| e-ISSN: 2948-460X | https://ibaf.usim.edu.my

BLOCKCHAIN AS A FINANCIAL NARRATIVE: CONSTRUCTING A MULTI-DIMENSIONAL DISCLOSURE INDEX AND EVIDENCE ON STOCK LIQUIDITY ACROSS THE U.S. STOCKS

Amin Haniff[1,a*], Asyraf Halim[1,b]

[1]*Faculty of Business and Economics, Universiti Malaya, Kuala Lumpur, Malaysia*
[1]*aminhaniff@um.edu.my*
[2]*asyrafhalim@um.edu.my*

*Corresponding Author: *aminhaniff@um.edu.my*

Abstract: *How does talking about blockchain impact the firm's economic values? Corporate blockchain disclosure has grown faster than the tools to measure it, leaving investors unable to separate substantive technological commitment from promotional "cheap talk". First, this paper constructs a firm-level Blockchain Disclosure Index (BDI) for United States public-listed firms over 2012–2025 using a hybrid natural-language-processing (NLP) pipeline that embeds filing sentences with a transformer model called the Sentences-BERT (SBERT). The index is validated as a measurement instrument: its four dimensions are empirically distinct and theory-carried rather than artefacts of the method, and firm-level disclosure intensity is systematic, with roughly three-quarters of its variance lying between firms and durable over time. Second, I then test whether the blockchain disclosure really moves the markets for the stable surviving firms throughout the period. At annual frequency there is no causal liquidity effect; Using the 10-K disclosures, the only robust association is positive and wrong-signed that is concentrated in small illiquid firms that consistent with reverse causality, and secondly, a dynamic system-GMM panel with valid instruments finds nothing. While the 8-K fillings at the disclosure event, price returns show no reaction, while abnormal illiquidity carries a stable negative coefficient, implicating a liquidity-improving that becomes significant once influential events are excluded, indicating a modest short-term liquidity benefit that operates through the liquidity channel rather than price.*

Keywords: blockchain, disclosure analysis, information asymmetry, stock liquidity, natural-language processing

1. Introduction

Listed firms increasingly describe blockchain and distributed-ledger activity in their regulatory filings, and the volume of that description has outpaced the market's ability to interpret it. A firm can announce a genuine, capital-intensive deployment, or it can adopt the vocabulary of innovation with no verifiable substance behind it, and both read similarly in an 8-K. This is the classic cheap-talk problem: when claims are costless to make, low-quality firms can mimic the language of high-quality ones (Crawford & Sobel, 1982), and investors cannot price what they cannot distinguish.

Existing measures of narrative disclosure do not resolve this problem for blockchain. Textual-analysis methods have quantified tone and readability in filings (Loughran & McDonald, 2011), and recent work has documented that speculative blockchain

announcements move prices (Cheng et al., 2019) and that disclosure can identify genuine adopters (Stratopoulos et al., 2022). What is missing is a validated, firm-level measure that separates the volume of blockchain disclosure from its quality, distinguishing substantive, verifiable commitment from symbolic posturing, along the theoretically meaningful dimensions of why, how, when, and who. Without such a measure, the questions that matter for disclosure transparency, information asymmetry, and market integrity cannot be asked cleanly.

The methodological opportunity is equally clear. Advances in natural-language processing, and transformer-based sentence embeddings make it possible to represent the meaning of disclosure at scale, rather than merely counting bag of words This combination of machine learning with auditable rules suits a setting where both semantic nuance and reproducibility matter.

This study pursues two objectives. First is the BDI construction: to build and validate a multi-dimensional Blockchain Disclosure Index (BDI) that captures both the intensity and the quality of corporate blockchain disclosure. The second is empirical: to test whether that disclosure affects stock liquidity relative to price. The contribution is therefore twofold, a machine-learning disclosure measure that is shown to be a valid and systematic firm-level construct, and the evidence on whether markets price the disclosure it captures.

