Natural Hazards and Financial Activity: Evidence from Solar Storms Impact on BTC Mining*

Panayotis G. Michaelides
National Technical University of Athens
& Hellenic Open University

Arsenios-Georgios Prelorentzos
Foundation of Economic and Industrial Research

Olivier Scaillet†
University of Geneva (GSEM)
& Swiss Finance Institute

Nikolas Topaloglou
Paris School of Business

Kien C. Tran
University of Lethbridge
& Center for Econometrics and Business Analytics

December, 2025

*We are grateful to Manthos Delis for his generous and insightful comments and suggestions. We also sincerely thank our dear friend and co-author, Mike Tsionas, for his insightful discussions on the econometric aspects of this research. Tragically, Mike Tsionas passed away about 2 years ago. May his soul rest in peace. Also, special thanks are due to Theodoros Daglis for kindly providing the dataset. The usual disclaimer applies.

†Corresponding author. Email: Olivier.Scaillet@unige.ch

Abstract

Digitalization of finance has rapidly transformed the global financial system by introducing new complex technologies. As digital finance grows more interconnected, we show in this work that solar geomagnetic storms, i.e., an exogenous space-weather shock, pose a systemic risk comparable to climate transition pressures, with measurable consequences for Bitcoin (BTC) mining. BTC mining, an energy-intensive process reliant on stable electricity and continuous network connectivity, faces increasing risks from solar storms, i.e., geomagnetic disturbances, capable of disrupting power grids and digital infrastructure. This study introduces a novel step-by-step ANN-based causality test, leveraging Artificial Neural Networks (ANNs) to examine the nonlinear and dynamic causal relationship between solar storms and BTC mining activity. In this context, we formally establish consistency, asymptotic normality and other useful properties of the proposed ANN estimator. Using daily data from 09/2014 to 01/2023, our multi-horizon analysis identifies statistically significant short-, medium-, and long-term effects, revealing that severe solar storms immediately destabilize networks, delaying transaction confirmations and reducing BTC activity. Furthermore, we quantify the economic significance of these disruptions. During 2015 alone, geomagnetic storms were associated with a cumulative reduction of nearly 3 million BTC in trading activity, equivalent to an estimated value of around 883 million. The findings position solar storms as a critical but overlooked risk factor in digital finance, with implications for cryptocurrency market stability, mining infrastructure resilience, and regulatory oversight. Our research stresses the need for risk monitoring, geographically diversified mining operations, and enhanced grid stability measures.

Keywords: Digital Finance, Bitcoin, AI, Artificial Neural Networks, Solar Storms

JEL Classification: C22, C50, C51, C58.

1 Introduction

Over the last decades, global financial systems have grown increasingly complex and interconnected, fueled by innovations such as securitization, secondary market trading, and digital technologies. These developments have enabled more efficient risk-sharing, but they also introduce new systemic vulnerabilities. For instance, interconnected financial networks can simultaneously facilitate risk-sharing and act as channels for cascading failures (Jalan et al. 2024). At the same time, climate-related uncertainties have emerged as critical factors influencing financial stability. Previous research has established links between climate regulations or policy uncertainty and loan pricing in credit markets. Extending this line of inquiry, we introduce a distinct exogenous shock from the area of space weather: solar storms. We argue that extreme geomagnetic events of this kind represent a special case of systemic risk, akin to climate transition risks, that can disrupt financial operations. In the aforementioned framework, assets with heightened exposure to tail-event risk suffer disproportionately large losses during extreme market downturns, stressing how rare shocks amplify systemic vulnerabilities (Oordt and Zhou, 2016). Just as regulatory climate risks have been shown to alter banks' strategies and pricing, solar shocks can similarly reshape how financial participants organize consortia, manage exposures, and price credit. While much of the extant literature focuses on transition risks and borrower-specific climate factors, we highlight the underexplored influence of abrupt physical environmental shocks on risk management behavior. In this spirit, our study shifts the emphasis toward this unconventional source of uncertainty, i.e., solar storms, providing new insights beyond traditional firm-level considerations. We specifically examine the implications of solar storm disruptions in a modern context, focusing on Bitcoin (BTC) mining as a case study within the digital finance ecosystem.

Recent studies in finance focus on how climate risks alter lending behavior and market pricing. For instance, a carbon tax can prompt banks to restrict lending to carbon-intensive sectors (Laeven and Popov, 2023). Banks also tend to charge higher interest rate spreads to "brown" firms relative to "green" borrowers, as documented by Ehlers et al. (2022), and Reghezza et al. (2022). In a similar vein, Kacperczyk and Peydró (2022) found that once banks commit to decarbonization goals, they cut back financing to polluting companies, even if prior lending relationships exist. Complementing this, Degryse et al. (2023) show that banks with a "green" orientation are more likely to offer favorable terms to environmentally responsible firms, especially those with strong ESG credentials. Indeed, climate-driven shifts in institutional portfolios can exert significant price pressure on affected firms, elevating their cost of equity (Skiadopoulos and Xue, 2025). Moreover, European banks have been observed to pass on climate regulatory risks to heavy emitters. Ginglinger and Moreau (2023) showed that such risks are indirectly transmitted to polluting firms through stricter loan covenants and impacts on corporate capital structure. After all, modern risk measures for extreme outcomes acknowledge that tail risk assessment is inherently tied to model ambiguity. Therefore, a decision-theoretic foundation for tail-risk metrics highlights the importance of accounting for statistical model uncertainty in extreme-event evaluation (Koun and Peng, 2016). Collectively, these studies examine climate risk transmitted from borrowers to banks through firms' emissions profiles and sustainability attributes. Further to this, our analysis considers how an exogenous physical climate-related shock, i.e., solar geomagnetic storms, can directly affect the markets, complementing a related strand of literature.

More precisely, prior studies found that extreme weather events can disrupt firm operations and shift economic activity (Huang et al., 2022), and banks accordingly adjust loan terms to incorporate the heightened credit risk from natural disasters (Correa et al., 2022). In fact, in the absence of explicit insurance against rare disasters, investors indirectly self-insure through the stock and money markets, leading to countercyclical risk premia. In other words, a severe incident could provoke similar defensive portfolio rebalancing, affecting asset prices and expected returns (Dieckmann, 2011). Related evidence by Dlugosz et al. (2024) suggests that a decentralized approach to rate-setting can help banks absorb regional liquidity shocks induced by climate disasters. Our analysis differs from these works in that, while they mostly address regulatory or transition risks stemming from terrestrial climate policy, we focus on space-weather shocks. Our insights not only enrich the understanding of financial market dynamics under extreme environmental disturbances, but also set the stage for examining such phenomena in the realm of digital finance, where analogous risks threaten the stability of decentralized networks like Bitcoin mining.

At the same time, the advent of new technologies has ushered in a new era for the global financial system. Digital finance, i.e., a comprehensive domain encompassing innovations such as Bitcoin, digital wallets, decentralized finance (DeFi), and a broad suite of financial technologies underpinned by artificial intelligence (AI), has fundamentally reconfigured traditional paradigms of financial intermediation. This ongoing digital revolution introduced novel dimensions to capital management, risk assessment, and liquidity dynamics. Amidst this change, the digital currency popularly known as Bitcoin (BTC) has emerged to the forefront of the digital finance revolution and the adoption of a decentralized and digital asset system.

At the core of the BTC economy is a mining process that requires significant computational power and energy for the sustainment of BTC operations and its blockchain (Matringe & Power, 2024). Mining serves a dual purpose: Confirming the transactions and, at the same time, providing the decentralised ledger with decentralised integrity, and creating more BTCs into circulation. The BTC network's architecture deliberately places limitations on block transaction capacity. This design feature creates a competitive environment where transaction initiators must compete by offering higher processing fees to incentivize miners to prioritize their transactions for inclusion in subsequent blocks (Hautsch et al, 2024; Basu et al., 2023), while mines compete for the acquisition of new hardware to improve the efficiency of their operations, within this payment-seeking environment (Capponi et al., 2023).

However, as the difficulty of mining has risen, so has the energy consumption required to perform the process. For instance, the scale of BTC mining's energy consumption is particularly evident in Iceland, where cryptocurrency mining facilities' power usage surpasses the total residential electricity consumption of the entire nation (Constantinides et al., 2018). In the meantime, numerous works have noted that BTC's price shocks are directly linked to the marginal cost of historical mining costs, i.e., power costs, emissions, and clean energy integration. A recent analysis reveals that cryptocurrency market dynamics are influenced by unique market-specific elements that can serve as predictive indicators (Liu & Tsyvinski, 2021). Of course, cryptocurrency markets have experienced episodic turmoil, such as pump-and-dump schemes, trigger brief but extreme surges in price, trading volume, and volatility that swiftly collapse, stressing the fragility of cryptocurrency markets under manipulation pressures (Griffin & Shams, 2020; Shin & Wang, 2025). Furthermore, BTC and other cryptocurrencies have started to interact with clean energy markets as both threats and opportunities for the transition to renewable energy sources (Naeem et al., 2023).

Although, despite the rapid evolution of cryptocurrency market and digital finance, several vulnerabilities associated with decentralized systems remain largely unexplored. Among the most under-appreciated risks is the impact of space weather, specifically, solar storms capable of disrupting terrestrial power grids and digital networks.

This study introduces a novel application of Artificial Neural Networks (ANNs) to investigate the causal link between geomagnetic disturbances and BTC mining efficiency. Such a model is warranted because when financial dynamics deviate from normality assumptions, nonlinear models can better capture predictable behavior, yielding more accurate risk assessments than static, unconditional models (Hsieh, 1993). More precisely, recent evidence in financial literature indicates that machine-learning methods, and particularly ANNs, demonstrate superior predictive performance relative to alternative approaches in return forecasting settings. For instance, Gu et al. (2020) show that ANN-based models outperform a broad set of traditional techniques in predicting stock returns, effectively capturing complex and intricate nonlinear interactions.

By analysing nearly a decade of daily data, our multi-horizon approach captures the immediate, medium-term, and long-run effects of solar storms on the Bitcoin ecosystem. Geomagnetic disturbances caused by solar storms can lead to power outages, damage satellites, and disrupt internet connectivity (Sigaeva et al., 2009, Eastwood et al., 2017, Daglis et al., 2019), all of which are essential for the continuous and reliable operation of BTC network. Contrary to other traditional economic sectors, cryptocurrency markets such as BTC are uniquely vulnerable to such disturbances as their two core functions, i.e., transaction processing and mining, rely on real-time computational processes requiring uninterrupted access to both stable electricity and network connectivity (Clark and Greenley, 2019, Corbet et al., 2020). One could argue that the decentralized and geographically dispersed nature of the system offers some resilience, and to an extent, this is true. However, mining is fundamentally a computational process through which transactions are validated and confirmed, and new coins are generated, making mining and transaction activity mutually dependent. Therefore, disruptions to electricity or internet access, whether affecting users initiating transactions or miners validating them, pose a threat to the overall stability of the BTC network. Even if some miners are temporarily unable to validate some transactions due to such disruptions, the burden shifts to the remaining active miners, increasing their workload, energy consumption needs, computational demands, and transaction confirmation times, ultimately undermining the system's operational efficiency and cost-effectiveness. Furthermore, despite BTC's decentralized architecture, approximately 72% of global mining activity is concentrated in just three countries: the United States, China, and Kazakhstan [1]. Thus, while solar storms do not affect mining through conventional economic channels, they do so through a technological-infrastructural channel that directly governs the operational dynamics of BTC mining performance.

The findings reveal that severe space-weather events trigger delays in blockchain confirmations, reduce mining efficiency, and increase operational costs. Over time, these disruptions compound, leading to infrastructure failures and regulatory concerns. More precisely, our findings reveal that geomagnetic disturbances, such as solar storms, have a statistically significant and time-varying causal impact on BTC mining activity, with effects manifesting immediately (short-term, $h = 1$–10), persisting through infrastructure-level adjustments (medium-term, $h = 11$–20), and lingering well into the operational and economic dynamics of mining ecosystems (long-term, $h = 21$–30). By applying a novel Step-by-Step ANN-based causality framework, rooted in the theory of universal function approximation, we uncover a complex temporal structure of influence that traditional linear models are unable to capture. Moreover, our ANN-based approach allows us to identify and unveil the intricate nonlinear causal relationships linking solar storms and geomagnetic disturbances to disruptions in power grids and internet connectivity, and consequently, to BTC mining activity. Within this framework, to assess the economic significance of these effects, we quantify the pecuniary losses associated with geomagnetic disturbances. As we notice, severe storms are heavily concentrated in the early years of BTC adoption, particularly in 2015, when 21 of the 33 major storms in our sample occur. During this period, storm days exhibit markedly lower activity, with cumulative reductions in BTC trading volume reaching nearly 3 million BTC in 2015 alone, corresponding to approximately $883 million at prevailing prices. Scenario-based impulse response simulations further reveal that extreme geomagnetic shocks generate substantial short-run contractions in BTC activity: a Dst shock of –200 nT, equal to an intense solar storm, suppresses market activity by $20–30 billion, while stronger shocks of –250 nT and –300 nT deepen these losses to $30–40 billion and $40–50 billion, respectively. Our results underscore the urgent need for robust risk management, diversified mining operations, and real-time monitoring tools to safeguard the stability of digital finance in an era increasingly influenced by environmental shocks.

The persistence of causality across multiple time horizons shows that the impact of solar storms on BTC mining is not confined to the immediate aftermath but unfolds progressively over time. This temporal spillover highlights the presence of delayed disruptions, ranging from infrastructure stress to market instability, that spill over through the entire cryptocurrency ecosystem. The economic repercussions of solar storms tend to unfold gradually, often materializing through increased electricity prices in affected regions, the costly replacement of damaged ASIC miners and power infrastructure, and the need for regulatory or operational adjustments to mitigate future risks. The persistence of these effects stresses the critical role of adaptive mining infrastructure in responding to space-weather shocks. By tracing these various horizon effects, our analysis offers new insight into how space-weather events propagate within decentralized digital systems, uncovering a critical

[1] `https://ccaf.io/cbnsi/cbeci/mining_map`

yet underexplored source of systemic vulnerability in modern financial infrastructures. Rare but potentially catastrophic events often lack historical data and involve complex interdependencies, making conventional statistical risk models unreliable (Sampson and Smith, 1982). Taken together, our approach provides a powerful diagnostic tool for mapping exogenous environmental shocks to endogenous instability in digital finance systems, highlighting a new frontier in the era of digital finance in both econometric methodology and systemic risk analysis.

The paper makes significant contributions to the relevant literature by introducing, for the first time in the literature, the Step-by-Step ANN-based Causality Test, a novel framework that integrates ANNs to address the inherent complexities of non-linear and dynamic relationships in financial systems. More precisely, our motivation for developing and employing an ANN-based causality methodology lies in the effort to model and reveal the inherent complexity and nonlinearity that link solar storms and geomagnetic disturbances to disruptions in power grids and network connectivity and, hence, to BTC mining, a topic that remains largely unexplored. Traditional linear causality tests, such as Granger causality, or Dufour & Renault tests, assume stable lag structures and linear dependencies, assumptions that are likely to be violated in systems characterized by intricate and dynamic relationships, such as the modern digital finance ecosystems. In this context, our state-of-the-art Step-by-Step approach, which leverages ANNs, allows us to identify and capture the nonlinear and multi-horizon effects of solar storms on BTC mining activity by providing a flexible, data-driven framework capable of modeling these dynamics without imposing rigid assumptions. Moreover, since this is among the first studies to empirically examine this causal link, our objective is not merely to detect whether a relationship exists, but to capture how and when it manifests across different time horizons and at varying intensities. Furthermore, we also perform robustness checks to validate our results across different specifications and activation functions, ensuring the reliability of the findings.

The proposed methodology extends traditional causality analysis by offering horizon-specific insights, capturing highly non-linear and delayed responses, feedback loops, and multi-level interactions often overlooked in existing approaches. By applying this framework to examine the impact of solar storms on Bitcoin (BTC) mining, the study provides the first exhaustive investigation into how solar storms influence the cryptocurrency ecosystem. Specifically, to formalize the mechanism through which geomagnetic disturbances affect BTC mining, Section 3.1 develops a compact structural model linking solar-storm intensity to mining infrastructure uptime, effective hash power, and ultimately daily BTC production. This structural formulation provides a transparent bridge between solar storms and the functioning of the BTC network, thereby motivating the ANN-based causality framework that follows. Beyond methodological advancements, this work provides critical insights for policymakers and industry stakeholders facing the risks and opportunities of an increasingly interconnected digital economy, while also setting a foundation for future research at the intersection of AI, digital finance, and financial economics.

The paper is structured as follows. Section 2 provides stylized facts and relevant background literature on both BTC mining and solar storms. Section 3 introduces the novel step-by-step causality model, leveraging ANNs. The empirical results are presented in Section 4, while Section 5 discusses the findings and provides policy implications. Finally, Section 6 gives some concluding remarks and provides avenues for future research. We gather the proofs of our inferential theory in Appendix A and relate to the results for ANNs available in the literature. Appendix B provides additional empirical results

2 Background Literature

2.1 Stylized Facts

The escalating challenges related to climate change present obstacles to the global economic and financial systems. Although the consequences of BTC mining have been researched extensively, little has been done to understand how external factors, and more specifically, solar climate, affect mining. Within this framework, there is a growing recognition of the complex relationship between these systems and the physical world. The Sun, being the main source of climate on the Earth, has a major impact on energy systems, communication networks, and other components of the digital finance environment. Flares and sunspots, for instance, can interfere with energy grids and communication, necessities for BTC mining. NASA's Solar Dynamics Observatory captured an X-class solar flare, i.e., a bright flash on the Sun's limb. Such intense solar eruptions can launch charged particles and magnetic plasma toward Earth, potentially disrupting electrical and electronic systems. Solar storms can induce electromagnetic effects on Earth. These geomagnetic storms have been known to damage satellites, knock out power grids, and interfere with communications. BTC mining, which relies on specialized hardware running 24/7, stable electrical power, and constant internet connectivity, is vulnerable to severe space weather.

In what follows, we examine how increased solar activity could affect mining hardware, power infrastructure, and network connectivity, review historical incidents, and discuss mitigation strategies. Even though the data of NASA/JPL-Caltech indicate the historical minimum of sunspot numbers of the modern solar cycle, significant warming effects on the planet have been observed, which express the multifaceted relationship between the Sun and the Earth's systems (Daglis et al., 2023a, 2023b). In this framework, a solar storm is an event that interferes with the magnetosphere and has severe negative effects. Solar storms induce geomagnetic currents in power grids and lead to damage to numerous pieces of equipment and blackouts. Solar storms also have the impact of interfering with various methods of telecommunication. In turn, there is the risk that disruptions in telecommunications may affect processes related to BTC mining. Due to the uncertainty of disruptions and the operating risks related to solar storms, both traditional and digital firms may face significant threats to their operations and profitability. Moreover, these vulnerabilities are consistent with recent insights from operations and supply-chain research, which emphasize that blockchain technologies, which have expanded and adopted rapidly across various industries, should be viewed as complex systems exposed to coordination failures, system bottlenecks, and external shocks (Lumineau et al., 2025).

