the hash: The mean difficulty at a given day of finding a hash that meets the protocol-designated requirement (i.e., the difficulty of finding a new block).

Unique addresses: The sum count of unique addresses that were active in the network (either as a recipient or originator of a ledger change) at a given day.

Commissions to miners (fees): The sum USD value of all fees paid by user that makes transactions at a given day. Fees do not include new issuance.

Hash rate: The mean rate at which miners are solving hashes at a given rate. Hash rate is the speed at which computations are being completed across all miners in the network.

Sum of blocks: The sum count of blocks created that interval that were included in the main (base) chain at a given day.

Average block size: The mean size (in bytes) of all blocks created at a given day.

Sum of transfers: The sum count of transfers at a given day. Transfers represent movements of native units from one ledger entity to another distinct ledger entity. Only transfers that are the result of a transaction and that have a positive (non-zero) value are counted.


Medium transfer size: The sum value of native units transferred divided by the count of transfers (i.e., the mean size of a transfer) between distinct addresses that interval.

8.2 LSMT with lagged price

Our LSTM performed worse in the *consolidation period*, 2021, than in the other two periods. In this exercise we include as a feature the price of bitcoin in $t-1$, along with the original 25 features. With this new feature, the predictive performance of the LSTM improves considerably in every period, but particularly in the *consolidation period*. **Figure 14** shows the result from the LSTM model once we include the lag of the price of bitcoin as a feature. It can be seen that the gap between actual price and predicted price in March and April of 2021 is much smaller now than in the original exercise (see **Figure 7**).

The RMSE is of 11%, in contrast with 21% in the main exercise. Our model can predict better once we include the lag of the price. The reason for not including this variable in the main exercise is that the bitcoin price lag will bias our interpretation of the SHAP values, since it will appear as the main determinant. But this will not help us in our analysis, which is to understand what are the potential determinants of the price of bitcoin.

Figure 14: Predicted and actual bitcoin price using lag price as feature. Expansion period



Source: Own authors (2022)