The findings are as follows. The measure is valid: its four dimensions are empirically distinct and carry theoretical ground rather than mechanical structure, demonstrated that the firm-level disclosure intensity is systematic and durable. Yet, its market consequences are more limited. This study found that at annual frequency there is no causal liquidity effect, and the one robust association is wrong-signed and consistent with reverse causality among small, illiquid firms. Around the disclosure event itself, however, there is suggestive evidence of the predicted short-term liquidity benefit, operating through liquidity rather than price. The remainder of the paper reviews the theoretical framework, describes the construction and empirical strategy, reports the results, and concludes.

2. Literature and Theoretical Framework

Disclosure is a signaling act. Signalling theory holds that informed parties convey private information to uninformed parties through observable and costly signals (Spence, 1973). Whereas voluntary-disclosure theory extends this to the firm's discretionary choice to release information beyond mandatory requirements (Verrecchia, 1983). Central to both is that not all signals are equal, meaning, a signal is informative only to the extent that it is costly or verifiable enough that a low-quality firm would not imitate it.

Additionally, the economic channel through which credible disclosure operates is information asymmetry. Thus, when investors cannot value a firm's assets, they face adverse selection, and market makers widen bid-ask spreads to protect against better-informed counterparties, reducing liquidity (Akerlof, 1970). Disclosure that reduces this asymmetry lowers the adverse-selection component of the spread and improves liquidity (Diamond & Verrecchia, 1991). This is the liquidity channel, and it is the primary mechanism the present study tests.

Previous empirical work has supported this channel across disclosure types (Healy et al., 1999). Schoenfeld (2017), using index-fund ownership as a source of plausibly exogenous variation finds that greater voluntary disclosure increases stock liquidity and isolating the disclosure effect from the strategic motives that usually confound it. In the environmental, social, and governance setting, Chen et al. (2023) show that stronger ESG disclosure improves

liquidity for Chinese firms, operating through institutional-investor preference and risk-alleviation channels. Across financial and non-financial disclosure alike, credible information that reduces asymmetry is associated with more liquid markets, and this literature measures liquidity by the Amihud (2002) illiquidity ratio,which captures the price impact of trading.

However, blockchain disclosure is a harder case for this prediction. Unlike the audited financial disclosure examined by Schoenfeld (2017) or the rated ESG disclosure examined by Chen et al. (2023), blockchain narrative is largely unverified and promotional (Akyildirim et al., 2020), and is prone to the cheap-talk problem. Therefore, whether it reduces asymmetry enough to move liquidity is an open empirical question. Answering it requires first separating the substantive content of the disclosure from its symbolic packaging, which is the measurement task this study takes up before turning to the market test.

Related to the previous literatures of disclosure measurement, whether a given disclosure reduces asymmetry depends on its quality, either substantive or noise. The "cheap-talk" literature shows that costless communication may be uninformative in equilibrium (Crawford & Sobel, 1982). Applied to blockchain, this motivates a distinction between substantive disclosure, which is anchored in named technologies, quantified metrics, and verifiable milestones and is therefore costly to fabricate, and symbolic disclosure, which relies on promotional language without verifiable content. Hence, a credible measure must separate the two, because theory predicts they carry different market consequences.

This study distinguishes this measurement task into two features. First, blockchain disclosure is multi-dimensional: a single filing may describe why a firm adopts, how it does so technically, how far implementation has advanced, and how the adoption fits its industry, and these conceptually separate questions cannot be captured by a one-dimensional tone or frequency count. Second, it is the substance of the disclosure, not merely its presence that theory predicts should matter, which requires a measure that identifies verifiable content rather than volume alone.

Prior textual-disclosure research provides the methodological precedent for quantifying narrative content (Loughran & McDonald, 2011), and blockchain-specific studies establish that such disclosure is economically consequential and empirically tractable (Cheng et al., 2019; Stratopoulos et al., 2022). None, however, delivers a validated, multi-dimensional, quality-separated measure of blockchain disclosure, nor tests its liquidity and price consequences on a broad panel of US firms. These two gaps define the study's objectives.