An example of solar activity-related telecommunications wire meltdowns can be found in the Carrington event of 1859. In order to illustrate the impact of such phenomena, consider the projected economic costs of a Carrington-level solar storm, which are estimated to range between 0.6 billion and 2.6 trillion in the United States, depending, of course, on the number of people potentially affected, ranging between 20–40 million (Eastwood et al., 2017). Another example could be linked to the Halloween Solar Storm of 2003, which affected 47 satellites, with one being lost for good, and ten others being affected for some time (Eastwood et al., 2017). The Halloween Solar Storm of 2003 exposed the risks facing satellite operations, requiring further operational costs of between 10 million and 50 million for the aviation sector alone (Johnson and Lee, 2004). On January 20, 1994, Canada's Anik E1 and E2 communications satellites failed due to solar activity, leaving TV, phone, and data services cut off across a huge area. For one satellite, the disruption lasted hours, and for the other lasted weeks, demonstrating that a solar storm can impose lengthy effects on the global network infrastructure. If a mining pool's coordination servers were on a satellite that went down, miners would have to failover to terrestrial links or mine solo until service returned. In fact, the most notorious geomagnetic storm that occurred on March 13, 1989, which is known as one of the most devastating solar phenomena in contemporary times, triggered such violent voltage fluctuation in the electricity grid of Hydro Québec that the whole system shut down, resulting into the expulsion of voltage from six (6) million people, and not to mention billions of US dollar damage. Just two minutes later, the whole network was disconnected (Green et al., 2016). Thus, due to an extreme solar event, earthbound communications can be disrupted (Sigaeva et al., 2009, Daglis et al., 2019). In this framework, Den Baeumen (2014) pointed out that such events as the Quebec blackout could lead to annual losses of between 2.4 and 3.4 trillion. Indeed, the blackout of electricity in Quebec in 1989 is an event that can remind us of the vulnerability of electrical grids, and by extension, the large losses that follow in various electricity-intensive processes like BTC mining. These are some of the examples of the need for further research on the topic (Johnson and Lee, 2004; Williams et al., 2011).

In order to fill this gap, this paper proposes a new methodological approach to analyze the causal link between solar storms and the digital finance sphere, namely BTC mining. This study also benefits from the development of artificial intelligence (AI), especially artificial neural networks (ANNs), to capture the highly nonlinear relationship between these variables. The conventional approaches to causality testing do not take into consideration the complexity of such systems, which are recursive in nature. The step-by-step causality test based on the ANNs proposed in this research is an extension of the seminal work by Renault and Dufour (1998), and Dufour et al. (2006), but with the use of ANNs.

It is of paramount importance to study the effects of AI on digital finance because it can serve as a revolutionary and disruptive technology. AI plays a significant role in making processes automated, enabling data-driven decision-making, and drawing useful conclusions from big data through the support of Machine Learning (ML) and other algorithms. Coordinating these processes is crucial to realizing the potential of the AI revolution and also to overcoming some problems associated with it (Brynjolfsson and McAfee 2014). Floridi et al. (2018) noted that non-normative technical issues that must be tackled include transparency, fairness, and accountability so that AI will only produce and develop positively. These discussions highlight the possibility of achieving the optimum utilization of the potential of AI in offering solutions to many of the global challenges.

2.2 BTC Mining

BTC mining is an essential activity in the global economy and the financial system, with rising energy requirements to perform this task. Mining is an activity of validating the transactions as well as adding to the list of records known as a blockchain. These are users who try to solve such cryptographic problems; the one who solves such a problem first gets the privilege to incorporate a new block in the blockchain; in return for his/her efforts, he/she is rewarded with newly created coins of cryptocurrencies (Nakamoto, 2008). Therefore, mining of BTC is a crucial task that is performed as a way to validate the transactions within the blockchain. As the ecosystem develops, the algorithms that have to be solved in order to successfully perform the mining demand even more and more computational power and resources, which results in an excessive consumption of electricity. Therefore, to carry out mining, there is a high demand for energy. In this regard, cryptocurrencies and especially blockchain technology (Corbet et al., 2020) show both the prospects and the problems for the energy sector (Clark and Greenley, 2019). However, no work has considered the solar storms that cause numerous problems to the mining process. Such storms influence the volume of financial transactions, as empirical evidence shows that solar climate impacts the energy sector (Daglis et al., 2023) and the telecommunication sector (Daglis et al., 2019). For mining to be effective, uninterrupted energy is needed for the miner's devices to solve the cryptographic issues. Therefore, considerable interaction in the energy sector, for various regions, can greatly hamper the mining process. Besides that, the internet and telecommunications are also crucial for the BTC ecosystem to work as it should. This was acknowledged in that earlier studies have revealed that BTC's certain attributes, like price fluctuations (Alexander and Heck, 2020), integration with other markets, and high volatility (Naeem et al, 2023) may affect and be affected by the broader financial system. Furthermore, it has been established that BTC exhibits volatility spillover effects on other related assets, indicating that BTC could affect financial stability (Elsayed et al., 2022).

Moreover, the movement of BTC is connected to how energy and technology companies move. The interrelation of the cryptocurrency projects is one of the factors that affect the volatility of the whole market, as there are observed similar trends in DeFi, NFTs, and large market capitalization cryptocurrencies like Bitcoin (Rahman et al., 2024). However, the speculation on the cryptocurrency market may result in erratic shifts in trade, affecting the stability of the whole financial system (Bourghelle et al., 2022, p. 135). Thus, solar storms affect BTC mining in terms of the number of transactions. Therefore, based on the link between BTC and the whole financial system, this complex relationship results in financial vulnerability. However, while solar climate is generally understood as the provision of energy to energy-intensive industries, little has been done to determine the effects of such intense solar storms on the BTC mining process. It is crucial to fill this gap in the literature since the understanding of this relationship between solar storms and BTC mining might be vital to anticipate the interruption in mining and the conclusions made regarding the BTC market.

2.3 Solar Storms

There is a dearth of literature on the subject due to the scarcity of data in this area (Eastwood et al. 2017). Most likely, the first person who stated that the activity of our Sun may influence economic and financial activity was Jevons (1878, 1879, 1882) in a famous article in Nature. Jevons's first arguments in his initial papers were that the Sun may influence the economy and society in general (Chizhevsky, 1924, 1938, 1976) and the 'Jevons hypothesis' was expanded by Garcia-Mata and Shaffner (1934) in the Quarterly Journal of Economics where one of the parameters that affected these results was the Sun (Tasnim et al. 2016, Daglis, et al. 2019).

To the best of our knowledge, no study has looked at the effect of solar storms on BTC mining, and thus, this study seeks to fill this gap. Specifically, in this subsection, we discuss how solar storms can affect the volume of BTC transactions. The technological and societal effects of solar activity vary depending on the atmospheric and geographical conditions of the planet, as well as surface and radiation levels in various parts of the world. A decrease in sunspot activity was observed since 1750 in the current solar cycle that started on January 4, 2008 (IPCC, 2013). At the same time, data obtained by NASA/JPL-Caltech suggest a warming effect from the Sun during this period (Daglis et al. 2023). In this regard, solar storms, which are closely associated with other solar activities, have important consequences. These storms impact satellite operations and power systems critical to daily transactions and market functioning, as well as communication systems. Furthermore, explosions on the surface of the Sun caused by fluctuations of magnetic fields transmit energy into space; this may hinder communication systems, besides affecting other technologies on Earth (Daglis et al. 2019).

A normal magnetic storm involves the formation of a large ring current within a few hours and a slow dissipation over several days. The magnetic field and plasmasphere may be deformed by solar events (Baket et al. 2005). Therefore, during a magnetic storm, energy particles, mainly protons and electrons, are emitted when the Earth's magnetic field is distorted (Asikainen et al. 2005), where these particles interfere with systems, endanger spacemen, and often interrupt radio signals (Xystouris et al. 2014). During a magnetic storm, the ring current shows a sharp negative shift, which indicates the motion of particles produced during the storm, and large-scale currents are mainly controlled by changes in solar activity (Daglis et al. 2019). Overall, as has been discussed above, solar storms can directly and indirectly affect BTC mining, involving numerous procedures and technologies (Green et al. 2016). In this regard, it is necessary that proper technological tools are available for the performance and handling of the massive number of electronic transactions that occur in firms across the globe. To perform such a time-stamping, GPS and timing devices are employed, both of which have gained popularity in recent years (Krausmann et al., 2014). To perform these tasks, BTC mining needs constant electricity and, of course, access to the internet, which solar storms may interfere with, thus causing an interruption (Ahluwalia, 2015).

3 Methodology

Building on BTC's mining process, and the nature of solar activity, we now compactly, mathematically formalize the mechanism that links storms to mining outcomes. The key observations are: (i) solar storms are external to the crypto economy at daily horizons; (ii) disruptions manifest as reductions in the effective uptime of energy and mining devices; and (iii) BTC production follows a block-arrival process whose rate scales with effective hash power, and inversely with protocol difficulty. To isolate this channel with minimal structure, we introduce relevant expressions. This parsimonious setup will deliver closed-form relationships between storm intensity and core outcomes and, after a standard log-linearization, will yield a regression-ready equation. We next present the compact mathematical model that will make the mechanism explicit and will lead directly into the empirical specifications estimated below.

3.1 Structural Model: From Solar Storms to BTC Mining

Let $X_t \in \mathbb{R}^k$ collect exogenous indicators of solar activity such (e.g. Dst) at day t. Let $U_t \in [0,1]$ denote the fraction of mining infrastructure effectively *available*, with no power / device outage. In this context, solar storms reduce availability: $U_t \;=\; 1 \;-\; \alpha' X_t \;+\; \varepsilon_t^U, \quad \alpha \in \mathbb{R}^k_{\geq 0},\; \mathbb{E}[\varepsilon_t^U | X_t] = 0.$

Let $\widetilde{H}_t$ be the *potential* total computational power devoted to mining, and is equal to: $H_t \;=\; U_t \cdot \widetilde{H}_t$. Now, D_t is the BTC *difficulty* at time t, i.e. a dimensionless index that scales how hard it is to find a valid block, and implies average inter-block time $\tau_t \propto D_t/H_t$, given computational power rate H_t. The protocol *retargets* D_t to keep the average block time constant.

Let the average inter-block time be equal to $\tau_t \;=\; \vartheta\, \frac{D_t}{H_t}$, where $\vartheta > 0$ is a units constant. Over a day of length Δ, the expected number of blocks is equal to $B_t \;=\; \frac{\Delta}{\tau_t} \;=\; \frac{\Delta}{\vartheta}\,\frac{H_t}{D_t} \;\equiv\; \Delta\kappa\,\frac{H_t}{D_t}$, where $\kappa \equiv 1/\vartheta$, for notational convenience. Hence, the quantity of BTC mined is: $Q_t \;=\; R_t\, B_t$, where R_t is the BTC reward per block which remains fixed between halvings. Taking conditional expectations and using the previous equation, yields $\mathbb{E}[Q_t \mid X_t, D_t, \widetilde{H}_t] \;=\; R_t\, \Delta\kappa\, \frac{U_t \widetilde{H}_t}{D_t}$.

Now, a log-linearization around mean values gives us the empirical link equation:

$$\underbrace{\log Q_t}_{\text{mined quantity}} \;\approx\; \underbrace{\log(\Delta\kappa)}_{\text{const}} + \underbrace{\log R_t - \log D_t}_{\text{protocol terms}} + \underbrace{\log \widetilde{H}_t}_{\text{fundamentals}} + \underbrace{\log U_t}_{\text{solar storm impact}} + \eta_t.$$

We have that: $\log U_t \approx -\alpha' X_t + \nu_t$, for small outages. Therefore:

$$\underbrace{\log Q_t}_{\text{mined quantity}} \;\approx\; \underbrace{\log(\Delta\kappa)}_{\text{const}} + \underbrace{\log R_t - \log D_t}_{\text{protocol terms}} + \underbrace{\log \widetilde{H}_t}_{\text{fundamentals}} + \underbrace{-\alpha' X_t}_{\text{solar storm impact}} + \; u_t.$$

This equation delivers a coherent equation of the mined quantity.

Now, to bring the previous equation to the data, while allowing persistence and feedback through difficulty D_t, we use ARDL (Autoregressive Distributed Lag)-type specifications. Let Z_t collect controls that are either exogenous or predetermined at daily frequency. We get the following quantity equation: $Y_t = a_Y + \sum_{r=1}^{p} \phi_r^{(Y)} Y_{t-r} + \beta_Y' X_t + \Gamma_Y' Z_t + u_t^{(Y)}$, where $Y_t = \log Q_t$

To implement a forward-looking, horizon-h causality test consistent with the mechanism, we use the following: $Y_{t+h}^{(y)} = a^{(y,h)} + \Pi^{(y,h)} Y_{t,p}^{(y)} + b^{(y,h)\prime} X_{t,q} + \Lambda^{(y,h)\prime} Z_{t,q} + \varepsilon_{t+h}^{(y)}$, where $Y_{t,p}^{(y)} = (Y_t^{(y)}, \dots, Y_{t-p}^{(y)})$ and $X_{t,q} = (X_t, \dots, X_{t-q})$. In the absence of controls, we have that $Y_{t+h}^{(y)} = a^{(y,h)} + \Pi^{(y,h)} Y_{t,p}^{(y)} + b^{(y,h)\prime} X_{t,q} + \varepsilon_{t+h}^{(y)}$.

In conclusion, this mechanism above delivers compact and testable predictions: solar activity X_t reduces effective uptime U_t, thereby lowering hash power H_t and the mined quantity Y_t, while increasing inter-block times τ_t. Guided by this mechanism, we now take the model to the data with time–series specifications in which (log) outcomes Y_t are regressed on contemporaneous and lagged X_t, and autoregressive terms to absorb persistence. Under the solar storm exogeneity of X_t at daily horizons, the coefficients trace causal steps, and the dynamic profiles reveal the speed and duration of the storm's impact on mining.

3.2 A Novel ANN Causality Test

At a fundamental level, the evolution of a dynamic system subject to exogenous shocks can be modelled as a discrete-time stochastic process governed by an unknown functional operator.

Let $\{Y_t\}_{t\in\mathbb{Z}}$ denote the response process of interest (e.g., Bitcoin transaction volume), and $\{X_t\}_{t\in\mathbb{Z}}$ denote the exogenous driver (e.g., solar storms). Then, the system can be described abstractly by $Y_t = \mathcal{G}_\theta(\mathcal{H}(X_t, X_{t-1}, \dots), \mathcal{L}(Y_{t-1}, Y_{t-2}, \dots)) + where \mathcal{H}(\cdot)$ and $\mathcal{L}(\cdot)$ are operators acting on the exogenous and endogenous inputs, $\mathcal{G}_\theta$ is a nonlinear response function parameterized by θ, and ε_t is a stochastic innovation term. In classical time series analysis, the operator $\mathcal{G}_\theta$ is often assumed to be linear, while in modern machine learning applications, it is flexibly modelled using universal function approximators such as Artificial Neural Networks (ANNs). In what follows, we will specialize this representation to develop a formal ANN-based causality framework suitable for capturing the complex dynamics between solar storms and Bitcoin mining activity.

So, consider a general nonlinear dynamic system in which the outcome variable Y_t evolves over time as a function of its own lags, the current and past values of an exogenous input X_t, and a stochastic innovation u_t. Formally, we get $Y_t = \mathcal{F}(Y_{t-1}, Y_{t-2}, \dots, Y_{t-p}, X_t, X_{t-1}, \dots, X_{t-q}) + u_t$, where $\mathcal{F}(\cdot)$ is an unknown, potentially nonlinear and high-dimensional function. In traditional econometric modelling, $\mathcal{F}$ is approximated by linear functions or low-order polynomials, which are often inadequate for capturing the complex interactions, delayed effects, and feedback loops inherent in systems like digital finance and decentralized infrastructure. ANNs, by contrast, provide a flexible, data-driven mechanism to approximate the unknown functional form $\mathcal{F}$, as they are universal function approximators capable of modelling arbitrary nonlinear mappings. This motivates the specification of an ANN-based structure, allowing us to estimate and test for dynamic causality in a more robust and generalizable way. This architecture is consistent with recent advances that show how deep learning frameworks can uncover latent nonlinear factor structures, significantly improving asset-pricing performance relative to traditional linear models (Feng et al. 2024).

In this context, to begin with a typical baseline model, consider a linear autoregressive distributed lag model for time series data:

$$Y_t = a + \sum_{r=1}^{p} \pi_r Y_{t-r} + \sum_{j=0}^{q} b_j X_{t-j} + u_t, \quad t = 1, \cdots, T, \tag{1}$$

where $Y_t \in \Re^m$ is a vector of m dependent variables, $X_t \in \Re^l$ is an l_m-vector of explanatory variables, $a \in \Re^m$ is an m-vector of constants, π_r and b_j are the coefficients for the lagged terms of Y_t and X_t, respectively, and u_t is an m-vector of random shocks which we assume to be normal with zero mean vector and covariance matrix Σ_u. The number of lags p and q is selected based on some relevant Information Criteria.

To test the causality on the horizon h, we construct a forward-looking model:

$$Y_{t+h} = a^{(h)} + Y_{t,p}\pi^{(h)} + b^{(h)} X_{t,q} + u_{t+h}, \quad h = 1, 2, ..., \tag{2}$$

where $Y_{t,p} = (Y_t, Y_{t-1}, \cdots, Y_{t-p})$ and $X_{t,q} = (X_t, X_{t-1}, \cdots, X_{t-q})$ represent the current and lagged values vectors of Y and X, and $\pi^{(h)} = \left(\pi_0^{(h)}, \pi_1^{(h)}, \cdots, \pi_p^{(h)}\right)'$ and $b^{(h)} = \left(b_0^{(h)}, b_1^{(h)}, \cdots, b_q^{(h)}\right)'$ captures the horizon-specific coefficients.

3.3 ANN Representation

Recently, ANNs have received increasing attention in economic and financial research because of their flexibility and capability to model non-linear relationships and dynamics. This flexibility accounts for the extensive early application of ANNs in Economics and Finance (e.g. Athanassopoulos and Curram, 1996; Fleissig et al., 2000; Tsitsiklis and Roy, 2001; White and Racine, 2001; Hansen and Nelson, 2003; Sfetsos and Siriopoulos, 2005; Michaelides et al., 2010; Michaelides et al., 2015). In tandem with this, ML approaches are also quickly emerging, especially in portfolio management and financial forecasting (Rundo et al., 2019; Ozbayoglu, 2020; Nazareth and Reddy, 2023). For example, Chen and Huang (2021) predicted S&P 500 with an impressive performance by employing ANNs, and Sermpinis et al. (2021) compared recurrent ANNs with other competitors to anticipate NASDAQ 100 and NIKKEI225 indices. In a similar context to Dutta et al. (2020), the authors used recurrent ANNs for forecasting BTC prices, thereby achieving better results than traditional methods such as ARMA, as experienced by Cao and Wang (2019) and Adebiyi et al. (2014) with the Hang Seng and NYSE financial indices, among others. Some earlier studies by Cao et al. (2011) and Yang et al. (2010) made comparisons between the ANN results of the linear model, including the CAPM as well as mixed non-linear specifications, where ANNs turn out to be more efficient on several occasions. Apart from finance, fields such as Data Science, Management Science, Econometrics, and Operations Research have increasingly embraced ANNs, demonstrating their versatility and applicability to a wide range of problems. In general, ANNs outperform conventional techniques in both modelling and forecasting, offering superior capabilities in capturing complex, non-linear relationships and dynamic interactions that traditional methods struggle to address.

Traditional causality tests, such as the famous Granger causality test, are effective in identifying whether one time series has predictive power over another, but often fail to capture the temporal dynamics and multi-step relationships inherent in complex systems. These limitations are particularly evident in scenarios where causality operates indirectly through intermediate variables or across multiple time horizons. Step-by-step causality testing addresses these shortcomings by analysing causal relationships at specific horizons, allowing for a more dynamic understanding of when and how causality emerges. This approach is especially crucial in systems characterized by non-linear dynamics, delayed responses, and feedback loops, features often observed in financial markets, cryptocurrency ecosystems, and external environmental influences. By breaking down the causality into discrete steps, this method provides deeper insights into the timing and pathway of causal interactions, offering a richer analytical framework for investigating complex economic and financial phenomena. The proposed approach builds on Dufour and Renault (1998) and Dufour et al. (2006), by extending the traditional linear model with ANN components to model nonlinearities and improve robustness across horizons.