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS

- 2040 GABRIEL JIMÉNEZ, DAVID MARTÍNEZ-MIERA and JOSÉ-LUIS PEYDRÓ: Who Truly Bears (Bank) Taxes? Evidence from Only Shifting Statutory Incidence.
- 2041 FELIX HOLUB, LAURA HOSPIDO and ULRICH J. WAGNER: Urban air pollution and sick leaves: evidence from social security data.
- 2042 NÉLIDA DÍAZ SOBRINO, CORINNA GHIRELLI, SAMUEL HURTADO, JAVIER J. PÉREZ and ALBERTO URTASUN: The narrative about the economy as a shadow forecast: an analysis using Banco de España quarterly reports.
- 2043 NEZIH GUNER, JAVIER LÓPEZ-SEGOVIA and ROBERTO RAMOS: Reforming the individual income tax in Spain.
- 2101 DARÍO SERRANO-PUENTE: Optimal progressivity of personal income tax: a general equilibrium evaluation for Spain.
- 2102 SANDRA GARCÍA-URIBE, HANNES MUELLER and CARLOS SANZ: Economic uncertainty and divisive politics: evidence from the *Dos Españas*.
- 2103 IVÁN KATARYNIUK, VÍCTOR MORA-BAJÉN and JAVIER J. PÉREZ: EMU deepening and sovereign debt spreads: using political space to achieve policy space.
- 2104 DARÍO SERRANO-PUENTE: Are we moving towards an energy-efficient low-carbon economy? An input-output LMDI decomposition of CO₂ emissions for Spain and the EU28.
- 2105 ANDRÉS ALONSO and JOSÉ MANUEL CARBÓ: Understanding the performance of machine learning models to predict credit default: a novel approach for supervisory evaluation.
- 2106 JAVIER ANDRÉS, ÓSCAR ARCE and PABLO BURRIEL: Market polarization and the Phillips curve.
- 2107 JUAN de LUCIO and JUAN S. MORA-SANGUINETTI: New dimensions of regulatory complexity and their economic cost. An analysis using text mining.
- 2108 DANILO LEIVA-LEON and LUIS UZEDA: Endogenous time variation in vector autoregressions.
- 2109 CORINNA GHIRELLI, ANDREA GONZÁLEZ, JOSÉ LUIS HERRERA and SAMUEL HURTADO: Weather, mobility and the evolution of the Covid-19 pandemic.
- 2110 KLODIANA ISTREFI, FLORENS ODENDAHL and GIULIA SESTIERI: Fed communication on financial stability concerns and monetary policy decisions: revelations from speeches.
- 2111 YUNUS AKSOY, HENRIQUE S. BASSO and CAROLYN ST AUBYN: Time Variation in Lifecycle Consumption and Income.
- 2112 JENNIFER PEÑA and ELVIRA PRADES: Price setting in Chile: micro evidence from consumer on-line prices during the social outbreak and Covid-19.
- 2113 NEZIH GUNER, YULIYA A. KULIKOVA and ARNAU VALLADARES-ESTEBAN: Does the added worker effect matter?
- 2114 RODOLFO G. CAMPOS and JACOPO TIMINI: Unequal trade, unequal gains: the heterogeneous impact of MERCOSUR.
- 2115 JAVIER QUINTANA: Import competition, regional divergence, and the rise of the skilled city.
- 2116 PATRICK MACNAMARA, MYROSLAV PIDKUYKO and RAFFAELE ROSSI: Marginal Tax Changes with Risky Investment.
- 2117 RODOLFO G. CAMPOS, JACOPO TIMINI and ELENA VIDAL: Structural gravity and trade agreements: does the measurement of domestic trade matter?
- 2118 ALESSANDRO MELCARNE, JUAN S. MORA-SANGUINETTI and ROK SPRUK: Democracy, technocracy and economic growth: evidence from 20 century Spain.
- 2119 ÁNGEL ESTRADA and DANIEL SANTABÁRBARA: Recycling carbon tax revenues in Spain. Environmental and economic assessment of selected green reforms.
- 2120 ALEJANDRO FERNÁNDEZ-CEREZO, BEATRIZ GONZÁLEZ, MARIO IZQUIERDO and ENRIQUE MORAL-BENITO: Firm-level heterogeneity in the impact of the COVID-19 pandemic.
- 2121 EDUARDO GUTIÉRREZ and CÉSAR MARTÍN MACHUCA: The effect of tariffs on Spanish goods exports.
- 2122 JACOPO TIMINI: Revisiting the 'Cobden-Chevalier network' trade and welfare effects.
- 2123 ALEJANDRO BUESA, JAVIER J. PÉREZ and DANIEL SANTABÁRBARA: Awareness of pandemics and the impact of COVID-19.
- 2124 ALICIA AGUILAR and DIEGO TORRES: The impact of COVID-19 on analysts' sentiment about the banking sector.
- 2125 SILVIA ALBRIZIO, IVÁN KATARYNIUK, LUIS MOLINA and JAN SCHÄFER: ECB euro liquidity lines.
- 2126 ANTHONY BALD, ERIC CHYN, JUSTINE HASTINGS and MARGARITA MACHELETT: The causal impact of removing children from abusive and neglectful homes.
- 2127 IRMA ALONSO, PEDRO SERRANO and ANTONI VAELO-SEBASTIÀ: The impact of heterogeneous unconventional monetary policies on the expectations of market crashes.