3. Methodology

3.1 Data and Sample

The sample comprises United States domestic public firms that disclosed blockchain or distributed-ledger activity in Form 8-K or Form 10-K filings, earliest from 2012 until 2025. Market and accounting data are drawn from LSEG Workspace. Among the sample, the analysis reported here is restricted to firms that survived over the sample period. Firms that delisted or underwent identity discontinuity are held out and the survivorship boundary discussed in the conclusion. The annual liquidity panel comprises 821 firms and 7,937 firm-years, of which 1,431 firm-years carry positive blockchain disclosure in the annual filing and the remainder are firm-years in which a firm filed a 10-K with no blockchain content, while the event sample comprises 1,226 usable blockchain 8-K events across 240 firms.

A 10-K filed with no blockchain content is coded as a true zero for that firm-year, not as missing, and verified against the firm's full filing history. The panel then captures the move

from non-disclosure into disclosure rather than conditioning on the act of disclosing. Event filings within three trading days of a quarterly or annual earnings announcement are excluded, so the measured reaction reflects the blockchain disclosure and not a confounding earnings surprise. To prevent a few extreme observations from driving the results, continuous variables are winsorised at the 1st and 99th percentiles, with values beyond these cut-points reset to the percentile value.

3.2 Constructing the Blockchain Disclosure Index

The BDI is built through a hybrid pipeline that combines transformer-based semantic representation with theory-defined dimensions and auditable classification rules as shown in Figure 1. Blockchain-relevant filings are decomposed into related sentences, and each sentence is embedded using a pre-trained Sentence-BERT model (Reimers & Gurevych, 2019) which maps it to a 768-dimensional semantic vector.

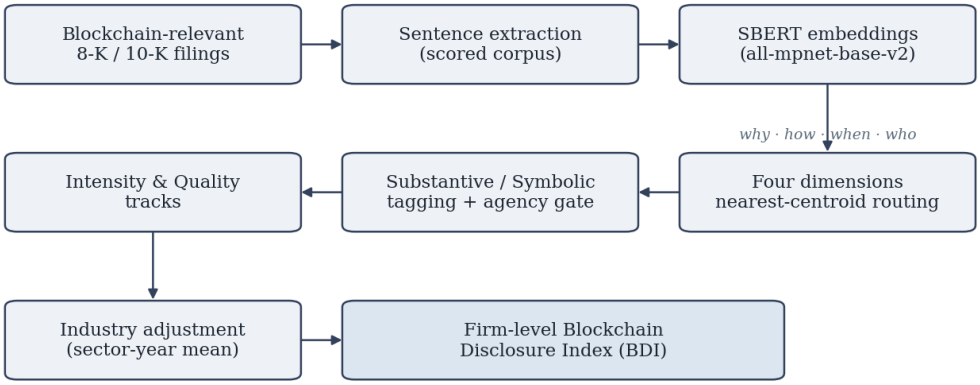

Figure 1. Construction of the Blockchain Disclosure Index (BDI)

Each sentence is then routed to one of four disclosure dimensions, defined a priori from theory rather than discovered inductively: Strategic Rationale (why a firm adopts), Technological Architecture (how), Implementation Status (when), and Industry Context (who). Each dimension is anchored by a small set of curated seed sentences whose embeddings define a dimension centroid and each sentence is assigned to the dimension whose centroid it lies closest to by cosine similarity. This nearest-centroid routing keeps the assignment transparent and reproducible in contrast to opaque unsupervised clustering.

The four dimensions are grounded in distinct literatures: strategic rationale in signalling and voluntary-disclosure theory, technological architecture in the blockchain-standards literature, implementation status in the diffusion-of-innovations tradition (Cole et al., 2019), and industry context in industry-structure analysis. Anchoring the dimensions in theory rather than discovering them inductively, is deliberate: it keeps the assignment auditable and reproducible, and empirical testing confirms that the resulting dimensions are genuinely distinct.

Moreover, signal quality is then tagged by a deterministic rule. A sentence is classified as substantive if it contains named technical entities, quantitative metrics, or specific verifiable milestones. Symbolic otherwise, subject to a firm-agency gate that excludes sentences referring to third parties rather than the firm's own activity. The index thus carries two orthogonal tracks: first, an intensity track measuring the volume of dimension-relevant disclosure, and second, a quality track separating substantive from symbolic content. Firm-year scores are industry-

adjusted by subtracting the sector-year mean, and the resulting substantive and symbolic intensities constitute the firm-level index carried into the empirical tests.