In conclusion, traditional causality tests often fall short when applied to complex, nonlinear systems like BTC mining and its interaction with external environmental shocks. These methods assume linear relationships and fixed-lag structures, overlooking the complex, dynamic, and recursive nature of financial and technological disruptions. Given the highly nonlinear dependencies between solar storms and digital infrastructure, a more flexible and adaptive approach is required. To address this gap, in what follows, we introduce a novel Step-by-Step ANN-based Causality Test, which extends conventional methodologies by leveraging ANNs to capture horizon-specific causal effects, delayed feedback mechanisms, and nonlinear interactions.

Therefore, to account for all non-linearities, we replace the linear terms for X_t and its lags with an ANN representation. ANNs are collections of transfer functions capable of modelling any relationship with a set of input variables $X' = [X_1, \cdots, X_q]$. The input variables are then combined linearly to form K intermediate variables/nodes $\{Z_1, \cdots, Z_K\}$ where $Z_k = X\beta_k$, for $k = 1, \cdots, K$, and $\beta_k \in \Re^{l(q+1)}$ are the weighting vectors. Consequently, the intermediate variables are combined nonlinearly as follows:

$$\sum_{k=1}^{K} \gamma_k \varphi(Z_k) \tag{3}$$

where $\varphi(\cdot)$ is an activation function, γ_k are unknown parameters, and K is the number of nodes (Kuan and White, 1994). Hence, equation (2) becomes:

$$Y_{t+h} = a^{(h)} + Y_{t,p}\pi^{(h)} + \sum_{k=1}^{K} \gamma_k^{(h)} \varphi(Z_k) + u_{t+h}, \quad h = 1, 2, ... \tag{4}$$

For simplicity, we use a feedforward ANN with one (1) hidden layer, and an activation function equal to $\varphi(\beta X_{t,q})$ where φ is a logistic activation function: $\varphi(z) = (1 + \exp(-z))^{-1}$, β is the weighting vector and $X_{t,q}$ are the inputs.

Note that the activation function is a critical component of an ANN, enabling it to capture highly nonlinear relationships and complex interactions between input and output variables. In the context of ANN-based causality models, the activation function determines how the weighted sum of inputs is transformed at each neuron before being passed to the next layer. In their seminal work, Kuan and White (1994) observed that conditions that are generally applied to the activation functions are boundedness, differentiability, and monotonicity. These functions are usually defined within a certain domain, such as $[0, 1]$. Of these, the most popular one in empirical research is the sigmoid function. This function is relevant because it is quite close to linear performance in the middle range and because it constrains the dependent variables to a finite maximum value of 1. Another popular activation function is the hyperbolic tangent function, which guarantees that the values of the dependent variables are within the range $[-1, 1]$. Other activation functions, which are also widely employed, consist of the Gaussian kernel and the Threshold Logical Unit (TLU) functions, being particularly relevant for Hopfield networks (Haykin, 1999). In this work, we employ the logistic activation function and also examine an alternative function to ensure robustness.

For estimation purposes, the following subsection presents the estimable models.

3.4 Estimable Model and Causality Testing

Firstly, we begin with the general model in Equation (1). Without loss of generality and for simplicity reasons, we assume: $p = q = 2$. Then, Equation (1) becomes:

$$Y_t = a + \pi_1 Y_{t-1} + \pi_2 Y_{t-2} + b_0 X_t + b_1 X_{t-1} + b_2 X_{t-2} + u_t, \quad t = 1, \cdots, T. \tag{5}$$

Next, to test causality at horizon h, we construct a forward-looking model:

$$Y_{t+h} = a^{(h)} + Y_{t,p}\pi^{(h)} + b^{(h)} X_{t,q} + u_{t+h}, \quad h = 1, 2, ..., \tag{6}$$

where $Y_{t,p} = (Y_t, Y_{t-1}, \cdots, Y_{t-p})$, $\pi^{(h)} = \left(\pi_0^{(h)}, \pi_1^{(h)}, \cdots, \pi_p^{(h)}\right)$, $b^{(h)} = \left(b_0^{(h)}, b_1^{(h)}, \cdots, b_q^{(h)}\right)$, $X_{t,q} = (X_t, X_{t-1}, \cdots, X_{t-q})$, and $u_{t+h} = (u_{1,t+h}, \cdots, u_{m.t+h})'$ for $t = 1, \cdots, T - h$ and $h < T$.

Subsequently, we introduce the ANNs term by reformulating the previous equation. Hence, Equation (6) becomes:

$$Y_{t+h} = a^{(h)} + Y_{t,p}\pi^{(h)} + \sum_{k=1}^{K} \gamma_k^{(h)} \varphi(Z_k) + u_{t+h}, \quad h = 1, 2, ... \tag{7}$$

In what follows, we fully develop the econometric model to be estimated, first in the case with one node and then with two nodes. In both cases, we assume that $p = q = 2$.

One node:

- For $h = 1$

$$\begin{aligned} Y_{t+1} &= a^{(1)} + \pi_0^{(1)} Y_t + \pi_1^{(1)} Y_{t-1} + \pi_2^{(1)} Y_{t-2} + \gamma_1^{(1)} \varphi \left(X_t \beta_{10}^{(1)} + X_{t-1} \beta_{11}^{(1)} + X_{t-2} \beta_{12}^{1)} \right) + u_{t+1} \\ &= a^{(1)} + \pi_1^{(1)} Y_{t-1} + \pi_2^{(1)} Y_{t-2} + \gamma_1^{(1)} \left[1 + \exp \left(X_t \beta_{10}^{(1)} + X_{t-1} \beta_{11}^{(1)} + X_{t-2} \beta_{12}^{(1)} \right) \right]^{-1} + u_{t+1}. \end{aligned} \tag{8}$$

- For $h = 2$

$$\begin{aligned} Y_{t+2} &= a^{(2)} + \pi_0^{(1)} Y_t + \pi_1^{(2)} Y_{t-1} + \pi_2^{(2)} Y_{t-2} + \gamma_1^{(2)} \varphi \left(X_t \beta_{10}^{(2)} + X_{t-1} \beta_{11}^{(2)} + X_{t-2} \beta_{12}^{(2)} \right) + u_{t+2} \\ &= a^{(2)} + \pi_1^{(2)} Y_{t-1} + \pi_2^{(2)} Y_{t-2} + \gamma_1^{(2)} \left[1 + \exp \left(X_t \beta_{10}^{(2)} + X_{t-1} \beta_{11}^{(2)} + X_{t-2} \beta_{12}^{(2)} \right) \right]^{-1} + u_{t+2}. \end{aligned} \tag{9}$$

Two nodes:

- For $h = 1$

$$\begin{aligned} Y_{t+1} &= a^{(1)} + \pi_0^{(1)} Y_t + \pi_1^{(1)} Y_{t-1} + \pi_2^{(1)} Y_{t-2} + \gamma_1^{(1)} \varphi \left(X_t \beta_{10}^{(1)} + X_{t-1} \beta_{11}^{(1)} + X_{t-2} \beta_{12}^{(1)} \right) \\ &\qquad + \gamma_2^{(1)} \varphi \left(X_t \beta_{20}^{(1)} + X_{t-1} \beta_{21}^{(1)} + X_{t-2} \beta_{22}^{(1)} \right) + u_{t+1}, \\ &= a^{(1)} + \pi_0^{(1)} Y_t + \pi_1^{(1)} Y_{t-1} + \pi_2^{(1)} Y_{t-2} + \gamma_1^{(1)} \left[1 + \exp \left(X_t \beta_{10}^{(1)} + X_{t-1} \beta_{11}^{(1)} + X_{t-2} \beta_{12}^{(1)} \right) \right]^{-1} \\ &\qquad + \gamma_2^{(1)} \left[1 + \exp \left(X_t \beta_{20}^{(1)} + X_{t-1} \beta_{21}^{(1)} + X_{t-2} \beta_{22}^{(1)} \right) \right]^{-1} + u_{t+1}. \end{aligned} \tag{10}$$

- For $h = 2$

$$\begin{aligned} Y_{t+2} &= a^{(2)} + \pi_0^{(1)} Y_t + \pi_1^{(2)} Y_{t-1} + \pi_2^{(2)} Y_{t-2} + \gamma_1^{(2)} \varphi \left(X_t \beta_{10}^{(2)} + X_{t-1} \beta_{11}^{(2)} + X_{t-2} \beta_{12}^{(2)} \right) \\ &\qquad + \gamma_2^{(2)} \varphi \left(X_t \beta_{20}^{(2)} + X_{t-1} \beta_{21}^{(2)} + X_{t-2} \beta_{22}^{(2)} \right) + u_{t+2}, \\ &= a^{(2)} + \pi_0^{(1)} Y_t + \pi_1^{(2)} Y_{t-1} + \pi_2^{(2)} Y_{t-2} + \gamma_1^{(2)} \left[1 + \exp \left(X_t \beta_{10}^{(2)} + X_{t-1} \beta_{11}^{(2)} + X_{t-2} \beta_{12}^{(2)} \right) \right]^{-1} \\ &\qquad + \gamma_2^{(2)} \left[1 + \exp \left(X_t \beta_{20}^{(2)} + X_{t-1} \beta_{21}^{(2)} + X_{t-2} \beta_{22}^{(2)} \right) \right]^{-1} + u_{t+2}, \end{aligned} \tag{11}$$

and similarly for $h = 3,\ 4,\ \cdots$.

Empirical estimation is based on Equations (8)–(11), which define a highly nonlinear model due to the inclusion of ANN components. Following the approach of Chan and Genovese (2001), we adopt a hybrid estimation strategy: we treat the nonlinear parameters—specifically the neural network weights $\beta_k^{(h)}$—as random draws, while estimating the remaining parameters using conventional econometric techniques. This estimation is operationalized through the following structured algorithm:

- *Step 1 (Initialization)*: Randomly draw the nonlinear weights $\beta_k^{(h)}$ from a uniform distribution.
- *Step 2 (Conditional Estimation)*: Conditional on the chosen $\beta_k^{(h)}$, estimate the linear parameters $\pi_r^{(h)}$, $\gamma_k^{(h)}$, and the intercept $a^{(h)}$ using ordinary least squares (OLS).
- *Step 3 (Model Evaluation)*: Compute the sum of squared residuals (SSR). Repeat Steps 1 and 2 for multiple draws of $\beta_k^{(h)}$, and retain the optimal set $\bar{\beta}_k^{(h)}$ that minimizes the SSR.
- *Step 4 (Re-Estimation)*: Given the optimal weights $\bar{\beta}_k^{(h)}$, re-estimate $\pi_r^{(h)}$, $\gamma_k^{(h)}$, and $a^{(h)}$ using OLS to obtain the final parameter estimates.
- *Step 5 (Model Selection)*: Repeat steps 1-4 for each forecast horizon $h = 1, 2, \ldots$, and for varying numbers of ANN nodes $k = 1, \ldots, K$, and use a relevant information criterion such as the Akaike Information Criterion (AIC) or the Bayesian Information Criterion (BIC) to determine the optimal model structure.

Consistency of Model Selection

An important practical aspect of the estimation procedure set out here concerns the selection of the number of hidden nodes K in the ANN architecture. The following lemma establishes that model selection via penalized likelihood criteria (such as AIC or BIC) is consistent under standard assumptions.

Lemma 1. *(Consistency of Model Selection)*
Suppose the true number of hidden nodes is finite and denoted by K^*. Assume that:

1. *The data $\{(Y_t, X_t)\}$ are stationary.*
2. *The information criteria penalize model complexity appropriately, with penalty terms growing slower than T but faster than any constant (e.g., $\log T$ as in BIC).*
3. *The true model minimizes the expected prediction error within the candidate class of models.*

Then, the selection of the number of nodes K, using penalized likelihood criteria, is consistent, that is $\mathbb{P}(\hat{K} = K^*) \to 1 \quad as \quad T \to \infty$, where $\hat{K}$ denotes the selected number of nodes.

Proof. See Appendix, Section A, Proof 1 □

Having established the consistency of model selection, we now proceed to conduct hypothesis testing on the causal structure identified by the ANN-based framework.

- *Step 6 (Causality Testing)*: Having defined the optimal model structure, for each forecast horizon h, test the null hypothesis of non-causality using relevant inference procedures, as described below.

$$\mathcal{H}_0^{(h)} : \text{The ANN component does not contribute significantly to predicting } Y_{t+h}.$$

This hypothesis implies that the coefficients associated with the ANN terms are jointly equal to zero. Let $\mathbf{\Gamma}$ denote the full parameter vector of the model, and let $\mathbf{R}$ be the restriction matrix selecting the ANN-related coefficients. Then, the hypothesis can be expressed in the linear restriction form:

$$\mathcal{H}_0^{(h)} : \mathbf{R\Gamma} = \mathbf{r}, \tag{12}$$

where $\mathbf{r}$ is a conformable vector of constants, typically zero. To test this restriction, we use the Wald statistic:

$$W = (\mathbf{R}\hat{\mathbf{\Gamma}} - \mathbf{r})'(\mathbf{R}\hat{\mathbf{\Omega}}\mathbf{R}')^{-1}(\mathbf{R}\hat{\mathbf{\Gamma}} - \mathbf{r}), \tag{13}$$

where $\hat{\mathbf{\Gamma}}$ is the estimated parameter vector, and $\hat{\mathbf{\Omega}}$ is a consistent estimate of its asymptotic covariance matrix. Under the null hypothesis $\mathcal{H}_0^{(h)}$, the test statistic W is asymptotically distributed as a chi-squared random variable with degrees of freedom equal to the number of restrictions imposed, i.e., $W \sim \chi^2_{(q)}$.

3.5 ANNs as Universal Approximators

As a universal function approximator, ANNs provide a global estimation framework that can model any highly complex, nonlinear relationship without imposing restrictive functional forms. This innovative framework enables a more precise examination of how solar storms influence BTC mining activity over different time horizons, providing deeper insights into the resilience and vulnerabilities of decentralized digital finance. In this subsection, we present the result that the ANN-based multistep causality model can uniformly approximate any continuous h-step causal forecast function over the domain $\mathcal{D}$.

Let $\{(Y_t, X_t)\}_{t\in\mathbb{Z}}$ be a stationary stochastic process, where $Y_t \in \mathbb{R}$ is a univariate response variable, and $X_t \in \mathbb{R}^d$ is a d-dimensional exogenous input. For each forecast horizon $h \in \mathbb{N}$, define the vector:

$$Z_t := (X_t, X_{t-1}, \ldots, X_{t-q}, Y_{t-1}, Y_{t-2}, \ldots, Y_{t-p}) \in \mathbb{R}^{d(q+1)+p}. \tag{14}$$

Assume that the conditional expectation function:

$$f_h(Z_t) := \mathbb{E}[Y_{t+h} \mid Z_t] \tag{15}$$

is a continuous mapping over a compact domain $\mathcal{D} \subset \mathbb{R}^{d(q+1)+p}$. Then, we have the following result.

Theorem 3.1. *(ANN Multistep Causality is a Universal Approximator)*
Under the above assumption and for any $\varepsilon > 0$, there exists a feedforward ANN with finitely many nodes $K \in \mathbb{N}$, using a non-constant, bounded and continuous activation function ϕ :

$\mathbb{R} \to \mathbb{R}$, *such that the ANN-based approximation* $\widehat{f}_h^{(K)}(Z_t)$ *satisfies:*

$$\sup_{Z_t \in \mathcal{D}} \left| f_h(Z_t) - \widehat{f}_h^{(K)}(Z_t) \right| < \varepsilon. \tag{16}$$

Proof. See Appendix, Section A, Proof 2 □

From the above theorem, since $f_h(Z_t)$ encodes the h-step-ahead predictive impact of past values of X_t and Y_t on future outcomes, it captures the underlying dynamic causal structure. The ANN, by virtue of its universal approximation capacity, can non-parametrically learn this causal relationship directly from data, without imposing assumptions.

The global approximation property of ANNs not only allows them to flexibly approximate complex functional forms, but also ensures that the estimator remains robust even when the true data-generating process is only approximately captured by the assumed ANN structure. This robustness to local misspecification is formalized in the following remark.

Lemma 2. *(Robustness to Local Misspecification)*
As a consequence of the universal approximation property, the ANN-based estimator remains asymptotically valid even under local deviations from the assumed functional form. In particular, the model misspecification is uniformly bounded if the true process lies within an ϵ*-neighbourhood of a model representable by the ANN architecture.*

Proof. See Appendix, Section A, Proof 3 □

Therefore, the additional approximation error introduced by model misspecification is uniformly bounded by ϵ, independent of the sample size T. Hence, even if the ANN is slightly misspecified relative to the true process, the estimation error introduced by this deviation becomes negligible as $T \to \infty$. Therefore, the overall estimation procedure remains consistent, and the asymptotic distribution derived under the correctly specified model still applies asymptotically.

Next, having discussed the flexibility and (local) robustness of the ANN framework, we now turn to an analysis of the asymptotic properties of the proposed estimator, including consistency and asymptotic normality.

3.6 Asymptotic Properties of the Estimator

To ensure valid inference in the proposed, novel ANN-based multistep causality framework, we examine the large-sample properties of the estimator for the parameter vector $\Gamma^{(h)}$ at each forecast horizon h. While the model is highly nonlinear, once the ANN weights $\beta_k^{(h)}$ are fixed through the estimation algorithm presented earlier, the model becomes intrinsically linear in parameters, allowing standard asymptotic arguments to be applied. The following results establish the asymptotic consistency and normality of the estimator, conditional on the fixed ANN architecture.

Consistency of the ANN Estimator

Before establishing asymptotic normality, it is necessary to prove the consistency of the ANN-based estimator. Consistency ensures that, as the sample size increases, the estimator converges in probability to the true parameter values, thus providing the foundation for valid asymptotic inference. The following lemma formally establishes that, under standard conditions, the OLS estimator of the ANN-based multistep forecasting model converges in probability to the true parameter vector.

Theorem 3.2. *(Consistency of the ANN Estimator)*

Consider the ANN multistep causality model:

$$Y_{t+h} = a^{(h)} + Y'_{t,p}\pi^{(h)} + \sum_{k=1}^{K} \gamma_k^{(h)} \phi(X'_{t,q}\beta_k^{(h)}) + u_{t+h}, \tag{17}$$

where the ANN weights $\{\beta_k^{(h)}\}$ are fixed after random initialization and optimization. Let us suppose that the following conditions hold:

1. *The process $\{(Y_t, X_t)\}$ is stationary.*
2. *The matrix $\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]$ exists and is positive definite.*
3. *The residuals u_{t+h} satisfy $\mathbb{E}[u_{t+h} \mid \mathcal{F}_{t-1}] = 0$ with finite second moment, where $\mathcal{F}_{t-1}$ denotes the information set at time $t-1$.*

Then, as $T \to \infty$, the OLS estimator $\hat{\Gamma}^{(h)}$ satisfies $\hat{\Gamma}^{(h)} \xrightarrow{p} \Gamma^{(h)}$.

Proof. See Appendix, Section A, Proof 4 □

Asymptotic Normality of the ANN Estimator

Building on the established consistency of the estimator, we will now derive its asymptotic distribution. Under standard regularity conditions, the ANN-based estimator satisfies the central limit theorem, converging in distribution to a multivariate normal law with a well-defined covariance matrix, thus enabling standard inference.