- 2128 MARÍA T. GONZÁLEZ-PÉREZ: Lessons from estimating the average option-implied volatility term structure for the Spanish banking sector.
- 2129 SIMÓN A. RELLA, YULIYA A. KULIKOVA, EMMANOUIL T. DERMITZAKIS and FYODOR A. KONDRASHOV: Rates of SARS-COV-2 transmission and vaccination impact the fate of vaccine-resistant strains.
- 2130 MATÍAS LAMAS and DAVID MARTÍNEZ-MIERA: Sectorial holdings and stock prices: the household-bank nexus.
- 2131 ALBERT BANAL-ESTAÑOL, ENRIQUE BENITO, DMITRY KHAMETSHIN and JIANXING WEI: Asset encumbrance and bank risk: theory and first evidence from public disclosures in Europe.
- 2132 ISABEL ARGIMÓN and MARÍA RODRÍGUEZ-MORENO: Business complexity and geographic expansion in banking.
- 2133 LUIS GUIROLA: Does political polarization affect economic expectations?: Evidence from three decades of cabinet shifts in Europe.
- 2134 CHRISTIANE BAUMEISTER, DANILO LEIVA-LEÓN and ERIC SIMS: Tracking weekly state-level economic conditions.
- 2135 SERGI BASCO, DAVID LÓPEZ-RODRÍGUEZ and ENRIQUE MORAL-BENITO: House prices and misallocation: The impact of the collateral channel on productivity.
- 2136 MANUEL ARELLANO, STÉPHANE BONHOMME, MICOLE DE VERA, LAURA HOSPIDO and SIQI WEI: Income risk inequality: Evidence from Spanish administrative records.
- 2137 ANGELA ABBATE and DOMINIK THALER: Optimal monetary policy with the risk-taking channel.
- 2138 MARTA BAÑBURA, DANILO LEIVA-LEÓN and JAN-OLIVER MENZ: Do inflation expectations improve model-based inflation forecasts?
- 2139 MÁXIMO CAMACHO, MARÍA DOLORES GADEA and ANA GÓMEZ LOSCOS: An automatic algorithm to date the reference cycle of the Spanish economy.
- 2140 EDUARDO GUTIÉRREZ, AITOR LACUESTA and CÉSAR MARTÍN MACHUCA: Brexit: Trade diversion due to trade policy uncertainty.
- 2141 JULIO A. CREGO and JULIO GÁLVEZ: Cyclical dependence in market neutral hedge funds.
- 2142 HERVÉ LE BIHAN, MAGALI MARX and JULIEN MATHERON: Inflation tolerance ranges in the new keynesian model.
- 2143 DIEGO COMIN, JAVIER QUINTANA, TOM SCHMITZ and ANTONELLA TRIGARI: Measuring TFP: the role of profits, adjustment costs, and capacity utilization.
- 2144 ROBERTO PASCUAL: Do analysts forecast differently in periods of uncertainty? An empirical analysis of target prices for Spanish banks.
- 2145 BEATRIZ GONZÁLEZ, GALO NUÑO, DOMINIK THALER and SILVIA ABRIZIO: Firm heterogeneity, capital misallocation and optimal monetary policy.
- 2201 RYAN BANERJEE and JOSÉ-MARÍA SERENA: Dampening the financial accelerator? Direct lenders and monetary policy.
- 2202 JUAN S. MORA-SANGUINETTI and ISABEL SOLER: La regulación sectorial en España. Resultados cuantitativos.
- 2203 JORGE E. GALÁN, MATÍAS LAMAS and RAQUEL VEGAS: Roots and recourse mortgages: handing back the keys.
- 2204 MONICA MARTINEZ-BRAVO and CARLOS SANZ: Inequality and psychological well-being in times of COVID-19: evidence from Spain.
- 2205 CORINNA GHIRELLI, DANILO LEIVA-LEÓN and ALBERTO URTASUN: Housing prices in Spain: convergence or decoupling?
- 2206 MARÍA BRU MUÑOZ: Financial exclusion and sovereign default: The role of official lenders.
- 2207 RICARDO GIMENO and CLARA I. GONZÁLEZ: The role of a green factor in stock prices. When Fama & French go green.
- 2208 CARLOS MONTES-GALDÓN and EVA ORTEGA: Skewed SVARs: tracking the structural sources of macroeconomic tail risks.
- 2209 RODOLFO G. CAMPOS, ILIANA REGGIO and JACOPO TIMINI: Thick borders in Franco's Spain: the costs of a closed economy.
- 2210 MARCO CELENTANI, MIGUEL GARCÍA-POSADA and FERNANDO GÓMEZ POMAR: Fresh start policies and small business activity: evidence from a natural experiment.
- 2211 JOSE GARCIA-LOUZAO, LAURA HOSPIDO and ALESSANDRO RUGGIERI: Dual returns to experience.
- 2212 ADRIÁN CARRO and PATRICIA STUPARIU: Uncertainty, non-linear contagion and the credit quality channel: an application to the Spanish interbank market.
- 2113 MARIO IZQUIERDO, ENRIQUE MORAL-BENITO, ELVIRA PRADES and JAVIER QUINTANA: The propagation of worldwide sector-specific shocks.
- 2214 HENRIQUE S. BASSO: Asset holdings, information aggregation in secondary markets and credit cycles.
- 2215 JOSÉ MANUEL CARBÓ and SERGIO GORJÓN: Application of machine learning models and interpretability techniques to identify the determinants of the price of bitcoin.