3.3 Empirical Strategy

The liquidity tests use the Amihud (2002) illiquidity ratio as the dependent variable, following the disclosure-liquidity literature (Schoenfeld, 2017; Chen et al., 2023) in which it is the standard measure of the price impact of trading. The static panel regresses annual log illiquidity on substantive and symbolic disclosure intensity with firm and sector-by-year fixed effects and firm-clustered standard errors Equation (1). The control set comprises firm size, leverage, return on assets, Tobin's Q, book-to-market, return volatility, turnover, and a count of blockchain 8-K filings:

$$\text{Ln}(ILLIQ)_{i,t} = \beta_0 + \beta_1 \cdot \text{ Sub }_{i,t} + \beta_2 \cdot \text{ Sym }_{i,t} + \sum \varphi_k \cdot \text{ Controls } + \eta_i + \theta_t + \varepsilon_{i,t} \quad (1)$$

where i indexes firms and t years; Ln(ILLIQ) is the natural logarithm of the firm's Amihud illiquidity; β_0 is the intercept; Sub and Sym are the firm's substantive and symbolic disclosure intensities, carrying coefficients β_1 and β_2; Controls is the control vector named above, with coefficients φ_k; η_i and θ_t are firm and sector-by-year fixed effects; and ε is the error term. The coefficient of interest is β_1, which theory predicts to be negative.

Since the sector-by-year fixed effects absorb the industry-year mean of disclosure, the regressors Sub and Sym enter as raw substantive and symbolic intensities rather than pre-centred values. Which mean, the industry adjustment is performed within the model rather than by subtracting a sector-year benchmark beforehand, which also absorbs sector-wide liquidity shocks in a given year.

Moreoever, since a static panel cannot separate the effect of disclosure from feedback running from past liquidity, then a dynamic system-GMM specification adds the lagged dependent variable and instruments disclosure with its own lags (Blundell & Bond, 1998), addressing reverse causality Equation (2):

$$\text{Ln}(ILLIQ)_{i,t} = \delta \cdot \text{Ln}(ILLIQ)_{i,t-1} + \beta_1 \cdot \text{Sub}_{i,t} + \beta_2 \cdot \text{Sym}_{i,t} + \sum \varphi_k \cdot \text{Controls } + \eta_i + \psi_t + \varepsilon_{i,t} \quad (2)$$

where all terms are as in Equation (1); the one-year lag of Ln(ILLIQ) enters with coefficient δ, which captures liquidity persistence; φ_k are the control coefficients; and ψ_t is a year fixed effect and disclosure is instrumented by its own lags.

In contrast of the panel data analysis, the event study measures the short-term reaction around each blockchain 8-K in two channels. In the price channel, the cumulative abnormal return over an event window w is the sum of daily abnormal returns, where the expected return is estimated from a Fama-French six-factor model (Fama & French, 2018) over a pre-event window as shown in the Equation (3):

$$CAR_i = \sum_{\tau \in w} \left(R_{i,\tau} - [\alpha_i + \beta_i' F_\tau] \right) \quad (3)$$

where CAR is the cumulative abnormal return for event i over the window w; τ indexes trading days within the window; R is the realised return on day τ; and α + β′F is the expected return, with α the intercept, β the vector of factor loadings and F the six Fama-French factor returns, all estimated over the pre-event window.

Instead, in the liquidity channel, abnormal illiquidity is the mean Amihud illiquidity over the event window less the firm's own pre-event baseline, so that each firm serves as its own control Equation (4):

$$AbILLIQ_i = \text{mean}_{\tau\in[-1,+1]} ILLIQ_{i,\tau} - \text{mean}_{\tau\in[-210,-11]} ILLIQ_{i,\tau} \quad (4)$$

where AbILLIQ is abnormal illiquidity for event i; ILLIQ is the firm's Amihud illiquidity on day τ; the window [−1, +1] is the three trading days centred on the filing; and [−210, −11] is the 200-day pre-event baseline. Each measure is then regressed cross-sectionally on event-level substantive and symbolic disclosure intensity with sector fixed effects, yielding the coefficients reported below.