Theorem 3.3. *(Asymptotic Normality of the ANN Estimator)*

Let $\{Y_t, X_t\}$ be a stationary stochastic process, and consider the ANN-based forecasting model at horizon h:

$$Y_{t+h} = a^{(h)} + Y'_{t,p}\pi^{(h)} + \sum_{k=1}^{K} \gamma_k^{(h)} \phi(X'_{t,q}\beta_k^{(h)}) + u_{t+h}, \tag{18}$$

where the ANN weights $\beta_k^{(h)}$ are treated as fixed. Define the regressor vector:

$$\mathcal{Z}_t^{(h)} = \left(1, Y'_{t,p}, \phi(X'_{t,q}\beta_1^{(h)}), \ldots, \phi(X'_{t,q}\beta_K^{(h)})\right)', \tag{19}$$

and let $\Gamma^{(h)} = \left(a^{(h)}, \pi^{(h)\prime}, \gamma_1^{(h)}, \ldots, \gamma_K^{(h)}\right)'$. Suppose:

1. *$\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]$ exists and is positive definite.*
2. *The error term u_{t+h} is a sequence with finite second moment.*

Then, as $T \to \infty$, it holds that:

$$\sqrt{T}(\hat{\Gamma}^{(h)} - \Gamma^{(h)}) \xrightarrow{d} \mathcal{N}(0, \Omega^{(h)}), \tag{20}$$

where

$$\Omega^{(h)} = \left(\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]\right)^{-1} \mathbb{E}[u_{t+h}^2\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}] \left(\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]\right)^{-1}. \tag{21}$$

Proof. See Appendix, Section A, Proof 5 □

This result guarantees that, despite the underlying highly nonlinear structure induced by the ANN components, the OLS estimator of the linear parameters attains asymptotic normality under standard regularity conditions, thereby enabling valid inference and the construction of asymptotically correct confidence regions. In econometric terms, this result formally establishes that, conditional on the fixed network architecture, the estimator is asymptotically normal with a well-defined covariance matrix, thereby justifying the application of standard asymptotic inference procedures.

Below, we extend the previous result by allowing the activation function to vary across a broader family of smooth transformations. This will ensure that the asymptotic properties of the ANN-based estimator continue to hold beyond the logistic case.

3.7 Robustness of Asymptotic Properties

While the previous subsection established the consistency and asymptotic normality of the ANN estimator under a fixed activation function, like the logistic sigmoid, it is crucial to verify that these large-sample properties continue to hold when the activation function varies. This generalization is particularly relevant for our empirical framework, as we later perform robustness checks, using an alternative but equally popular activation function, i.e. the hyperbolic tangent function $\tanh(z)$. The following result demonstrates that the ANN estimator retains its asymptotic validity under any activation function, satisfying standard smoothness and boundedness conditions, thereby strengthening the theoretical foundation of the robustness analysis in the Online Appendix.

Theorem 3.4. *(Asymptotic Validity of the ANN-based estimator under Smooth Activation Functions) Let $\{Y_t, X_t\}$ be a stationary time series process, and consider the ANN-based model with activation function $g : \mathbb{R} \to \mathbb{R}$. Suppose the following:*

(i) The activation function $g(\cdot)$ is continuous, bounded, and differentiable;

(ii) The model is correctly specified or the true conditional mean function can be approximated arbitrarily well by the ANN model using $g(\cdot)$;

(iii) The regressors have finite second moments, and the error u_{t+h} satisfies $\mathbb{E}[u_{t+h}^2] < \infty$;

(iv) The matrix of regressors used in estimation satisfies full rank conditions.

Then the ANN estimator $\hat{\Gamma}^{(h)}$, obtained by minimizing the sum of squared residuals using $g(\cdot)$, and defined as the solution to:

$$\hat{\Gamma}^{(h)} = \arg\min_{\Gamma} \frac{1}{T} \sum_{t=1}^{T} \left(Y_{t+h} - \mathcal{Z}_t^{(h)\prime} \Gamma^{(h)} \right)^2 \tag{22}$$

is consistent and asymptotically normal:

$$\sqrt{T} \left(\hat{\Gamma}^{(h)} - \Gamma^{(h)} \right) \xrightarrow{d} \mathcal{N}(0, \Omega^{(h)}), \tag{23}$$

where

$$\Omega^{(h)} = \left(\mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}] \right)^{-1} \mathbb{E}[u_{t+h}^2 \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}] \left(\mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}] \right)^{-1}. \tag{24}$$

Proof. See Appendix, Section A, Proof 6 □

This formal mathematical result confirms that the ANN estimator retains its desirable asymptotic properties across a broad range of smooth activation functions. This generalization formally justifies the robustness analysis in the Online Appendix, where we employ the hyperbolic tangent function $\tanh(z)$ as an alternative to the logistic sigmoid. Since both functions satisfy the required conditions of smoothness, boundedness, and differentiability, inference remains valid regardless of the specific transformation used.

However, to complement the asymptotic normality result established above, in Theorem 3.4, it is crucial to ensure that the second-order approximation, used in the Taylor expansion, is well-behaved. More precisely, we need to verify that the sample Hessian of the least squares (LS) objective function converges to its population analogue. The following theorem establishes this result.

Lemma 3. *(Convergence of the Sample Hessian on the ANN-based regressor structure) Let $\mathcal{Z}_t^{(h)}$ be the vector of regressors defined in the ANN-based model:*

$$Y_{t+h} = \mathcal{Z}_t^{(h)\prime}\Gamma^{(h)} + u_{t+h}, \tag{25}$$

where $\mathcal{Z}_t^{(h)}$ includes smooth bounded nonlinear transformations such as $g(X_t^\top \beta_k)$, with $g(\cdot)$ continuous and differentiable. Suppose the following assumptions hold:

(i) $\{\mathcal{Z}_t^{(h)}\}$ *is stationary;*

(ii) $\mathbb{E}[\|\mathcal{Z}_t^{(h)}\|^2] < \infty$;

(iii) The regressor matrix $\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]$ *is positive definite.*

Then the sample Hessian of the least squares (LS) objective function:

$$\nabla^2 Q_T(\Gamma) = \frac{2}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime} \tag{26}$$

converges in probability to its population counterpart:

$$\nabla^2 Q_T(\Gamma) \xrightarrow{p} 2\,\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}] \quad \text{as } T \to \infty. \tag{27}$$

Proof. See Appendix, Section A, Proof 7 □

4 Empirical Analysis

4.1 Data and Variables

The *Dst* index is widely recognized as a critical measure to quantify solar storms, which result from the acceleration of charged particles in space and lead to significant disruptions in Earth's magnetosphere. As noted by Temerin and Li (2006), the *Dst* index effectively captures the influence of large-scale electric currents induced by solar weather fluctuations, making it an invaluable tool for understanding solar-terrestrial interactions. Its ability to track the intensity and evolution of magnetic disturbances establishes it as a reliable proxy for identifying and analysing solar storms.

Moreover, researchers such as Mursula et al. (2011) and Daglis et al. (2019) in relevant studies in the field have emphasized its role in studying dynamic solar processes, as it directly reflects variations in solar activity and their terrestrial impacts. Beyond its geophysical origins, the Dst index has become increasingly important for interdisciplinary research. More precisely, prior studies, including Temerin and Li (2006), highlight its relevance for assessing terrestrial disruptions such as power outages, communication failures, and disturbances in satellite-based infrastructure—all of which are foundational to the stability of modern financial systems and digital networks. As a time-varying indicator of storm intensity, the Dst index offers a valuable tool for modelling the real-time impact of solar activity on economic and technological systems. Given the recurring potential of solar activity to disrupt critical systems, the Dst index is the most reliable and widely accepted metric to monitor and analyze solar storms and geomagnetic storms (Sierra-Porta et al., 2024).

The data for the Dst index were obtained from NASA's Space Physics Data Facility[2], while the data on BTC transaction volume were retrieved from Yahoo Finance[3]. Specifically, the data are daily and span the period from 17/09/2014 to 10/01/2023. More precisely, the data for the Dst index were originally provided in an hourly format and were converted to a daily format using the average daily value. Working at daily frequency enables us to precisely align the timing of geomagnetic events with observable fluctuations in BTC transactions and blockchain confirmation rates.

Table 1 presents the descriptive statistics for Dst index and the BTC volume, while Figure 1 depicts the daily values of the Dst index and Figure 2 illustrates the evolution of daily BTC transactions along with the days on which geomagnetic storms were observed. As noticed, the Dst index takes both positive and negative values. Notably, we should point out that the more negative values the Dst index takes, the more intense the solar and geomagnetic storms are (Alves et al., 2006).

Table 1: Descriptive statistics

Variable	**Obs.**	**Mean**	**SD**	**Skew**	**Kurt**	**p25**	**p50**	**p75**
BTC Volume	3038	$1.633e^{10}$	$1.994e^{10}$	2.748	30.428	$1.023e^{08}$	$7.553e^{09}$	$2.787e^{10}$
Dst	3038	-9.37	13.502	-1.75	10.038	-16	-7	0

The descriptive statistics presented in Table 1 offer a first glimpse into the markedly different statistical properties of the two central variables under investigation. The distribution of the BTC transaction volume exhibits substantial right skewness and excess kurtosis, indicating the presence of extreme values and strong tail behaviour, which is an expected characteristic in high-frequency digital asset markets. In contrast, the Dst index displays a negative mean and is left-skewed, confirming that extreme solar storms are relatively rare but impactful when they do occur. These statistical features point to the inherent volatility in digital financial activity and the sporadic but intense nature of space-weather phenomena. In this context,

[2]https://omniweb.gsfc.nasa.gov/
[3]https://finance.yahoo.com/

they provide an empirical motivation for adopting a nonlinear, flexible modelling framework capable of capturing asymmetric responses and fat-tailed behaviours, such as the Artificial Neural Network (ANN) structure developed in this study.

Figure 1: Daily values of the Dst Index from September 17, 2014 to January 10, 2023

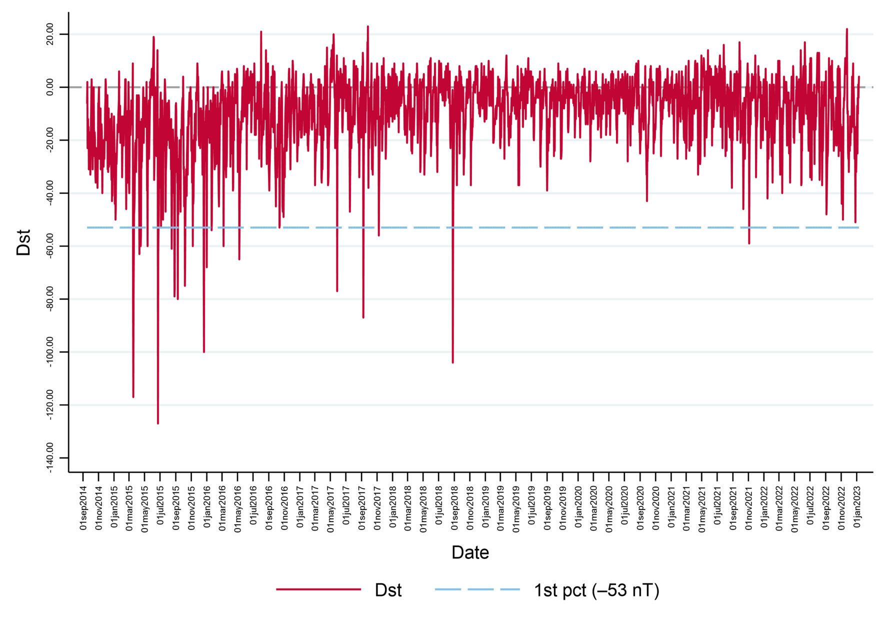

Figure 2: Daily values of the BTC transaction volumes from September 17, 2014 to January 10

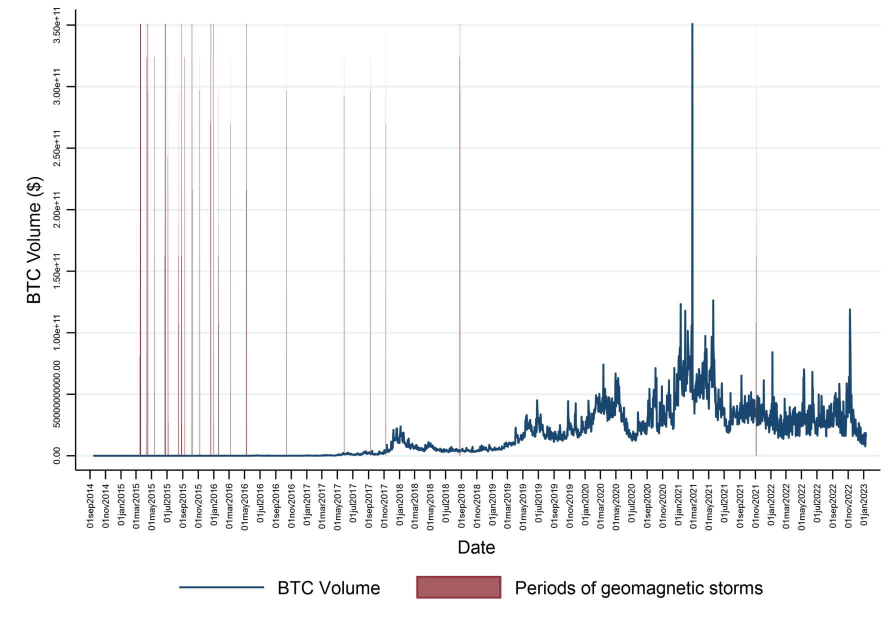

Note: The periods of geomagnetic storms correspond to days on which the Dst index is ≤ -53 nT (1st percentile) (see Sabeha et al., 2025).

Figure 1 excellently illustrates the dynamic evolution of solar storms over the observed period, with spikes marking the onset of intense geomagnetic disturbances, observed primarily during the early years of our sample (2014–2018). These intense, negative dips in the Dst index coincide with periods of heightened solar activity. The underlying time series reveals the sporadic yet severe nature of solar storms, reinforcing the hypothesis that such external shocks can severely disrupt systems that rely on continuous energy and connectivity, such as Bitcoin mining. From an economic and financial perspective, this visual evidence stresses the importance of real-time monitoring and predictive modeling to manage the vulnerabilities of digital infrastructures to exogenous shocks.

4.2 Empirical Results

In the following, Tables 4–6 present the estimated coefficients of our model across all examined horizons, while Tables 8–10 below provide the optimal values for the weight vector $\beta_k^{(h)}$.

In what follows, we examine the exact time pattern of the causal relationship between solar storms and BTC transaction volumes by employing the state-of-the-art step-by-step neural network causality test for horizons $h = 1, \cdots, 30$. The selected period range fully captures the time frame during which a solar storm may affect the Earth's geomagnetic field, potentially disrupting electricity grids and telecommunications systems, both of which are crucial for

BTC transactions and, hence, BTC mining.

According to our empirical results presented in Table 2, we notice that a very strong and statistically significant causal relationship between solar storms (expressed through the Dst index) and the process of Bitcoin (BTC) mining (expressed through the BTC transaction volume) exists across all examined horizons ($h = 1$ to $h = 30$). Specifically, the p-values are very low, and the F-statistics remain high, which confirms the rejection of the null hypothesis of no causality. For longer horizons ($h = 20$ to $h = 30$), the causal link persists, implying persistent effects of solar activity in BTC mining.

The fact that the causality strongly persists across multiple time horizons indicates that the effects of solar storms on BTC mining are not just immediate but can also have delayed implications and consequences. This finding provides a deeper understanding of the solar storms impact on the cryptocurrency markets, and by extension, on the contemporary broader digital financial system, an aspect that has been overlooked in existing research.

More precisely, an important temporal asymmetry emerges in the dynamic causal structure: the transmission of solar storms to BTC mining activity is notably stronger during the first half of the monthly forecasting horizon (i.e., h=1 to h=15). This front-loaded response likely reflects the immediate technological and operational vulnerabilities triggered by solar storm events. Specifically, in the early stages following a solar storm, critical infrastructure, such as mining rigs, power systems, and network nodes, faces voltage fluctuations and electromagnetic interference, which impairs computational throughput and blockchain synchronization. From an economic standpoint, the early-month period may also coincide with energy procurement cycles, mining difficulty adjustments, or liquidity constraints faced by smaller mining operators, amplifying sensitivity to external shocks. As time progresses, systems may adjust through tools like automatic backups, shifting workloads, or deliberately reducing activity, which can help reduce the impact. This diminishing pattern across horizons highlights the system's capacity for short-term absorption of shocks but also points to its increased vulnerability in the immediate aftermath of solar storms.

In more detail, BTC mining activity reacts immediately after a solar storm (short horizons, $h = 1$ to $h = 10$). Block confirmations may be delayed due to network instability or power outages, resulting in disruptions in trading and mining. These disruptions can, in turn, cause fluctuations in BTC volume and impact the overall performance of the cryptocurrency market.

In the medium term ($h = 11$ to $h = 20$), the effects of solar activity on BTC mining continue to be felt as power grid instability and internet disruptions remain a factor, while the effect of solar storms on BTC mining is still visible in the long run ($h = 21$ to $h = 30$), with transaction volume changing noticeably even weeks after the initial event.

Hence, the structure of our model aligns with the fact that a solar storm may affect Earth's geomagnetic field for several hours to multiple days (Mishra & Teriaca, 2023; Piersanti et al., 2020).

Table 2: Multi-step Causality Results

h	**F-stat**	**p-value**	**Causality**	h	**F-stat**	**p-value**	**Causality**
1	11.58	0.0007	Yes	2	23.10	0.0000	Yes
3	14.39	0.0002	Yes	4	20.88	0.0000	Yes
5	16.95	0.0000	Yes	6	15.46	0.0001	Yes
7	19.27	0.0000	Yes	8	22.87	0.0000	Yes
9	13.78	0.0002	Yes	10	16.34	0.0001	Yes
11	11.49	0.0007	Yes	12	15.40	0.0001	Yes
13	17.30	0.0000	Yes	14	14.71	0.0001	Yes
15	8.28	0.0040	Yes	16	10.96	0.0009	Yes
17	8.49	0.0036	Yes	18	13.32	0.0003	Yes
19	5.40	0.0202	Yes	20	7.87	0.0051	Yes
21	4.64	0.0313	Yes	22	3.91	0.0480	Yes
23	9.80	0.0018	Yes	24	5.07	0.0245	Yes
25	6.69	0.0097	Yes	26	7.69	0.0056	Yes
27	7.99	0.0047	Yes	28	9.06	0.0026	Yes
29	12.77	0.0004	Yes	30	10.59	0.0012	Yes

The delayed economic consequences from such events can be pinpointed to increased electricity costs in the affected regions, the need to replace damaged ASIC miners and power unit equipment and the requirement of regulatory or operational adjustments in order to avoid future risks. The duration of this impact highlights the importance of resilient mining infrastructure and those adaptive measures aimed at overcoming disruptions brought about by solar storms.

Based on these findings, Table 3 presents the yearly distribution of geomagnetic disturbances, corresponding to days on which the Dst index falls below -53 nT (1st percentile)[4]. As we notice, of the 33 such events during the examined period, 30 occurred between 2015–2017, with 21 of them taking place in 2015 alone. Thereafter, only sporadic events appear, with no occurrences between 2019–2020 and only one in 2021, indicating that geomagnetic activity dropped significantly.

Table 3: Days of geomagnetic disturbances per year

	Year									
	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Days	0	21	6	3	2	0	0	1	0	0

At the same time, as shown in Figure 2, these events predominantly occurred during a period when the cryptocurrency market in general, and BTC in particular, were still in the early

[4]Following Sabeha et al. (2025), Dst index values lower than -50 nT indicate a moderate storm, while values below -100 nT correspond to a strong storm.

stages of development. Trading volumes were extremely low compared to the heightened levels observed after 2018, when BTC achieved broader adoption and market participation expanded significantly.

Hence, to address the concern that storm days are concentrated in BTC's early adoption period, when BTC's market size and trading activity were substantially lower, and to avoid any time-period bias arising from this early low-volume regime, we focus our subsequent analysis on the period 2015–2017.

The results of the two-sample t-test (Table 4) reveal a substantial and statistically significant difference in BTC trading volume between days with normal geomagnetic conditions (non-storm days) and days with geomagnetic storms. On non-storm days, average trading volume is 238,588 BTC, whereas on storm days it falls to 170,087 BTC, a reduction of 68,501 coins. This difference is consistent with the previous findings that intense geomagnetic disturbances reduce market activity and hence BTC mining. For instance, as illustrated in Figure 3, during 2015, where we notice the majority of geomagnetic storms, the cumulative reduction in trading activity amounts to approximately 3 million BTC, which corresponds to an estimated value of around $883.4 million in losses.