Each specification is checked against standard diagnostics. In the static panel, the firm fixed effects are jointly significant and are retained, and multicollinearity is immaterial, with every variance-inflation factor below five. The residuals are cross-sectionally dependent (Pesaran, 2015), so the panel is re-estimated with Driscoll-Kraay standard errors as a robustness check (Driscoll & Kraay, 1998). In the event study, influential observations are identified by Cook's distance (Cook, 2000) and the cross-sections re-estimated with them removed. The dynamic panel is validated on two counts: the Hansen test does not reject the over-identifying restrictions, so the instruments are valid, and the Arellano-Bond test shows no second-order serial correlation (Arellano & Bond, 1991).

4. Results and Discussion

4.1 The BDI is a Valid and Systematic Measure (RO1)

The four disclosure dimensions are empirically distinct. The pairwise cosine similarity between dimension centroids ranges from 0.59 to 0.69, below the 0.70 threshold that would indicate redundancy, so the dimensions occupy separate directions in the semantic space (Table 1, Panel A). Crucially, this structure is theory-carried, not manufactured by the routing: in a decoy-null test that randomly relabels the anchor seeds fifty times, the true labelling produces centroids far more separable than any random relabelling (z = −10.6), establishing that the dimensional structure derives from the theoretical anchors rather than the method.

Disclosure intensity is a systematic firm attribute. The intraclass correlation of the firm-level measures lies between 0.72 and 0.82, so roughly three-quarters of the variation in blockchain disclosure is stable between-firm variation rather than filing-level noise (Table 1, Panel B). It is also durable: a firm's position in the disclosure distribution persists across years (rank correlation 0.70 to 0.77), and 82 to 85 per cent of firms remain consistently above or below their sector peers throughout the sample. The measure is therefore valid and systematic, satisfying the first objective, and is suitable as an independent variable in the market tests.

An unexpected structural finding accompanies this validation. The disclosure space does not resolve into discrete, density-separable clusters; density-based clustering fails to recover the four dimensions, and the typical sentence lies only moderately close to its nearest centroid. Blockchain disclosure is therefore better understood as a semantic continuum along which

theory-defined directions can be measured than as a set of naturally occurring categories, which is precisely why theory-anchored routing, rather than unsupervised clustering, is the appropriate method. Among the four dimensions, industry context is the most contestable, overlapping at its boundary with strategic rationale, and its per-dimension scores are treated with corresponding caution; the aggregate index is unaffected, since it sums across dimensions.

Table 1. Validation of the Blockchain Disclosure Index

Panel A. Dimensional distinctiveness	
Pairwise centroid cosine, six pairs	0.59 – 0.69
Decoy-null separability of real vs 50 random relabeling	z = −10.6
Panel B. Systematic firm-level variation	
Intraclass correlation between-firm share	0.72 – 0.82
Rank stability, first vs last year of the (Spearman)	0.70 – 0.77
Above/below-peer sign consistency	82% – 85% of firms

Notes. Panel A establishes discriminant validity; the decoy-null test relabels the anchor seeds at random to confirm the dimensional structure is theory-carried rather than an artefact of the method. Panel B establishes reliability on the panel of firms observed in two or more years. Full construct-validity evidence is reported in the underlying study.

4.2 Causal Liquidity Effect at Annual Frequency (RO2)

The static panel appears at first to show an effect, but it is the wrong sign and does not survive scrutiny. The substantive-intensity coefficient on illiquidity is positive and significant (Table 2), meaning that firm-years with more substantive blockchain disclosure are, if anything, less liquid, the opposite of the theoretical prediction. This coefficient strengthens rather than weakens when standard errors are corrected for cross-sectional dependence using Driscoll-Kraay, so it is not an artefact of understated precision. It is, however, an artefact of concentration: the top one per cent of disclosing firm-years hold nearly half of all substantive content, and removing these fifteen observations collapses the coefficient to insignificance. These heavy disclosers are small, volatile, and already illiquid, which points to reverse causality, illiquid and speculative firms disclose more, rather than disclosure worsening liquidity.

The profile of the heavy disclosers reinforces this reading. Relative to other disclosing firms they are markedly smaller, more volatile, higher-turnover, and already more illiquid before any disclosure is measured. While the profile of small, speculative firms for which heavy blockchain narrative and poor liquidity are joint symptoms rather than cause and effect. Event-time evidence is consistent: illiquidity is elevated in the years before a firm ramps up substantive disclosure, pointing to illiquidity preceding disclosure rather than following from it.

The dynamic panel confirms the absence of a causal effect. Estimated by system GMM with valid instruments (the Hansen test does not reject instrument validity and there is no second-order serial correlation), the substantive coefficient is small and insignificant, and liquidity persistence is moderate and stable. Once feedback from past liquidity is purged, no relationship between disclosure and liquidity remains in either direction.

Table 2. Annual Liquidity Panel and Dynamic-Panel Test

Specification	Substantive	Symbolic	Diagnostic
Panel A. Static panel (dependent variable: log Amihud illiquidity)			
Firm-clustered standard errors	+0.029** (0.025)	−0.010*** (0.005)	—
Driscoll–Kraay standard errors	+0.029*** (0.005)	−0.010*** (0.000)	CD test rejects independence
Top 1% of disclosers excluded	+0.005 (0.75)	—	Effect collapses (tail-driven)
Panel B. Dynamic panel (system GMM)			
Coefficient	−0.016 (0.82)	−0.004 (0.77)	Hansen 0.38; AR(2) 0.13; $\delta = 0.43$

*Notes. Dependent variable is log Amihud illiquidity on the 7,724-firm-year stable panel; p-values in parentheses. A positive coefficient denotes higher illiquidity (lower liquidity); the theory predicts a negative substantive coefficient. The wrong-signed positive coefficient strengthens under Driscoll–Kraay standard errors, yet collapses when the fifteen most substantive firm-years are removed, consistent with reverse causality. The GMM null, obtained under valid instruments, indicates no causal effect. p-values in parentheses; ***, **, * denote significance at the 1%, 5%, 10% levels.*

4.3 A Suggestive Short-Term Liquidity Benefit at the Event (RO2)

Interestingly, the event study tells a more nuanced story than the annual null. In the price channel, cumulative abnormal returns show no reaction to either substantive or symbolic disclosure, and this null is robust to the exclusion of influential events (Table 3). The market does not reward blockchain disclosure with abnormal returns.

The liquidity channel is different. The substantive coefficient on abnormal illiquidity is negative in every window, the theory-predicted direction of a disclosure that eases information asymmetry, but it is insignificant in the full sample. Applying a standard influential-observation rule (Cook's distance) removes 34 noisy events, and the coefficient, which barely changes in magnitude, becomes significant. The effect is therefore being masked by a handful of influential events rather than created by them: a modest, stable, short-term liquidity benefit emerges once they are excluded. Because this appears in liquidity and not in price, it is consistent with the study's channel-primacy premise, that blockchain disclosure, where it matters, resolves uncertainty rather than revising valuation. The result is reported as suggestive and specification-contingent, warranting confirmation rather than a strong claim.

Table 3. Event-Study Cross-Sections: Price and Liquidity Channels

Substantive coefficient (γ_1)	Full sample	Influential events excluded
Cumulative abnormal return [−1, +1]	+0.002 (0.75)	null unchanged (0.79)
Abnormal illiquidity [−1, +1]	−25.16 (0.22)	−22.56** (0.012)

Notes. Event-level cross-sections on 1,226 blockchain 8-K events; p-values in parentheses. A negative abnormal-illiquidity coefficient denotes reduced illiquidity (improved liquidity) around the event. Influential events are identified by Cook's distance; the substantive coefficient barely changes in magnitude while its precision improves, indicating a masked rather than a manufactured effect. The

*symbolic coefficient is insignificant across windows in both channels. p-values in parentheses; ***, **, * denote significance at the 1%, 5%, 10% levels.*

5. Discussion and Conclusion

5.1 Discussion and Implications

This study set out to measure corporate blockchain disclosure and to test its market consequences. On the first objective it succeeds: the Blockchain Disclosure Index is a multi-dimensional that are empirically distinct and theory-carried and systematic firm-level measure that separates substantive commitment from symbolic narrative. This is the paper's primary contribution, an auditable, reproducible instrument for a form of disclosure that has so far resisted measurement.