Table 4: t-test of BTC trading volume on Storm vs. Non-Storm days (2015–2017)

Group	**Obs.**	**Mean**	**SD**	**Std. Err.**
0	1,066	238,588	213,221	6,531
1	30	170,087	154,877	28,277
Difference (0-1)		68,501		39,225
t-statistic	1.7464			
p-value (one-sided, Diff >0)	0.0405			

Note: 1 indicates days with Dst ≤ -53 (Storm days).

In addition, Figures 4–6 illustrate the dynamic response of BTC trading volume under three different scenarios of extreme geomagnetic storms that differ in the magnitude of the imposed Dst shock. According to Zhang et al. (2007), between 1996–2005 there were 18 severe storms with Dst value below -200 nT, while Gopalswamy et al. (2005) report that the strongest geomagnetic storm during that period occurred on 20 November 2003, reaching a Dst value of -472 nT. Moreover, the geomagnetic storm of 14 March 1989 recorded a Dst value of -589 nT (Temerin & Li, 2006), whereas the Carrington storm of 1–2 September 1859 may have reached an equivalent Dst value below -1500 nT (Tsurutani et al., 2003).

Within this framework, the first scenario examines a geomagnetic shock of -200 nT, the second scenario applies a shock of -250 nT, and the third scenario evaluates an even more severe storm with a shock of -300 nT.

Figure 3: Cumulative BTC losses (in units) in 2015

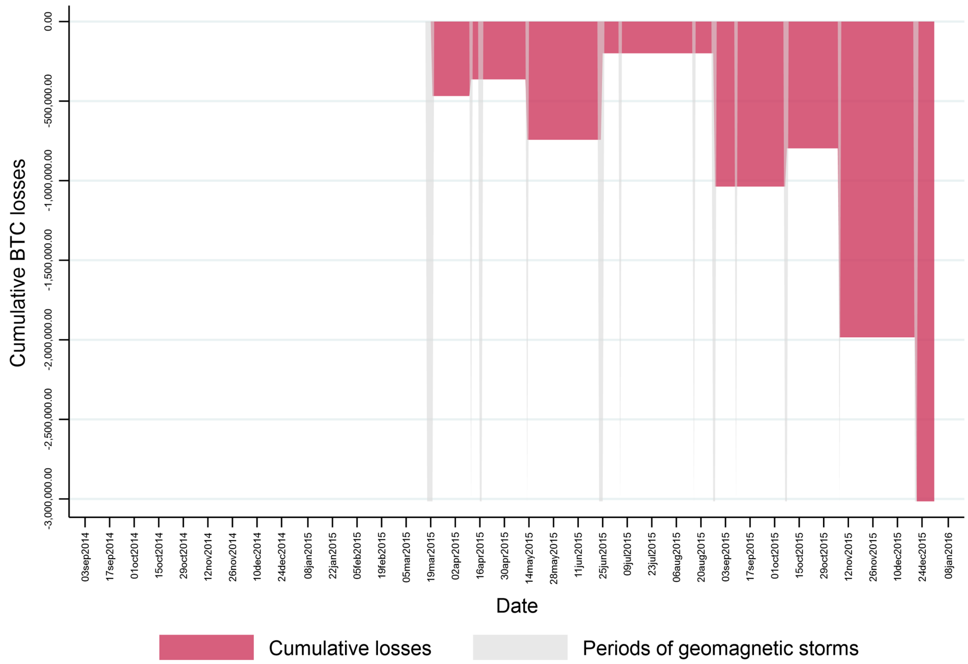

In the first scenario (Figure 4), where a -200 nT Dst shock is applied, the simulated responses reveal a consistently negative effect on the BTC trading volume. In the shorter-horizon specifications (Horizons 5, 10 & 15) the BTC volume faces a mild decline which gradually intensifies with BTC volume exhibiting a sharper fall, leading to a reduction in BTC trading volume by approximately \$20–30 billion. Although this negative effect fades out in longer-horizon specifications, the results reveal that even in the relatively mild scenario of an extreme-storm, the BTC activity experiences an intense decline.

In the second scenario (Figure 5), where the magnitude of the geomagnetic disturbance increases to -250 nT, the BTC volume exhibits an even more pronounced reduction. Within the shorter-horizon specifications, the initial decline observed in Scenario 1 is amplified, with the drop in BTC activity occurring more rapidly and with greater intensity, reaching approximately \$30–40 billion. In the longer-horizon specifications, the negative effect diminishes as the market adjusts to this external shock.

Finally, in the third scenario (Figure 6), where the geomagnetic disturbance intensifies further to -300 nT, the resulting impact becomes even more substantial, with the decline deepening more sharply than in the previous two scenarios. The reduction in BTC activity reaches \$40–50 billion, demonstrating the increasingly disruptive nature of stronger geomagnetic shocks.

Figure 4: Dynamic response of BTC trading volume to a −200 nT Dst shock

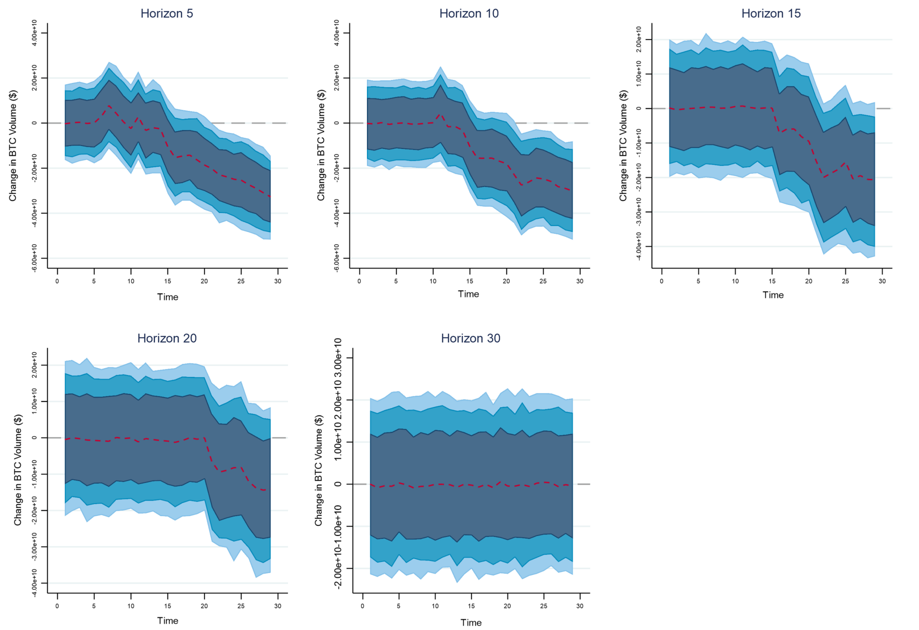

Note: The GIRFs were computed by employing a linear approximation of the proposed ARDL step-by-step ANN-based causality test. The shaded areas show (from darkest to lightest) the 75%, 90% and 95% confidence intervals, respectively.

Figure 5: Dynamic response of BTC trading volume to a −250 nT Dst shock

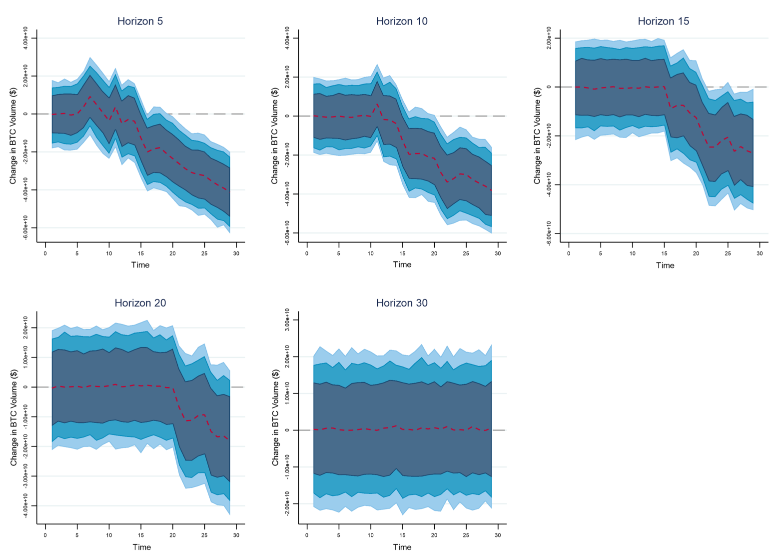

Note: The GIRFs were computed by employing a linear approximation of the proposed ARDL step-by-step ANN-based causality test. The shaded areas show (from darkest to lightest) the 75%, 90% and 95% confidence intervals, respectively.

Figure 6: Dynamic response of BTC trading volume to a −300 nT Dst shock

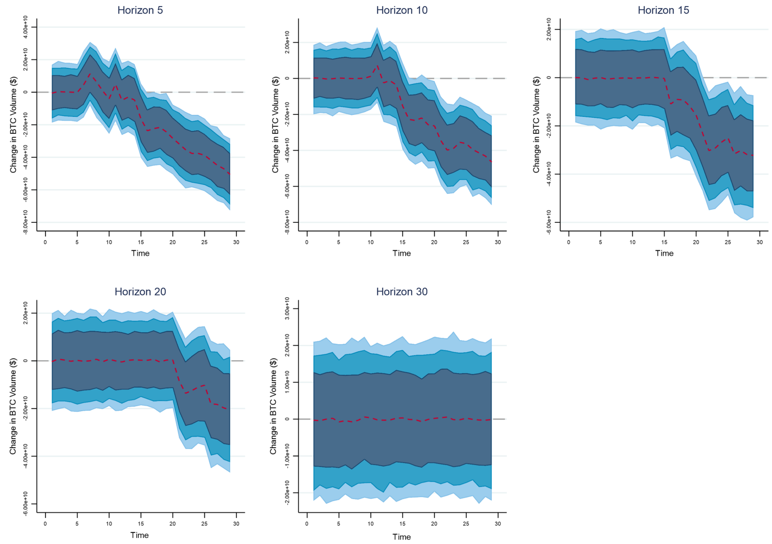

Note: The GIRFs were computed by employing a linear approximation of the proposed ARDL step-by-step ANN-based causality test. The shaded areas show (from darkest to lightest) the 75%, 90% and 95% confidence intervals, respectively.

Taken together, the three scenarios provide clear evidence that BTC activity is highly sensitive to geomagnetic disturbances, as BTC volume consistently declines following a negative Dst shock. While the −200 nT shock already produces a sizeable drop in BTC activity, the −250 nT and −300 nT scenarios demonstrate progressively larger and more rapid reductions. Although the negative effects tend to moderate at longer horizons, the immediate and medium-run declines in BTC volume remain substantial across all cases. Overall, our results suggest that extreme geomagnetic storms can significantly suppress BTC market, underscoring the relevance of space-weather dynamics as a non-traditional but meaningful source of exogenous stress for cryptocurrency markets.

To conclude, having revealed a causal relationship between solar storms and BTC mining, the different responses of BTC volume across the various horizon specifications may indicate that while some mining facilities recover rapidly from disruptions caused by solar storms, others experience prolonged recovery periods. This could be attributed to location/geographical factors, infrastructure resilience and reliance on secondary power sources. Moreover, difficulties in predicting the transaction volume of BTC during these periods may also be attributed to the temporary ménage of mining operations or deliberate decrease of computational power as miners take precautionary actions to safeguard their infrastructure from possible damage. Finally, to assess the robustness of the novel ANN-based causality model, we re-estimate the model using an alternative activation function, namely, the hyperbolic tangent function $\tanh(z)$. See the ***Online Appendix***. This robustness check confirms that the ANN-based causality results are not sensitive to the specific choice of activation function, thereby strengthening the credibility of our findings.

5 Policy Implications and Discussion

The results show that an Artificial Neural Network (ANN)-based causality test is very suitable for analysing the effect of solar storms on BTC mining. While the various geomagnetic events do not correlate with equal mining disruptions, the past few years have seen unprecedented sensitivity due to the fragile nature of these new digital financial instruments like BTC. With the nonlinear interactions in mind, an ANN framework enables us to present situations where the increase of solar storm intensity amplifies the solar storm consequences. In this framework, the proposed state-of-the-art ANN causality model enables us to accommodate the complexity, reveal the complex relationships, and also recognize instances where a slight variation in solar activity has noteworthy effects in contemporary cryptocurrency markets.

The multi-step ANN causality test indicated that effects in the short term ($h = 1$ to $h = 10$) were immediate and clear, as BTC transaction volume responded quickly. At the same time, the medium-term effects ($h = 11$ to $h = 20$) showed the different ways of recovery among mining operations, which highlights ANN's ability to capture the adaptive behaviour, geographic diversity and availability of backup power among the various mining infrastructures. In addition, the long-term effects ($h = 21$ to $h = 30$) were still significant, suggesting that some economic and operational consequences of solar storms linger for weeks after the event. This shows why ANNs are more suitable for handling relationships characterized by delayed feedback loops, and memory effects, such as the duration it takes for mining infrastructures to recover from a power deficiency, how long it takes for the network to become stabilized after long breaks in connectivity, and how do miners adjust their behaviour when energy prices respond to these unpredictable situations.

Additionally, the ANN model effectively captures volatility clustering and asymmetry, where minor solar storms may cause negligible disruptions, while extreme events trigger sharp, nonlinear reactions. Mining disruptions are neither uniformly distributed nor equally severe; they may be minimal in some regions while catastrophic in others. Furthermore, the economic and market responses to these disruptions exhibit nonlinear dynamics, as a sudden reduction in hash power can lead to significant fluctuations in Bitcoin prices, which, in turn, influence mining incentives and overall market stability. In this context, ANN specification is able to pick up how responses differ in various contexts and severity, in contrast to traditional linear models, which incorrectly assume constant effect sizes. Hence, the proposed state-of-the-art Step-by-Step Neural Network Causality Test proved to be a very strong tool which can help us to understand complex, nonlinear, dynamic relationships between variables in financial and technological systems.

Therefore, since BTC has been shown to exhibit volatility spillover effects on related assets, thereby influencing overall financial stability (Elsayed et al., 2022), this study aimed to address a critical yet overlooked factor in recent research. By employing a novel causality test, we highlight the potential impact of solar storms on the stability of the modern digital financial system, underscoring the need for further exploration of external environmental risks in the context of digital finance. Fluctuating BTC price, accompanied by high volatility (Alexander and Heck, 2020), due to solar storms, can lead to higher uncertainty and present significant difficulties for financial institutions and investors. The potential volatility spillovers to other financial instruments, especially during such events, highlight the ability to manage risks in the other financial instruments in the broader financial system (Elsayed et al., 2022). In addition, it points to the existence of different channels of interconnectedness between solar storms, BTC and other markets (Lin et al., 2021; Zhang et al., 2021) that may produce contagion effects, especially in the presence of market risks and economic turbulence. Nowadays, BTC is almost completely intertwined with the traditional financial system; such disruptions in its volume of transactions have far-reaching consequences, which lead to financial volatility (Maghyereh and Abdoh, 2020; Elsayed et al., 2022). To address uncertainty in financial infrastructure, one can construct “robust” worst-case networks consistent with limited data: such models reveal sparse critical interbank links (often between banks with highly correlated asset portfolios) that maximize contagion impact. Identifying these worst-case networks provides regulators a practical tool for systemic-risk monitoring and helps target interventions to bolster network resilience under ambiguous conditions (Hu et al., 2024). Therefore, policymakers and central bankers should pay a lot of attention to the risks that BTC’s financial characteristics can pose to financial stability in the global financial arena, in a constantly changing digital finance era. Consistent with the above, evidence shows that cryptocurrency assets under robust regulatory oversight attain high efficiency levels that are comparable to traditional markets. In the meantime, noncompliant assets exhibit markedly higher inefficiencies and investor risk. This clearly suggests that strengthening compliance and oversight can materially enhance market stability in the face of rare shocks (Nimalendran et al., 2025).

The findings of this study have important implications for decision-makers in the digital finance sector. First, the strong and persistent impact of geomagnetic storms on BTC mining activity showcases the need for enhanced operational resilience among firms dependent on energy-intensive computational processes. In this context, managers of mining operations should consider diversifying geographical locations, adopting dynamic load-balancing strategies, and investing in solar-shielded hardware or backup power systems in order to mitigate exposure to space-weather risk. Meanwhile, solar storm disruptions to BTC miners and data centre operators can be proactive steps to mitigate their risks. One of the key strategies is early warning and responsive operations, since one can take advantage of the fact that space weather forecasts exist. Solar storms are detected by agencies such as NOAA’s Space Weather Prediction Center, and they issue alerts (G1 to G5 scale). These alerts should be integrated into the operational planning of miners. For example, during the prediction of a severe G4/G5 geomagnetic storm, a mining farm may decide to secure essential equipment and disconnect the sensitive hardware from the grid to prevent surge damage. This could allow a controlled shutdown that would be, possibly for a few hours or days, to prevent the seemingly forced outages and equipment failures. Much like having an emergency plan for hurricanes or earthquakes, protocols in place (preparedness) decrease how much scrambling is involved when a solar storm does strike.

Second, financial institutions and insurers underwriting crypto-assets or blockchain-dependent services may need to incorporate space-weather indices, such as the Dst or solar flare intensity, as exogenous risk factors in stress-testing and smart contract design. Given the observed lagged responses in network hash rates and transaction volumes, strategic hedging mechanisms could play a vital role in preventing disruptions across digital finance systems. Electrical hardening of mining facilities could also be an important long-term measure. This includes installing robust surge protectors, lightning arrestors and GIC blockers on the facility's grid connections. Miners can't protect the whole power grid, but they can install devices on their own substations and transformers to dissipate induced currents. Uninterruptible power supplies (UPS), and power conditioners can filter out voltage spikes or frequency deviations, and this is what mining farms do to ride through minor fluctuations. UPS battery backup or diesel generators should power critical systems not only during normal outages but also to support an orderly shutdown in case the grid drops. Essential control electronics can be shielded from sudden electromagnetic pulses by building them inside enclosures of Faraday cages or electromagnetic shielding. For example, shielded racks could be used for network routers, mining pool servers, and monitoring PCs.

Moreover, regulators and grid operators may need to work closely with the digital asset industry to ensure that coordinated contingency planning is in place. As decentralized financial technologies scale further, the systemic relevance of physical infrastructure shocks like geomagnetic disturbances will grow, underscoring the importance of interdisciplinary risk assessment and real-time monitoring. Finally, the geographic distribution of mining farms could be an essential measure. Bitcoin mining is distributed across many countries and latitudes. Because this decentralization exists, it is highly unlikely that a solar storm is going to wipe out all the miners. Some will go down in high-risk zones, but others will keep going. This is something an individual mining company can do as well, by diversifying its locations, running some hash power in North America, and some in Southeast Asia or South America, to hedge against a regional space weather event.

6 Conclusion

As BTC continues to emerge as a central player in the global economy, its influence on the broader financial system has grown rapidly. However, the increasing interconnectedness between BTC and other financial operations has made it particularly vulnerable to external shocks, including environmental disruptions such as solar storms. These risks, if not properly understood and mitigated, can amplify systemic vulnerabilities and propagate instability across financial markets. While previous research has explored the integration of cryptocurrencies with conventional financial markets, the specific effects of external environmental factors on BTC remain underexplored. This research addresses this gap by investigating how solar storms impact BTC trading volumes in both normal and volatile market conditions.