On the second objective the evidence is deliberately hedged. Blockchain disclosure does not robustly move markets for surviving US firms at annual frequency; the only robust annual association is wrong-signed and consistent with reverse causality among small, speculative, illiquid firms, and a correctly identified dynamic panel finds no causal effect. Yet at the disclosure event a different picture emerges: no price reaction specifically for this surviving and stable firms, but a modest and theory-consistent short-term liquidity benefit from substantive disclosure that surfaces once influential events are controlled. Where the market responds at all, it does so by resolving uncertainty, the liquidity, rather than revising value the price, consistent with the channel-primacy view of disclosure.

This pattern stands in informative contrast to the disclosure-liquidity literature. Voluntary financial disclosure (Schoenfeld, 2017) and ESG disclosure (Chen et al., 2023) have both been shown to improve liquidity by reducing information asymmetry. Blockchain disclosure, at least in aggregate and at annual frequency, does not. The contrast is not a contradiction but a qualification: the disclosures studied in that literature are audited or independently rated and therefore carry verifiable content (Krueger et al., 2024), whereas blockchain narrative is largely unverified and promotional, so the market's apparent indifference is consistent with a setting in which much of the disclosed volume is cheap talk rather than credible signal. Where credibility is highest, in substantive disclosure evaluated in the tight event window, a liquidity benefit does begin to emerge, which locates blockchain disclosure at the weaker, less-verifiable end of the disclosure-liquidity relationship rather than outside it.

The findings speak to disclosure transparency and information asymmetry in a blockchain setting. Substantive, a verifiable blockchain disclosure carries a mild and short-lived liquidity benefit, while symbolic disclosure is discounted, evidence that sophisticated markets can, at short horizons separate costly signals from cheap talk. At the same time, the absence of an annual effect cautions against overstating the market value of blockchain narrative, much of which is promotional volume that markets appear to ignore.

For an audience concerned with financial technology and market integrity, the message is measured. The tools now exist to quantify the substance of technological disclosure at scale and to hold promotional narrative to account, which is itself valuable for transparency and investor protection. The market evidence indicates that, at least among established firms, blockchain disclosure is priced cautiously if at all: substantive content earns a small short-term liquidity benefit and symbolic content is largely ignored. This is consistent with a market that is neither misled by hype nor strongly moved by substance at the annual horizon, and it locates whatever value blockchain disclosure carries in the resolution of uncertainty rather than the revision of fundamental value.

5.2 Limitations and Future Research

The clearest limitation of this study is also the entry point to its continuation. The analysis covers a single arm of a larger sample: the 821 firms that survived the sample period intact. Two further arms are deliberately held out and form the subject of subsequent research. The delisted arm, 405 firms, contains firms that disclosed blockchain activity and were later removed from the exchange or failed, and the identity-discontinuity arm, 177 firms, contains firms whose corporate identity changed around their blockchain disclosure, including the renamings and pivots that first drew regulatory attention to this setting (Akyildirim et al., 2020).

These are not peripheral cases. They are the firms in which the market consequences of blockchain disclosure, and the cheap-talk-to-failure pathway the measure is built to detect, should be most visible, and their exclusion is the reason the present market null is a property of surviving firms rather than of blockchain disclosure in general. Applying the identical pipeline to these arms, and comparing disclosure and its market effects across all three, is the natural completion of the study and the direction of ongoing work.

Three further boundaries qualify the findings. First, the sample is US-domestic and spans 2012 to 2025, so the results speak to that market and period rather than to blockchain disclosure globally. Second, the measure observes the content of disclosure but not managerial intent: it separates substantive from symbolic language, yet it cannot establish whether substantive-looking disclosure was strategically constructed to appear credible, so the distinction is a property of the text rather than a claim about motive. Third, the event-window liquidity result is contingent on the treatment of influential observations and is reported as suggestive; confirming it, and determining whether the short-term liquidity benefit persists or reverses, is part of the fuller event-study analysis the multi-arm extension will support.