This study has introduced the Step-by-Step ANN-based Causality Test, providing a novel framework for understanding complex and dynamic relationships in financial systems. By incorporating ANNs, the proposed methodology went beyond traditional causality analysis, capturing horizon-specific insights and addressing non-linear interactions, feedback loops, and delayed effects. Developed for the case of solar storms and their impact on BTC mining, this approach has revealed significant and heterogeneous effects across different mining intensities, highlighting vulnerabilities in the cryptocurrency ecosystem. These findings stress the systemic risks posed by environmental shocks and emphasize the need for proactive policies and robust risk management strategies to safeguard financial stability. The study also demonstrated the broader potential of AI-driven methods in addressing the challenges of the digital economy.

Moving forward, further studies could explore additional, yet overlooked, issues. Research focusing on network effects within interconnected financial assets, the cascading impact of energy and telecommunications disruptions, and the broader implications for financial stability would provide critical insights. Moreover, understanding the regulatory and policy measures needed to manage these risks is essential, particularly as the adoption of digital finance continues to accelerate. Future research could also analyse the role of investor sentiment in amplifying volatility and spillovers in the cryptocurrency market. Investigating how central banks and regulators can address the effects of BTC fluctuations on other financial products would help craft targeted strategies to contain systemic risks. Examining the interactions between BTC volatility and macroeconomic factors, as well as the risk management practices used to hedge against BTC price fluctuations, could offer valuable guidance for market participants.

Appendix

Section A: Proofs

Consistency of model selection is a fundamental requirement for valid inference. In our setting, we select the number of hidden nodes $\hat{K}$ via a penalized criterion such as BIC. It is well-understood that criteria like the Bayesian Information Criterion (BIC) are strongly consistent in identifying the true model size under regular conditions. More precisely, if the true number of nodes K is finite, BIC will select the correct K with probability approaching 1 as $T \to \infty$. This result traces back to Schwarz (1978) and subsequent works and ensures that overfitted models (with $K > K$) are penalized and underfitted models ($K < K$) incur larger errors, so eventually the correct architecture is chosen. This property is analogous to the classical consistency of order selection in autoregressions (e.g. the Hannan–Quinn criterion) in simpler time-series contexts. Thus, our Lemma 1 is supported by the same intuition and prior results on consistent model selection.

Proof 1: Lemma 1. Since the data are stationary, the sample prediction error converges uniformly to its expected value. The penalty term associated with model complexity ensures that overfitted models (with $K > K^*$) are penalized sufficiently heavily in large samples. On the other hand, underfitted models (with $K < K^*$) incur larger prediction errors. Hence, as $T \to \infty$, the penalized criterion is minimized at K^*. □

Another key property is the universal approximation capability of ANNs, which characterizes the flexibility of our ANN-based causality test. The Universal Approximation Theorem states that a feedforward network with a single hidden layer can approximate any continuous function on a compact domain. This foundational result was established by several independent works, including Cybenko (1989), Hornik et al. (1989), and Funahashi (1989), who showed that using a non-constant, bounded activation (such as logistic or tanh), a one-hidden-layer ANN is dense in the space of continuous functions.

In simple terms, any smooth nonlinearity between the inputs and output can be approximated arbitrarily well by some neural networklink. This justifies our Theorem: any nonlinear forecasting function $f_h(\cdot)$ can be approximated to within ε-accuracy by a network of adequate size. Moreover, later research provided quantitative rates for this approximation.

Proof 2: Theorem 3.1. Under stationarity, the conditional expectation function:

$$f_h(Z_t) = \mathbb{E}[Y_{t+h} \mid Z_t = z] \tag{28}$$

exists and is continuous over the compact input domain $\mathcal{D} \subset \mathbb{R}^{d(q+1)+p}$. Hence, $f_h \in \mathcal{C}(\mathcal{D})$, the space of continuous real-valued functions on $\mathcal{D}$.

Now, by the *Universal Approximation Theorem* (Cybenko, 1989; Hornik et al., 1989; Funahashi, 1989), a feedforward ANN with nonconstant, bounded, and continuous activation function ϕ (e.g. sigmoid) is dense in $\mathcal{C}(\mathcal{D})$. Thus, one can show that, for every $\varepsilon > 0$, there exists a neural network of the form:

$$\widehat{f}_h^{(K)}(Z_t) = \sum_{k=1}^{K} \gamma_k^{(h)} \phi\left(\langle Z_t \rangle\right), \tag{29}$$

with $\beta_k^{(h)} \in \mathbb{R}^{d(q+1)+p}$, $\gamma_k^{(h)} \in \mathbb{R}$, such that:

$$\sup_{Z_t \in \mathcal{D}} \left| f_h(Z_t) - \widehat{f}_h^{(K)}(Z_t) \right| < \varepsilon. \tag{30}$$

Therefore, the ANN-based multistep causality model globally approximates the true forecasting function within any desired level of accuracy. □

These results imply that our ANN estimator has great flexibility and is robust to misspecification: even if the true functional form is complex, there exists an ANN that is arbitrarily close to it. Consequently, as noted in our Lemma 2, if the true data-generating process (DGP) lies within a small neighborhood of an ANN model class, the estimator will still perform well. This robustness is analogous to expanding sequence of function classes that can approximate the truth arbitrarily closely.

Proof 3: Lemma 2. Let $f_0 : \mathcal{D} \to \mathbb{R}$ denote the true but unknown data-generating process, where $\mathcal{D} \subseteq \mathbb{R}^d$ is compact, and suppose that f_0 is continuous on $\mathcal{D}$. By the universal approximation theorem (Hornik et al., 1989; Cybenko, 1989), for any given $\epsilon > 0$, there exists a feedforward ANN function $f_{\text{ANN}}(\cdot)$ such that:

$$\sup_{x \in \mathcal{D}} |f_0(x) - f_{\text{ANN}}(x)| < \epsilon. \tag{31}$$

Now, consider the observed model:

$$Y_{t+h} = f_0(Z_t) + u_{t+h}, \tag{32}$$

where Z_t collects all relevant regressors, and u_{t+h} satisfies $\mathbb{E}[u_{t+h} \mid \mathcal{F}_{t-1}] = 0$ and has finite variance. Typically, u_{t+h} denotes the error term, and $\mathcal{F}_{t-1}$ denotes the information set generated by $\{Y_s, X_s\}_{s \leq t-1}$.

Suppose that estimation proceeds under the approximating model:

$$Y_{t+h} = f_{\text{ANN}}(Z_t) + \varepsilon_t \tag{33}$$

or

$$f_0(Z_t) + u_{t+h} = f_{\text{ANN}}(Z_t) + \varepsilon_t. \tag{34}$$

Therefore:

$$\varepsilon_t = (f_0(Z_t) - f_{\text{ANN}}(Z_t)) + u_{t+h} \tag{35}$$

or

$$\varepsilon_t - u_{t+h} = (f_0(Z_t) - f_{\text{ANN}}(Z_t)). \tag{36}$$

Since f_{ANN} approximates f_0 uniformly within ϵ over $\mathcal{D}$, it follows that:

$$|\varepsilon_t - u_{t+h}| \leq \epsilon \quad \text{for all} \quad t, \tag{37}$$

and the result follows. □

With universal approximation and robustness in place, we turn to the consistency of the ANN estimator itself. Consistency means that as $T \to \infty$, the estimator converges in probability to the true parameter or function. This property has been established in earlier studies for ANN-based models. Notably, White (1990) proved that a multilayer feedforward network can learn the true regression mapping, i.e., the network weights obtained by minimizing prediction error will converge to the optimal values, provided the model is correctly specified. Even when the model is only approximately correct, the estimator targets the best approximation (in a least-squares sense) of the true function. This is a standard result in estimation theory (cf. White, 1982). In fact, White (1982) showed that maximum likelihood estimators remain consistent even under mild misspecification. In our context, because an adequately large ANN can approximate the true function, the parameters coincide with the actual truth, yielding consistency. For instance, White (1992) extended consistency results to ANN-quantile estimators for dependent (time-series), further confirming that network-based estimation retains consistency under dependence and nonlinearity. We also note that consistency of neural predictors is not limited to regression settings. Faragó and Lugosi (1993), for example, proved that a simple one-hidden-layer network can achieve strong universal consistency in classification. All these insights from simpler or related contexts lend credence to our following results, as the sample size grows, the ANN-based causality estimator, including the selected network architecture, will converge to the correct model.

Proof 4: Theorem 3.2. The OLS estimator can be written as:

$$\hat{\Gamma}^{(h)} = \left(\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \right)^{-1} \left(\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} Y_{t+h} \right). \tag{38}$$

Expanding Y_{t+h} from the model:

$$Y_{t+h} = \mathcal{Z}_t^{(h)\prime} \Gamma^{(h)} + u_{t+h}, \tag{39}$$

we have:

$$\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} Y_{t+h} = \sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \Gamma^{(h)} + \sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h}. \tag{40}$$

Thus:

$$\hat{\Gamma}^{(h)} = \Gamma^{(h)} + \left(\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \right)^{-1} \left(\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h} \right). \tag{41}$$

By Assumption 1 and the Law of Large Numbers, we have:

$$\frac{1}{T} \sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \xrightarrow{p} \mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}], \tag{42}$$

where the limiting matrix is positive definite by Assumption 2.

Moreover, under Assumption 3, the Law of Large Numbers for the errors implies:

$$\frac{1}{T} \sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h} \xrightarrow{p} 0. \tag{43}$$

Therefore:

$$\hat{\Gamma}^{(h)} \xrightarrow{p} \Gamma^{(h)}. \tag{44}$$

□

Finally, we consider the asymptotic normality of the ANN estimator. Asymptotic normality means that the estimation error, suitably scaled (typically by $\sqrt{T}$), follows an approximately normal distribution for large T. This property is crucial for constructing confidence intervals and hypothesis testing. In classical parametric models, asymptotic normality is guaranteed by regularity conditions, i.e., via the Central Limit Theorem (CLT), applied to score functions. In the context of ANNs, which are nonparametric estimators, establishing asymptotic normality is more involved, but prior research has shown it under certain smoothness conditions. A seminal result by Chen and White (1999) demonstrated that ANN estimators can achieve $\sqrt{T}$-consistency and asymptotic normality. This implies that, despite the complex, nonlinear nature of ANNs, their estimated parameters behave like those of ordinary regressions in large samples. This insight is exploited in the Theorem below.

Intuitively, once the network architecture is sufficiently rich to approximate the true function well, the remaining estimation uncertainty is due to sampling error, which averages out to normality as T grows. Our proof of asymptotic normality builds on standard techniques, such as Taylor expansions and dominance conditions. The conclusion is that the ANN-based causality test statistic will have a well-defined Gaussian limit distribution, enabling valid inference. In summary, the theoretical properties we establish, are strongly supported by previous results which have shown that ANN-based models can satisfy these properties in a variety of settings.

Proof 5: Theorem 3.3. Given the consistency of $\hat{\Gamma}^{(h)}$ established in Theorem 3.2, it suffices to derive the asymptotic distribution. By the Central Limit Theorem for stationary sequences, we have:

$$\sqrt{T}\left(\frac{1}{T}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}u_{t+h}\right)\xrightarrow{d}\mathcal{N}(0,\Sigma_u^{(h)}),\tag{45}$$

where

$$\Sigma_u^{(h)}=\lim_{T\to\infty}\operatorname{Var}\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}u_{t+h}\right).\tag{46}$$

By using the standard decomposition:

$$\sqrt{T}(\hat{\Gamma}^{(h)}-\Gamma^{(h)})=\left(\frac{1}{T}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}\right)^{-1}\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}u_{t+h}\right),\tag{47}$$

the desired asymptotic normality follows with asymptotic covariance matrix $\Omega^{(h)}$ as defined in (21). $\square$

Proof 6: Theorem 3.4. Let the regressor vector be defined as follows:

$$\mathcal{Z}_t^{(h)}=\left(1,Y_{t-1},Y_{t-2},\,g(X_t^\top\beta_1),\ldots,g(X_t^\top\beta_K)\right)'.\tag{48}$$

1. Consistency

By Assumption (ii), the ANN model with activation $g(\cdot)$ can approximate the true conditional mean arbitrarily well. Under stationarity, the Law of Large Numbers gives:

$$\frac{1}{T}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}\xrightarrow{p}\mathbb{E}[\mathcal{Z}_t^{(h)}\mathcal{Z}_t^{(h)\prime}]\quad\frac{1}{T}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}Y_{t+h}\xrightarrow{p}\mathbb{E}[\mathcal{Z}_t^{(h)}Y_{t+h}].\tag{49}$$

Hence, it follows that:

$$\hat{\Gamma}^{(h)}\xrightarrow{p}\Gamma^{(h)}.\tag{50}$$

2. Normality

Define the residual as follows: $u_{t+h}=Y_{t+h}-\mathcal{Z}_t^{(h)\prime}\Gamma^{(h)}$. The estimation problem is a nonlinear least squares (NLS) problem with fixed nonlinear components $g(\cdot)$, linear in the coefficients $\Gamma^{(h)}$:

$$Q_T(\Gamma)=\frac{1}{T}\sum_{t=1}^{T}\left(Y_{t+h}-\mathcal{Z}_t^{(h)\prime}\Gamma\right)^2.\tag{51}$$

The gradient and Hessian of the objective function are:

$$\nabla Q_T(\Gamma)=-\frac{2}{T}\sum_{t=1}^{T}\mathcal{Z}_t^{(h)}(Y_{t+h}-\mathcal{Z}_t^{(h)\prime}\Gamma),\tag{52}$$

$$\nabla^2 Q_T(\Gamma) = \frac{2}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}. \tag{53}$$

By the first-order Taylor expansion (via the multivariate mean value theorem [MVT]), there exists a point $\tilde{\Gamma}^{(h)}$ on the segment between $\hat{\Gamma}^{(h)}$ and $\Gamma^{(h)}$ such that:

$$0 = \nabla Q_T(\hat{\Gamma}^{(h)}) = \nabla Q_T(\Gamma^{(h)}) + \nabla^2 Q_T(\tilde{\Gamma}^{(h)})(\hat{\Gamma}^{(h)} - \Gamma^{(h)}). \tag{54}$$

Substituting the expressions above, we obtain:

$$0 = -\frac{2}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h} + \left(\frac{2}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}\right)(\hat{\Gamma}^{(h)} - \Gamma^{(h)}) + o_p(T^{-1/2}). \tag{55}$$

Rearranging terms and dividing by $\sqrt{T}$, we obtain:

$$\sqrt{T}(\hat{\Gamma}^{(h)} - \Gamma^{(h)}) = \left(\frac{1}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}\right)^{-1} \left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h}\right) + o_p(1). \tag{56}$$

By the central limit theorem (CLT):

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} u_{t+h} \xrightarrow{d} \mathcal{N}(0, \Sigma^{(h)}), \quad \Sigma^{(h)} = \mathbb{E}[u_{t+h}^2 \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]. \tag{57}$$

By the law of large numbers:

$$\frac{1}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \xrightarrow{p} \mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]. \tag{58}$$

Therefore, it follows that:

$$\sqrt{T}(\hat{\Gamma}^{(h)} - \Gamma^{(h)}) \xrightarrow{d} \mathcal{N}(0, \Omega^{(h)}), \tag{59}$$

where the variance matrix is:

$$\Omega^{(h)} = \left(\mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]\right)^{-1} \mathbb{E}[u_{t+h}^2 \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}] \left(\mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]\right)^{-1}. \tag{60}$$

□

Proof 7: Lemma 4. The sample Hessian of the least squares (LS) criterion function is:

$$\nabla^2 Q_T(\Gamma) = \frac{2}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}. \tag{61}$$

Based on Assumptions (i) and (ii), by the law of large numbers for stationary sequences, we have:

$$\frac{1}{T}\sum_{t=1}^{T} \mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime} \xrightarrow{p} \mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]. \tag{62}$$

Multiplying both sides by 2 yields:

$$\nabla^2 Q_T(\Gamma) \xrightarrow{p} 2\,\mathbb{E}[\mathcal{Z}_t^{(h)} \mathcal{Z}_t^{(h)\prime}]. \tag{63}$$

Assumption (iii) ensures that the limit is positive definite, which guarantees both invertibility of the matrix and uniqueness of the minimizer in the least squares (LS) problem. □

Section B: Empirical results

In Table A1, we examine the stationarity properties of the time series employing the Phillips-Perron unit root test. According to our results, both variables, i.e., BTC volume and Dst index, are found to be stationary, implying that their statistical properties remain unchanged over time. This confirms the validity of the causality test results, as non-stationary time series could have led to spurious results.

Table A1: Stationarity test results (Phillips-Perron)

Variable	**p-value**	**Stationarity**
BTC Volume	0.000	Yes
Dst	0.000	Yes

Note: The test includes a trend term

Since stationarity has been confirmed, determining the optimal lag length is important for accurately capturing the dynamic relationship between solar storms and BTC transactions, and, consequently, the BTC mining process. Our findings in Table A2 suggest that the optimal lag length is 14, as indicated by the lowest AIC value.

Table A2: Lag Length Selection

Lag	**AIC**	**Lag**	**AIC**
0	58.286	1	56.260
2	56.115	3	56.023
4	55.992	5	55.981
6	55.951	7	55.936
8	55.937	9	55.939
10	55.941	11	55.939
12	55.938	13	55.931
14	55.919*	15	55.921
16	55.922	17	55.923
18	55.923	19	55.923
20	55.924		

In what follows, Tables A3 – A5 present the estimated coefficients of our model across all examined horizons, while Tables A6 and A7 provide the optimal values for the weight vector $\beta_k^{(h)}$.

Table A3: BTC Volume - Estimated Coefficients (Horizon 1-10)

	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$	$h = 9$	$h = 10$
$L0$	0.314***	0.214***	0.243***	0.236***	0.150***	0.233***	0.267***	0.170***	0.165***	0.144***
$L1$	0.113***	0.179***	0.165***	0.072***	0.189***	0.198***	0.088***	0.112***	0.091***	0.076***
$L2$	0.145***	0.140***	0.046**	0.162***	0.182***	0.061***	0.082***	0.073***	0.057**	0.074***
$L3$	0.097***	0.015	0.128***	0.146***	0.038*	0.050**	0.034	0.032	0.049**	0.097***
$L4$	-0.014	0.105***	0.125***	0.020	0.034	0.012	0.007	0.034	0.081***	0.136***
$L5$	0.111***	0.127***	0.023	0.039*	0.013	0.006	0.038*	0.084***	0.138***	0.093***
$L6$	0.091***	-0.002	0.010	-0.016	-0.005	0.011	0.055**	0.120***	0.073***	0.011
$L7$	-0.031	-0.008	-0.037*	-0.027	-0.002	0.031	0.095***	0.059**	-0.004	0.048**
$L8$	-0.002	-0.028	-0.016	0.004	0.039*	0.105***	0.068***	-0.000	0.055**	0.059**
$L9$	-0.030	-0.014	0.007	0.038*	0.106***	0.072***	0.003	0.053**	0.063***	-0.001
$L10$	-0.004	0.013	0.044**	0.113***	0.077***	0.010	0.061***	0.066***	0.004	0.050**
$L11$	0.012	0.047**	0.118***	0.076***	0.008	0.064***	0.070***	0.004	0.054**	0.069***
$L12$	0.040**	0.121***	0.076***	0.004	0.065***	0.071***	0.002	0.052**	0.070***	0.055**
$L13$	0.106***	0.073***	-0.003	0.053**	0.068***	-0.000	0.039*	0.060***	0.052**	-0.000
$L14$	0.032*	-0.013	0.035*	0.037*	-0.009	0.028	0.039*	0.022	-0.011	0.026
Const	-1.9e+8	2.0e+8	2.2e+8	2.4e+9***	3.8e+8	4.8e+8**	3.0e+9***	6.3e+8**	6.8e+8***	2.3e+9***

Note: ***: $p < 0.01$; **: $p < 0.05$; *: $p < 0.1$

Table A4: BTC Volume - Estimated Coefficients (Horizon 11-20)

	$h = 11$	$h = 12$	$h = 13$	$h = 14$	$h = 15$	$h = 16$	$h = 17$	$h = 18$	$h = 19$	$h = 20$
$L0$	0.112***	0.116***	0.155***	0.206***	0.188***	0.140***	0.181***	0.189***	0.119***	0.147***
$L1$	0.088***	0.126***	0.176***	0.129***	0.079***	0.140***	0.139***	0.068***	0.123***	0.124***
$L2$	0.116***	0.163***	0.117***	0.056**	0.118***	0.125***	0.051**	0.105***	0.117***	0.077***
$L3$	0.149***	0.102***	0.038	0.088***	0.096***	0.031	0.079***	0.085***	0.063**	0.016
$L4$	0.092***	0.028	0.074***	0.079***	0.014	0.067***	0.071***	0.046*	0.003	0.022
$L5$	0.027	0.074***	0.077***	0.016	0.070***	0.076***	0.047*	0.005	0.022	0.028
$L6$	0.060**	0.063***	0.001	0.049**	0.054**	0.032	-0.014	0.001	0.013	0.021
$L7$	0.049**	-0.015	0.029	0.033	0.013	-0.030	-0.018	-0.005	0.010	0.037
$L8$	-0.005	0.042*	0.050**	0.021	-0.023	-0.012	0.002	0.023	0.047*	0.166***
$L9$	0.047**	0.062***	0.042*	-0.019	-0.012	0.005	0.027	0.056**	0.176***	0.059**
$L10$	0.067***	0.045*	-0.018	-0.005	0.012	0.031	0.063***	0.182***	0.059**	-0.010
$L11$	0.051**	-0.008	0.011	0.018	0.032	0.066***	0.186***	0.071***	-0.001	0.028
$L12$	0.000	0.023	0.037	0.035	0.062***	0.187***	0.071***	0.010	0.040	0.026
$L13$	0.026	0.045*	0.046**	0.058**	0.181***	0.068***	0.005	0.042*	0.032	0.079***
$L14$	0.053**	0.064***	0.096***	0.169***	0.046**	0.001	0.033	0.040*	0.095***	0.096***
Const	3.0e+9***	3.3e+9***	5.4e+8*	2.3e+9***	7.7e+8***	2.0e+9***	2.0e+9***	9.4e+8***	9.6e+8***	1.2e+9***

Note: ***: $p < 0.01$; **: $p < 0.05$; *: $p < 0.1$

Table A5: BTC Volume - Estimated Coefficients (Horizon 21-30)

	$h = 21$	$h = 22$	$h = 23$	$h = 24$	$h = 25$	$h = 26$	$h = 27$	$h = 28$	$h = 29$	$h = 30$
$L0$	0.153***	0.126***	0.086***	0.098***	0.101***	0.099***	0.121***	0.235***	0.146***	0.084***
$L1$	0.088***	0.049**	0.069***	0.072***	0.074***	0.100***	0.209***	0.082***	0.040	0.086***
$L2$	0.034	0.058**	0.061**	0.064**	0.093***	0.200***	0.073***	0.017	0.070***	0.058**
$L3$	0.037	0.045*	0.050**	0.078***	0.186***	0.059**	0.002	0.039	0.036	0.084***
$L4$	0.029	0.040	0.070***	0.176***	0.050**	-0.009	0.027	0.013	0.068***	0.082***
$L5$	0.039	0.069***	0.179***	0.050**	-0.008	0.024	0.012	0.069***	0.083***	0.162***
$L6$	0.054**	0.165***	0.040	-0.019	0.016	0.001	0.056**	0.057**	0.149***	0.077***
$L7$	0.149***	0.029	-0.027	0.005	-0.010	0.046*	0.042	0.125***	0.065**	0.004
$L8$	0.041	-0.022	0.009	-0.006	0.052**	0.052**	0.139***	0.079***	0.009	-0.017
$L9$	-0.012	0.010	-0.005	0.055**	0.058**	0.145***	0.091***	0.018	-0.015	-0.021
$L10$	0.014	-0.003	0.058**	0.061**	0.147***	0.092***	0.019	-0.010	-0.016	0.033
$L11$	0.007	0.059**	0.059**	0.149***	0.099***	0.026	0.000	-0.008	0.034	0.045*
$L12$	0.070***	0.062**	0.150***	0.099***	0.031	0.009	0.005	0.041	0.043*	0.115***
$L13$	0.063**	0.144***	0.094***	0.028	0.010	0.008	0.048*	0.042*	0.108***	0.099***
$L14$	0.152***	0.085***	0.019	0.002	0.011	0.057**	0.063***	0.108***	0.083***	0.007
Const	2.6e+9***	2.2e+9***	1.2e+9***	1.3e+9***	3.1e+9***	1.4e+9***	1.3e+9***	2.9e+9***	3.5e+9***	2.8e+9***

Note: ***: $p < 0.01$; **: $p < 0.05$; *: $p < 0.1$

Table A6: Weights in Neural Network (Horizon 1–15)

	h=1	**h=2**	**h=3**	**h=4**	**h=5**	**h=6**	**h=7**	**h=8**	**h=9**	**h=10**	**h=11**	**h=12**	**h=13**	**h=14**	**h=15**
β_0	-0.2144	-0.2711	0.1318	-0.0415	-0.2859	-0.0657	-0.1893	0.1091	-0.0381	0.047	-0.2907	-0.0664	0.1625	-0.2307	-0.0153
β_1	-0.263	-0.1003	-0.211	0.1994	0.0579	-0.278	-0.1349	0.2661	-0.0725	0.0074	-0.0134	-0.0751	0.3143	-0.129	0.1722
β_2	-0.2935	-0.0258	0.0548	0.169	0.2823	0.1617	-0.0194	-0.1361	0.0501	-0.1194	0.027	-0.2812	0.183	-0.3187	0.0286
β_3	-0.0642	-0.1701	-0.3067	-0.0999	0.0154	0.2351	0.2197	0.2194	-0.0071	-0.3103	-0.0423	0.0768	-0.0555	0.0166	0.1395
β_4	-0.1371	0.2487	0.2633	-0.2583	-0.0234	0.087	-0.1866	-0.0541	0.3124	-0.1761	0.0026	-0.2937	0.1351	0.3012	0.0477
β_5	0.2602	0.0327	0.2901	-0.1048	0.1211	0.3188	-0.2104	0.2389	0.2638	0.0631	-0.191	0.2065	-0.1151	-0.227	0.1051
β_6	0.328	0.2607	0.1681	0.047	0.1086	-0.126	-0.3208	0.1367	-0.171	-0.3206	0.2709	0.0852	0.1755	-0.1882	-0.1908
β_7	0.2713	0.0113	0.0739	-0.217	-0.0334	0.0047	-0.2228	0.223	0.1326	0.2923	-0.2664	-0.2382	-0.068	-0.1305	0.295
β_8	-0.2434	-0.1223	0.2253	-0.3204	-0.0964	0.1007	-0.3011	0.1577	0.0494	-0.0792	-0.137	-0.0394	0.1077	0.1296	-0.0346
β_9	0.1617	0.1731	-0.0034	0.2924	0.2222	0.1144	-0.2769	-0.294	-0.0262	-0.1855	-0.2609	-0.2347	-0.0594	0.2901	-0.2303
β_{10}	0.0453	0.313	0.2152	-0.2196	0.243	0.3196	-0.2196	-0.0677	0.2469	-0.2792	-0.0128	0.0571	-0.1761	-0.2837	0.2448
β_{11}	-0.0825	0.2665	0.2343	-0.1886	0.2363	-0.3038	-0.0585	0.0652	0.0342	-0.1289	-0.0431	0.0112	0.2459	-0.0849	-0.0512
β_{12}	0.2881	0.2423	-0.1735	0.0135	-0.2311	-0.0157	-0.1136	0.2786	0.0814	0.3195	0.2766	0.0921	-0.0236	0.1439	0.0968
β_{13}	-0.183	0.044	-0.1435	0.1375	-0.1587	0.3196	-0.2192	-0.0156	-0.1842	0.1519	-0.1252	-0.0697	-0.2169	0.1391	-0.1701
β_{14}	-0.0051	-0.1101	0.0068	-0.1744	0.246	0.184	0.2305	-0.0142	0.0526	0.1209	-0.2688	-0.0347	-0.2294	0.1377	-0.1729

Note: The values of $\beta_k^{(h)}$ have been randomly drawn from a uniform distribution within the range [–0.3, 0.3].

Table A7: Weights in Neural Network (Horizon 16–30)

	h=16	**h=17**	**h=18**	**h=19**	**h=20**	**h=21**	**h=22**	**h=23**	**h=24**	**h=25**	**h=26**	**h=27**	**h=28**	**h=29**	**h=30**
β_0	-0.2505	0.0765	-0.0754	0.0674	0.1766	-0.2096	-0.2642	0.2492	0.3198	0.0073	0.282	-0.1454	-0.1689	0.0226	0.2517
β_1	-0.0587	-0.1573	0.2465	0.1776	0.2328	-0.0266	0.1989	0.2944	-0.0828	-0.3074	-0.2274	-0.303	0.0646	0.1801	0.2033
β_2	0.2064	-0.1943	0.0068	0.1596	-0.1821	-0.0174	-0.0725	0.1327	0.1552	-0.0889	0.0786	0.0791	0.1909	0.1583	0.078
β_3	-0.3218	-0.1203	0.0671	-0.0184	-0.299	0.0503	-0.323	0.0728	0.0203	-0.1254	0.165	-0.2524	0.2744	-0.0302	-0.0359
β_4	0.1994	0.0564	-0.0856	0.1336	0.0988	-0.011	-0.2411	-0.0682	-0.1569	0.1346	-0.2686	-0.3083	0.0562	-0.1472	-0.0962
β_5	-0.1683	0.1184	-0.1752	-0.2578	0.3024	-0.1297	0.1702	-0.2165	0.2677	-0.3144	-0.2499	0.2814	-0.1555	-0.0094	-0.0042
β_6	-0.2068	-0.1335	-0.0314	0.2673	0.2478	-0.1672	-0.007	-0.0992	-0.077	0.2586	0.1498	0.3081	-0.2862	-0.2965	-0.195
β_7	-0.0164	-0.2573	0.2765	0.1239	-0.2996	0.0966	-0.0206	-0.1583	-0.2859	0.2578	-0.0776	0.2003	-0.0187	-0.1703	0.1296
β_8	0.1001	0.1107	0.1774	-0.0442	0.3117	0.0936	-0.0604	-0.0391	-0.2832	-0.2483	0.2228	0.2543	-0.2917	0.1562	-0.0322
β_9	-0.2901	0.0904	-0.1021	-0.0273	0.0608	-0.0781	0.2363	0.1249	0.116	-0.1551	0.1883	0.3047	-0.0083	-0.2865	-0.2768
β_{10}	0.0347	0.0376	0.1352	-0.0881	0.1869	0.1975	0.1677	0.0067	-0.1504	-0.0623	0.2663	-0.2017	-0.0676	-0.0615	0.0536
β_{11}	-0.163	0.0978	-0.0296	-0.294	-0.2281	-0.0317	0.1263	0.1609	0.1948	0.0203	0.0036	0.1633	0.1093	-0.0825	0.0891
β_{12}	0.3132	-0.1708	0.0049	0.075	-0.0931	-0.1294	-0.0441	0.0045	0.313	-0.2731	0.2334	-0.0441	-0.1524	-0.2158	-0.143
β_{13}	0.3058	0.2089	-0.2643	-0.2087	0.2651	-0.2256	-0.202	0.1443	0.232	0.0339	0.3139	0.0805	-0.2953	-0.0126	-0.1499
β_{14}	0.3062	0.1775	0.2406	0.0022	0.2276	-0.1645	-0.2008	0.2918	0.1267	-0.1132	0.0777	0.2305	-0.1006	-0.3013	-0.2121

Note: The values of $\beta_k^{(h)}$ have been randomly drawn from a uniform distribution within the range [-0.3, 0.3].

References

[1] Abid, I., Ayadi, R., Guesmi, K. et al. (2022). A new approach to deal with variable selection in neural networks: an application to bankruptcy prediction. *Ann Oper Res*, 313, 605–623.

[2] Adebiyi, A. A., Adewumi, A. O., & Ayo, C. K. (2014). Comparison of ARIMA and Artificial Neural Networks Models for Stock Price Prediction. *Journal of Applied Mathematics*, 614342.

[3] Ahluwalia, H. (2015). Space Weather Effects on RF Satellite Communication and Measures for Mitigation. *International Journal of Computer Applications*, 113(15), 17-20.

[4] Alexander, C., & Heck, D. F. (2020). Price discovery in Bitcoin: The impact of unregulated markets. *Journal of Financial Stability*, 50, 100776.https://doi.org/10.1016/j.jfs.2020.100776.

[5] Allen, W. A., & Wood, G. (2006). Defining and achieving financial stability. *Journal of Financial Stability*, 2(2), 152–172. https://doi.org/10.1016/j.jfs.2005.10.001.

[6] Alves, M. V., Echer, E., & Gonzalez, W. D. (2006). Geoeffectiveness of corotating interaction regions as measured by Dst index. *Journal of Geophysical Research: Space Physics*, 111, A07S05.

[7] Athanassopoulos, A.D., & Curram, S.P. (1996). A comparison of data envelopment analysis and artificial neural networks as tools for assessing the efficiency of decision-making units. *Journal of the Operational Research Society*, 47(8), 1000–1016.

[8] Asikainen, T., Mursula, K., Kerttula, R., Friedel, R., Baker, D., Soraas, F., Hoseph, F.F., & Bernard, B.J. (2005), Global view of energetic particles during a major magnetic storm. In: *The Inner Magnetosphere: Physics and Modeling: Geophysical Monograph Series*, vol. 155, 2005.

[9] Auer, R., & Tercero-Lucas, D. (2022).Distrust or speculation?The socioeconomic drivers of U.S. cryptocurrency investments. *Journal of Financial Stability*, 62, 101066.https://doi.org/10.1016/j.jfs.2022.101066.

[10] Basu, S., Easley, D., O'Hara, M., & Sirer, E. G. (2023). StableFees: A predictable fee market for cryptocurrencies. *Management Science*, 69(11), 6508–6524.

[11] Baxa, J., Horváth, R., & Vašíček, B. (2013). Time-varying monetary-policy rules and financial stress: Does financial instability matter for monetary policy? *Journal of Financial Stability*, 9(1), 117–138. https://doi.org/10.1016/j.jfs.2011.10..

[12] Berndsen, R. J., León, C., & Renneboog, L. (2018). Financial stability in networks of financial institutions and market infrastructures. *Journal of Financial Stability*, 35, 120–135. https://doi.org/10.1016/j.jfs.2016.12.007.

[13] Bourghelle, D., Jawadi, F., & Rozin, P. (2022). Do collective emotions drive bitcoin volatility? A triple regime-switching vector approach. *Journal of Economic Behavior & Organization*, 196, 294–306. https://doi.org/10.1016/j.jebo.2022.01.026.

[14] Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

[15] Cannon, P., Angling, M., Barclay, L., Curry, C., Dyer, C., Edwards, R., Greene, G., Hapgood, M., Horne, R., Jackson, D., Mitchell, Cathryn, Owen John, Richards Andrew, Rogers Christopher, Ryden Keith, S.S., Sweeting, M., Tanner, R., Thomson, A., & Underwood, C. (2013), *Extreme space weather: impacts on engineered systems and infrastructure*, London, Royal Academy of Engineering, 24pp.

[16] Cao, M., & Wei, J., (2004a). Weather derivatives valuation and market price of weather risk. *Journal Futures Markets*, 24, 1065–1089.

[17] Cao, M., & Wei, J., (2004b). Stock market returns: a note on temperature anomaly. *Journal of Banking & Finance*, 29, 1559–1573.

[18] Cao, Q., Parry, M.E. & Leggio, K.B. (2011) The three-factor model and artificial neural networks: predicting stock price movement in China. *Annal of Operational Research*, 185, 25–44.

[19] Capponi, A., Olafsson, S., & Alsabah, H. (2023). Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Management Science*, 69(11), 6455–6481.

[20] Cao, J., & Wang, J. (2019). Stock price forecasting model based on modified convolution neural network and financial time series analysis. *International Journal of Communication Systems*, 32(12), e3987. https://doi.org/https://doi.org/10.1002/dac.398.

[21] Castrén, O., Kavonius, I. K., & Rancan, M. (2022). Digital currencies in financial networks. *Journal of Financial Stability*, 60, 101000. https://doi.org/10.1016/j.jfs.2022.101000.

[22] Chan, N. H., & Genovese, R. C. (2001). A comparison of linear and nonlinear statistical techniques in performance attribution. *IEEE Transactions on Neural Networks*, 12(4), 922–928.

[23] Chang, S.-C., Chen, S.-S., Chou, K.R., & Lin, Y.-H. (2008), Weather and intraday patterns in stock returns and trading activity, *Journal of Banking & Finance*. 32, 1754–1766.

[24] Chen, Y.-C., & Huang, W.-C. (2021). Constructing a stock-price forecast CNN model with gold and crude oil indicators, *Applied Soft Computing*, 112, 107760.

[25] Chen, iaohong and H. White (1999), Improved rates and asymptotic normality for nonparametric neural network estimators, *IEEE Transactions on Information Theory*, vol. 45, no. 2, pp. 682-691, March.

[26] Chizhevsky, A. (1924), *Physical Factors of the Historical Process*, Kaluga, 1924.

[27] Chizhevsky, A. (1938), *Les Epidemies et les perturbations electro-magnetiques du milieu exterieur*, Paris, Hippocrate, 1938.

[28] Chizhevsky, A. (1976), *The Terrestrial Echo of Solar Storms*, 1976.

[29] Clark, C.E., & Greenley, H.L. (2019). *Bitcoin, Blockchain, and the Energy Sector*. Congressional Research Service. Link: https://sgp.fas.org/crs/misc/R45863.pdf.

[30] Constantinides, P., Henfridsson, O., & Parker, G.G. (2018). Introduction: Platforms and Infrastructures in the Digital Age. *Information Systems Research*, 29(2), 381–400. https://www.jstor.org/stable/48747824.

[31] Corbet, S., Larkin, C., Lucey, B., Meegan, A., & Yarovaya, L. (2020). Cryptocurrency reaction to FOMC Announcements: Evidence of heterogeneity based on blockchain stack position. *Journal of Financial Stability*, 46, 100706.https://doi.org/10.1016/j.jfs.2019.100706.

[32] Correa, R., He, A., Herpfer, C., Lel, U. (2022). The rising tide lifts some interest rates: climate change, natural disasters, and loan pricing. *International Finance Discussion Papers* 1345, Board of Governors of the Federal Reserve System.

[33] Cybenko, G. (1989). *Approximation by superpositions of a sigmoidal function.* Mathematics of Control, Signals and Systems, 2(4), 303–314.

[34] Daglis, T., Konstantinos K.N., Michaelides, P.G. (2019). Solar events and economic activity: Evidence from the US Telecommunications industry (1996–2014), *Physica A: Statistical Mechanics and its Applications*, Volume 534, 2019, 120805.

[35] Daglis, T., Konstantakis, K., Xidonas, P., Michaelides, P.G. & Zopounidis, C. (2023). Solar events and the US energy sector: a novel sectoral spillover GVAR approach introducing indirect GIRFs (IGIRF). *Annals of Operations Research*, 1-27. 10.1007/s10479-023-05471-7.

[36] Das, D., & Dutta, A. (2020). Bitcoin's energy consumption: Is it the Achilles heel to miner's revenue? *Economics Letters*, 186, 108530. https://doi.org/10.1016/j.econlet.2019.108530.

[37] Degryse, H., Goncharenko, R., Theunisz, C., Vadasz, T., (2023). When green meets green. *Journal of Corporate Finance* 78, 102355.

[38] Dieckmann S. Rare Event Risk and Heterogeneous Beliefs: The Case of Incomplete Markets. *Journal of Financial and Quantitative Analysis*. 2011;46(2):459-488.