Taken together, the study delivers a tested measure of corporate blockchain disclosure of version 1.0, and a well-identified account of its market consequences among surviving firms. It frames those consequences not as a final verdict but as the first arm of a broader investigation into how markets read the financial narrative that firms build around blockchain.

6. Acknowledgments

Alhamdulillah. All praises to the Allah, God Almighty. I am grateful to my supervisor and mentor, Dr Asyraf Abdul Halim, for his direct guidance throughout the research journey, for the opportunities he has opened to me, and for the financial support of a graduate research assistantship. I thank my co-supervisor, Associate Professor Dr Mohd Edil Abd Sukor, for his recommendations and constructive suggestions across the course of the research. I am indebted to Universiti Malaya for the resources, coursework, workshops, and bootcamps that shaped this work. I thank both of the proposal defense panel, in particular Dr Siti Hawa binti Yusof and Dr Mohammad Ali Tareq, together with Associate Professor Dr Rozaimah Zainudin, whose recommendations and suggested modifications sharpened the scope and direction of the study. I also thank Universiti Sains Islam Malaysia for organising and hosting the iBAF conference, and Maybank Islamic for the sponsorship that enabled students to present their work there. I thank INCEIF University for the opportunity, as an alumnus, to access the LSEG Workspace database drawn on in this study.

References

Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. The Quarterly Journal of Economics, 84(3), 488–500.

Akyildirim, E., Corbet, S., Cumming, D., Lucey, B., & Sensoy, A. (2020). Riding the wave of crypto-exuberance: The potential misusage of corporate blockchain announcements. Technological Forecasting and Social Change, 159, 120191.

Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. Journal of Financial Markets, 5(1), 31–56.

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. The Review of Economic Studies, 58(2), 277–297.

Blundell, R., & Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. Journal of Econometrics, 87(1), 115–143.

Chen, M.-t., Yang, D.-p., Zhang, W.-q., & Wang, Q.-j. (2023). How does ESG disclosure improve stock liquidity for enterprises — Empirical evidence from China. Environmental Impact Assessment Review, 98, 106926.

Cheng, S. F., De Franco, G., Jiang, H., & Lin, P. (2019). Riding the blockchain mania: Public firms’ speculative 8-K disclosures. Management Science, 65(12), 5901–5913.

Cole, R., Stevenson, M., & Aitken, J. (2019). Blockchain technology: Implications for operations and supply chain management. Supply Chain Management: An International Journal, 24(4), 469–483.

Cook, R. D. (2000). Detection of influential observation in linear regression. Technometrics, 42(1), 65–68.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. Econometrica, 50(6), 1431–1451.

Diamond, D. W., & Verrecchia, R. E. (1991). Disclosure, liquidity, and the cost of capital. The Journal of Finance, 46(4), 1325–1359.

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. The Review of Economics and Statistics, 80(4), 549–560.

Fama, E. F., & French, K. R. (2018). Choosing factors. Journal of Financial Economics, 128(2), 234–252.

Healy, P. M., Hutton, A. P., & Palepu, K. G. (1999). Stock performance and intermediation changes surrounding sustained increases in disclosure. Contemporary Accounting Research, 16(3), 485–520.

Krueger, P., Sautner, Z., Tang, D. Y., & Zhong, R. (2024). The effects of mandatory ESG disclosure around the world. Journal of Accounting Research, 62(5), 1795–1847.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. The Journal of Finance, 66(1), 35–65.

Pesaran, M. H. (2015). Time series and panel data econometrics. Oxford University Press.

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP), 3982–3992.

Schoenfeld, J. (2017). The effect of voluntary disclosure on stock liquidity: New evidence from index funds. Journal of Accounting and Economics, 63(1), 51–74.

Spence, M. (1973). Job market signaling. The Quarterly Journal of Economics, 87(3), 355–374.

Stratopoulos, T. C., Wang, V. X., & Ye, H. (2022). Use of corporate disclosures to identify the (non)adopters of blockchain. Accounting Horizons, 36(3), 197–220.
Verrecchia, R. E. (1983). Discretionary disclosure. Journal of Accounting and Economics, 5, 179–194.