[39] Dlugosz, J., Gam, Y. K., Gopalan, R., Skrastins, J., 2024. Decision-making delegation in banks. *Management Science*, 70, 3281–3301.

[40] Dufour, J.-M., & Renault, E. (1998). Short-run and long-run causality in time series: theory. *Econometrica*, 66 (5): 1099–1125.

[41] Dufour, J.-M., Pelletier, D. & Renault, E. (2006). Short run and Long run Causality in Time series: Inference. *Journal of Econometrics*, 132 (2): 337-362.

[42] Dwyer, G.P. (2015). The economics of Bitcoin and similar private digital currencies. *Journal of Financial Stability*, 17, 81–91. https://doi.org/10.1016/j.jfs.2014.11.006.

[43] Eastwood, J.P., Biffis, E., Hapgood, M.A., Green, L., Bisi, M.M., Bentley, R.D., Wicks, R., McKinnell, L.-A., Gibbs, M., & Burnett, C. (2017) The economic impact of space weather: Where do we stand? *Risk Analysis*, 37 (2), 206-218.

[44] Ehlers, T., Mojon, B., Packer, F., 2022. Carbon pricing and the cost of capital: Evidence from syndicated loans. *Journal of Banking & Finance*, 137, 106451.

[45] Elsayed, A.H., Gozgor, G., & Lau, C.K.M. (2022). Risk transmissions between bitcoin and traditional financial assets during the COVID-19 era: The role of global uncertainties. *International Review of Financial Analysis*, 81, 102069. https://doi.org/10.1016/j.irfa.2022.102069.

[46] Feng G, He J, Polson NG, Xu J. (2024), Deep Learning in Characteristics-Sorted Factor Models. *Journal of Financial and Quantitative Analysis*; 59(7):3001-3036.

[47] Fleissig, A., Kastens, T., & Terrell, D. (2000). Evaluating the semi-nonparametric fourier, aim, and neural networks cost functions. *Economics Letters*, 68, 235–244.

[48] Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707.

[49] Funahashi, K. (1989). *On the approximate realization of continuous mappings by neural networks*. Neural Networks, 2(3), 183–192.

[50] Garcia-Mata, C. & Shaffner, F.I. (1934). Solar and Economic Relationship: A Preliminary Report.*Quarterly Journal of Economics*, 49 (1) (1934) 1–51.

[51] Ginglinger, E., Moreau, Q., 2023. Climate risk and capital structure. *Management Science*, 69, 7492–7516.

[52] Goetzmann, W.N., & Zhu, N. (2005). Rain or Sunshine: Where is the Weather Effect? *European Financial Management*, 11, 559–578.

[53] Gomber, P., Kauffman, R.J., Parker, C., & Weber, B.W. (2018). On the Fintech Revolution: Interpreting the Forces of Innovation, Disruption, and Transformation in Financial Services. *Journal of Management Information Systems*, 35(1), 220–265. https://doi.org/10.1080/07421222.2018.1440766.

[54] Gopalswamy, N., Yashiro, S., Michalek, G., Xie, H., Lepping, R. P., & Howard, R. A. (2005). Solar source of the largest geomagnetic storm of cycle 23. *Geophysical Research Letters*, 32(12), L12S09.

[55] Griffin, J. M., & Shams, A. (2020). Is Bitcoin Really Untethered? *Journal of Finance*, 75(4), 1913–1964.

[56] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, 33(5), 2223–2273.

[57] Hansen, J.V., & Nelson, R.D. (2003). Forecasting and recombining time-series components by using neural networks. *Journal of the Operational Research Society*, 54, 307–317.

[58] Hautsch, N., Scheuch, C., & Voigt, S. (2024). Building trust takes time: Limits to arbitrage for blockchain-based assets. *Review of Finance*, 28(4), 1345-1381. https://doi.org/10.1093/rof/rfae004.

[59] Haykin, S. (1999). *Neural Networks*. Prentice-Hall, New Jersey.

[60] Hirshleifer, D. & Shumway, T. (2003), Good Day Sunshine: Stock Returns and the Weather. *Journal of Finance*, LVIII, 1009-1032.

[61] Hornik, K., Stinchcombe, M., & White, H. (1989). *Multilayer feedforward networks are universal approximators*. Neural Networks, 2(5), 359–366.

[62] Hsieh DA. Implications of Nonlinear Dynamics for Financial Risk Management. *Journal of Financial and Quantitative Analysis*, 28(1):41-64

[63] Hu, H.Z. & Xu, H. (2024), Global-GMDs: The global map of geomagnetic disturbances, *SoftwareX*, 25, 101619, ISSN 2352-7110, https://doi.org/10.1016/j.softx.2023.101619.

[64] Huang, R., Li, J., 2020. Do natural disasters affect bank lending? evidence from the us banking sector. *Journal of Financial Stability*, 46, 100707.

[65] Huberman, G., Leshno, J. D., & Moallemi, C. (2021). Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *Review of Economic Studies*, 88(6), 3011-3040. https://www.jstor.org/stable/27164142.

[66] Huynh, T. L. D., Shahbaz, M., Nasir, M. A., & Ullah, S. (2022). Financial modelling, risk management of energy instruments and the role of cryptocurrencies. *Annals of Operations Research*, 313(1), 47–75. https://doi.org/10.1007/s10479-020-03680-y.

[67] Farago, A. and G. Lugosi (1993), Strong universal consistency of neural network classifiers, *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1146-1151, July

[68] Feihong Hu, Daniel Mitchell, and Stathis Tompaidis (2024), Robust Financial Networks, *Operations Research*, 72(5): 1827–1842.

[69] Jalan, Akhil and Chakrabarti, Deepayan and Sarkar, Purnamrita (2024), Incentive-Aware Models of Financial Networks, *Operations Research*, 72, 6, 321-2336.

[70] Jana, R.K., Ghosh, I., Das, D., & Dutta, A. (2021). Determinants of electronic waste generation in Bitcoin network: Evidence from the machine learning approach. *Technological Forecasting and Social Change*, 173, 121101. https://doi.org/10.1016/j.techfore.2021.

[71] Jevons, W.S. (1878), Commercial crises and sun-spots, *Nature*, iv, 33–37.

[72] Jevons, W.S.(1879), Sun-spots and commercial crises, *Nature*, xix, 88–590.

[73] Jevons, W.S. (1882), The solar-commercial cycle, *Nature* xxvi, 226–228.

[74] Johnson, D., & Lee, S. (2004). The Economic Effects of the 2003 Halloween Solar Storm on the Aviation Industry, *Space Weather Economics*, 7(3), 198-207.

[75] Kacperczyk, Marcin T. and Peydro, Jose-Luis (2024), Carbon Emissions and the Bank-Lending Channel. *European Corporate Governance Institute – Finance Working Paper* No. 991/2024.

[76] Karmakar, S., Demirer, R.,& Gupta, R. (2021).Bitcoin mining activity and volatility dynamics in the power market. *Economics Letters*, 209, 110111. https://doi.org/10.1016/j.econlet.2021.110111.

[77] Kizys, R., Rouatbi, W., Umar, Z., & Zaremba, A. (2024). Air temperature and sovereign bond returns. *Financial Markets, Institutions & Instruments*, 1–31. https://doi.org/10.1111/fmii.12192.

[78] Klomp, J., & De Haan, J. (2009). Central bank independence and financial instability. *Journal of Financial Stability*, 5(4), 321–338. https://doi.org/10.1016/j.jfs.2008.10.001.

[79] Kou, S., & Peng, X. (2016). On the Measurement of Economic Tail Risk. *Operations Research*, 64(5), 1056–1072. https://doi.org/10.1287/opre.2016.1539

[80] Krausmann, E., Felton, C., Murtagh, W., Mitchison, N., & Lechner, S. (2014), *Space Weather and Financial Systems: Findings and Outlook*, JRC Science and policy reports.

[81] Laeven, L., Popov, A., 2023. Carbon taxes and the geography of fossil lending. *Journal of International Economics* 144, 103797.

[82] Lin, M.-Y., & An, C.-L. (2021). The relationship between Bitcoin and resource commodity futures: Evidence from NARDL approach. *Resources Policy*, 74, 102383. https://doi.org/10.1016/j.resourpol.2021.102383.

[83] Liu, Y., & Tsyvinski, A. (2021). Risks and Returns of Cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727. https://www.jstor.org/stable/48616705.

[84] Lu, J., & Chou, K.R. 2012. Does the weather have an impact on returns and trading activities in order-driven stock markets? Evidence from China. *Journal of Empirical Finance*, 19, 79–93.

[85] Lumineau, F., Shang, G., Swaminathan, J. M., Tsoukalas, G., Wagner, S. M., & Zhao, J. L. (2025). Charting the Future of Blockchain in Operations and Supply Chain Management: Opportunities and Challenges. *Journal of Operations Management*, 71(7), 886–892.

[86] Maghyereh, A.,& Abdoh, H. (2020). Tail dependence between Bitcoin and financial assets: Evidence from a quantile cross-spectral approach. *International Review of Financial Analysis*, 71, 101545. https://doi.org/10.1016/j.irfa.2020.101545.

[87] Marchionne, F., Pisicoli, B., & Fratianni, M. (2022). Regulation, financial crises, and liberalization traps. *Journal of Financial Stability*, 63, 101060. https://doi.org/10.1016/j.jfs.2022.101060.

[88] Marx, R.N. (2014). Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars, *Journal of Financial Economics*, 112, 137–146.

[89] Matringe, N., & Power, M. (2024). Memories lost: A history of accounting records as forms of projection. *Accounting, Organizations and Society*, 112, 101514. https://doi.org/10.1016/j.aos.2023.101514.

[90] Menkveld, A.J., & Yueshen, B.Z. (2019). The Flash Crash: A Cautionary Tale About Highly Fragmented Markets. *Management Science*, 65(10), 4470–4488. https://www.jstor.org/stable/48760765.

[91] Michaelides, P.G., Vouldis, A.T., & Tsionas, E.G. (2010), Globally flexible functional forms: The neural distance function, *European Journal of Operational Research*, 206 (2), 456-469.

[92] Michaelides, P.G., Tsionas, E.G., Vouldis, A.T., & Konstantakis, K.N. (2015), Global approximation to arbitrary cost functions: a Bayesian approach with application to US banking. *European Journal of Operational Research*, 241, 148–160.

[93] Michaelides, P.G., Tsionas, E.G., & Konstantakis, K.N. (2016). Non-linearities in financial bubbles: Theory and Bayesian evidence from S&P500. *Journal of Financial Stability*, 24, 61–70. https://doi.org/10.1016/j.jfs.2016.04.007.

[94] Minford, P., Ou, Z., Wickens, M., & Zhu, Z. (2022). The eurozone: What is to be done to maintain macro and financial stability? *Journal of Financial Stability*, 63, 101064. https://doi.org/10.1016/j.jfs.2022.101064.

[95] Mishra, W., & Teriaca, L. (2023). Propagation of coronal mass ejections from the Sun to the Earth. *Journal of Astrophysics and Astronomy*, 44, 20.

[96] Mursula, M., Holappa, L. Karinen, A. (2011), Uneven weighting of stations in the dst index. *J. Atmos. Sol.-Terr. Phys.*, 73, 316–322.

[97] Naeem, M.A., Nguyen, T.T.H., Karim, S., & Lucey, B.M. (2023). Extreme downside risk transmission between green cryptocurrencies and energy markets: the diversification benefits. *Finance Research Letters*, 58, 104263.

[98] Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. Bitcoin.org. Available: https://bitcoin.org/bitcoin.pdf.

[99] Nazareth N, & Reddy, Y.V.R. (2023). Financial applications of machine learning: A literature review. *Expert Systems with Applications*, 219,1 Article 19640.

[100] Nimalendran M, Pathak P, Petryk M, Qiu L. (2025), Informational Efficiency of Cryptocurrency Markets. *Journal of Financial and Quantitative Analysis*, 60(3):1427-1456.

[101] van Oordt MRC, Zhou C. (2016), Systematic Tail Risk. *Journal of Financial and Quantitative Analysis*. 2016;51(2):685-705.

[102] Panagiotidis, T., Papapanagiotou, G., & Stengos, T. (2024). A Bayesian approach for the determinants of bitcoin returns. *International Review of Financial Analysis*, 91, 103038.

[103] Park, S., Jang, K., & Yang, J.-S.(2021). Information flow between Bitcoin and other financial assets. *Physica A: Statistical Mechanics and Its Applications*, 566, 125604. https://doi.org/10.1016/j.physa.2020.125604.

[104] PCC (2013), Climate change 2013: The physical science basis. contribution of working group I to the Fifth assessment report of the intergovernmental panel on climate change, in: Stocker TF, D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.). Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1535.

[105] Piersanti, M., De Michelis, P., Del Moro, D., Tozzi, R., Pezzopane, M., Consolini, G., Marcucci, M. F., Laurenza, M., Di Matteo, S., Pignalberi, A., Quattrociocchi, V., & Diego, P. (2020). From the Sun to Earth: Effects of the 25 August 2018 geomagnetic storm. *Annales Geophysicae*, 38, 703–724.

[106] Rahman, M.R., Naeem, M.A., Yarovaya, L., & Mohapatra, S. (2024). Unravelling systemic risk commonality across cryptocurrency groups. *Finance Research Letters*, 105633.

[107] Reghezza, A., Altunbas, Y., Marques-Ibanez, D., Rodriguez d'Acri, C., Spaggiari, M., 2022. Do banks fuel climate change? *Journal of Financial Stability* 62, 101049.

[108] Rundo F., Trenta F., di Stallo, A.L. (2019). Machine Learning for quantitative finance applications: A Survey. *Applied Sciences*, 9(24):1-20.

[109] Sabeha, H. F., Hassanin, A., Minta, F. N., Elnawawy, M., & Mahrous, A. (2025). Dst index estimation from cosmic ray and solar wind data during solar cycle 24 using artificial neural networks. *Journal of Atmospheric and Solar-Terrestrial Physics*, 269, 106462.

[110] Sampson, Allan R. and Smith, Robert L. (1982), Assessing Risks Through the Determination of Rare Event Probabilities, *Operations Research*, 30, 5, 839-866.

[111] Sarkodie, S.A., Ahmed, M.Y., & Leirvik, T. (2022). Trade volume affects bitcoin energy consumption and carbon footprint. *Finance Research Letters*, 48, 102977.

[112] Schwarz, G. (1978), *The Annals of Statistics*, Vol. 6, No. 2, March, pp. 461-464.

[113] Sermpinis, G., Karathanasopoulos, A., Rosillo, R., & de la Fuente, D. (2021). Neural networks in financial trading. *Annal of Operational Research*, 297, 293–308.

[114] Sfetsos, A., & Siriopoulos, C. (2005). Time series forecasting of averaged data with efficient use of information. *IEEE Transactions on Systems, Man, and Cybernetics–A*, 35 (5), 738–745.

[115] Sharma, K., Dwivedi, Y.K. & Metri, B. (2022). Incorporating causality in energy consumption forecasting using deep neural networks. *Annal of Operational Research*. https://doi.org/10.1007/s10479-022-04857-3.

[116] Shin Li T, D, Wang B. Cryptocurrency Pump-and-Dump Schemes. *Journal of Financial and Quantitative Analysis*. Published online 2025:1-38.

[117] Siddique, M. A., Nobanee, H., Karim, S., & Naz, F. (2023). Do green financial markets offset the risk of cryptocurrencies and carbon markets?. *International Review of Economics & Finance*, 86, 822-833.

[118] Sierra-Porta, D., Petro-Ramos, J. D., Ruiz-Morales, D. J., Herrera-Acevedo, D. D., García-Teheran, A. F., & Tarazona Alvarado, M. (2024). Machine learning models for predicting geomagnetic storms across five solar cycles using Dst index and heliospheric variables. *Advances in Space Research*, 74(8), 3483–3495.

[119] Sigaeva, A., Nechaev, O.Y., Panasyuk, M.I., Bruns, A.V., & Troshichev, O.A. (2009), Thermal neutrons' response to the GLEs, *Advances in Space Research*, 43 (4) (2009) 649–653.

[120] Skiadopoulos, G., and Cheng, X. (2025), Climate-Triggered Institutional Price Pressure: Does it Affect Firms' Cost of Equity?, *Journal of Financial and Quantitative Analysis*, September 18, 2025, forthcoming

[121] Stege, N., Wegener, C., Basse, T., & Kunze, F. (2021) Mapping swap rate projections on bond yields considering cointegration: an example for the use of neural networks in stress testing exercises. *Annal of Operational Research*, 297, 309–321.

[122] Symitsi, E., & Chalvatzis, K.J. (2018). Return, volatility and shock spillovers of Bitcoin with energy and technology companies. *Economics Letters*, 170, 127–130. https://doi.org/10.1016/j.econlet.2018.06.012.

[123] Tasnim, H., Khalighi, M.-A., Bourennane, S., Pierre, L. & Opderbecke, J. (2016), Investigation of solar noise impact on the performance of underwater wireless optical communication links, *Opt. Express*, 24 (2016) 25832–25845.

[124] Temerin, M., & Li, X. (2006), Dst model for 1995–2002, *J. Geophys. Res.*, 111, A04221.

[125] Tsitsiklis, J.N., & Roy, B.V. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4), 694–703.

[126] Tsurutani, B. T., Gonzalez, W. D., Lakhina, G. S., & Alex, S. (2003). The extreme magnetic storm of September 1–2, 1859. *Journal of Geophysical Research*, 108, 1268.

[127] Von Peter, G. (2009). Asset prices and banking distress: A macroeconomic approach. *Journal of Financial Stability*, 5(3), 298–319. https://doi.org/10.1016/j.jfs.2009.01.001.

[128] Williams, R., Patel, H., & Thompson, L. (2011). Strategies for Mitigating the Economic Impact of Solar Storms on Satellite Operations", *International Journal of Space Policy*, 15(4), 2011, 275-289.

[129] White, H. (1982). Maximum Likelihood Estimation of Misspecified Models,*Econometrica* 50, no. 1: 1–25.

[130] White, H. (1990), Connectionist Nonparametric Regression: Multilayer Feedforward Networks Can Learn Arbitrary Mappings, *Neural Networks* 3, 535–550.

[131] White, H., and Racine, J. (2001). Statistical inference, the bootstrap, and neural-network modeling with application to foreign exchange rates. *IEEE Transactions on Neural Networks*, 12(4), 657–673.

[132] Yang, J., Cabrera, J., & Wang, T. (2010). Nonlinearity, data-snooping, and stock index ETF return predictability. *European Journal of Operational Research*, 200 (2):498–507.

[133] Yazıcı, A.F., Olcay, A.B., & ArkalıOlcay, G. (2023). A framework for maintaining sustainable energy use in Bitcoin mining through switching efficient mining hardware. *Technological Forecasting and Social Change*, 190, 122406. https://doi.org/10.1016/j.techfore.2023.122406.

[134] Yusoff, K. (2009). Excess, Catastrophe, and Climate Change. *Environment and Planning D: Society and Space*, 27(6), 1010-1029. https://doi.org/10.1068/d7407.

[135] Zhang, J., Richardson, I. G., Webb, D. F., Gopalswamy, N., Huttunen, E., Kasper, J. C., Nitta, N. V., Poomvises, W., Thompson, B. J., Wu, C.-C., Yashiro, S., & Zhukov, A. N. (2007). Solar and interplanetary sources of major geomagnetic storms (Dst ≤ -100 nT) during 1996–2005. *Journal of Geophysical Research*, 112, A10102.

[136] Zhang, Y.-J., Bouri, E., Gupta, R., & Ma, S.-J. (2021). Risk spillover between Bitcoin and conventional financial markets: An expectile-based approach. *The North American Journal of Economics and* , 55, 101296. https://doi.org/10.1016/j.najef.2020.101296.