Liquidity Collapse and the Propagation of Price Shocks*

C Stuart Adams† Arzé Karam‡ Christopher McCarty§ Jack Waller¶

March 14, 2026

Abstract

Extreme intraday price movements often occur even when trading pressure is not unusually large. We show that such events arise when market liquidity temporarily loses its ability to absorb order flow locally, allowing trading pressure to propagate through rapid sequences of price adjustments. We formalize this mechanism in a Depth-Driven Cascade (DDC) model in which the propagation horizon of trading pressure depends on observable liquidity conditions. The model represents short-horizon price dynamics using a quantum-inspired propagation kernel whose effective depth is mapped to top-of-book liquidity and order-flow imbalance. When liquidity is abundant, the model collapses to classical diffusion; when liquidity deteriorates, propagation persistence increases and the conditional return distribution becomes heavy-tailed and asymmetric. Using intraday data for CME E-mini S&P 500 futures from 2010 to 2022, we show that propagation fragility rises sharply during microstructure-driven dislocations—including the May 2010 flash crash and the August 2015 ETF event—but remains near its baseline during macro-driven volatility episodes such as the March 2020 COVID-19 crisis. Conditional on identical order-flow imbalance, high-fragility states exhibit substantially higher probabilities of subsequent extreme returns. These findings indicate that liquidity-driven propagation constitutes a distinct mechanism for extreme price movements that is not captured by volatility, jump, or order-flow-based models.

JEL Classification: G12, G14, C58.

*We appreciate the insightful discussions and suggestions from Xavier Gabaix. Many thanks also to Agostino Capponi, Anna Pavlova, Esen Onur, Kostantin Sokolov, Alexander Wagner and the principals at October 2024 Financial Stability Steering Group Deep Dive Liquidity Microstructure in Financial Conduct Authority, seminar participants at QRFE 2nd Workshop on Quantitative Finance, at Central Bank of Ireland in February 2025 and Commodity Futures and Trading Commission in March 2025 for their comments and suggestions. All errors are ours.

†Durham University, Department of Physics (c.s.adams@durham.ac.uk)

‡Corresponding author: Durham University, Department of Finance (arze.karam@durham.ac.uk)

§Durham University (christopher.mccarty1904@gmail.com)

¶University of Kingston London, Department of Mathematics (k2407357@kingston.ac.uk)

1. Introduction

Extreme price movements in modern financial markets often occur within minutes and are frequently followed by rapid reversals. Prominent examples include the May 6, 2010 flash crash in U.S. equity markets, the October 15, 2014 Treasury flash rally, and the August 25, 2015 ETF dislocation. These episodes share a striking feature: very large price movements occur over extremely short horizons even though the underlying trading pressure is not historically extreme. This observation raises a fundamental question: why do moderate trading imbalances sometimes generate disproportionately large price movements?

Standard asset-pricing frameworks typically attribute extreme returns to unusually large shocks. In diffusion-based models, volatility rises when uncertainty increases; in jump–diffusion models, large price movements arise from rare jump arrivals; and in self-exciting intensity models, clusters of trades generate bursts of volatility. Across these frameworks, the size or arrival rate of shocks varies, but the mechanism through which trading pressure translates into prices remains essentially unchanged.

In this paper we document a different mechanism. Extreme price movements can arise when market liquidity temporarily loses its ability to absorb trading pressure locally. When liquidity at the best quotes deteriorates, incoming orders are less likely to be absorbed immediately at prevailing prices. Instead, directional trading pressure generates sequences of rapid price adjustments within short calendar intervals. These rapid adjustment sequences allow moderate order-flow imbalances to produce large realized price movements.

The key economic distinction is therefore between shock scaling and propagation amplification. Shock scaling occurs when unusually large shocks generate large price changes through a standard impact relation. Propagation amplification arises when ordinary trading pressure produces large price movements because the market's capacity to absorb that pressure deteriorates. When liquidity becomes thin, identical order-flow imbalance can generate substantially larger price adjustments because trading pressure propagates through successive price adjustments before liquidity replenishes.

To formalize this mechanism, we develop a Depth-Driven Cascade (DDC) model of short-horizon price dynamics. The central variable in the model is propagation depth, which represents the number of internal price-adjustment steps occurring within a given calendar interval. Propagation depth is determined by observable liquidity conditions. In the empirical implementation we proxy liquidity using top-of-book depth, defined as the volume available at the best bid and ask prices. When liquidity is abundant, propagation depth is close to one and price dynamics resemble classical diffusion. When liquidity deteriorates, propagation depth increases, allowing directional pressure to persist across multiple adjustment steps and generating heavy-tailed and asymmetric

return distributions.

The model implements this mechanism through a propagation kernel that tracks directional pressure across successive price-adjustment steps. Importantly, the framework nests classical Brownian diffusion as a limiting case when liquidity is abundant. Extreme returns therefore emerge not from larger shocks but from changes in the mechanism through which trading pressure propagates through prices. In this sense, liquidity deterioration alters the price formation process itself, transforming otherwise normal trading pressure into large realized price movements. We evaluate the model using intraday data for the CME E-mini S&P 500 futures contract from 2010 through 2022. The E-mini market provides a natural laboratory for studying liquidity-driven propagation because it is among the most liquid markets in the world yet has experienced several flash-style dislocations. We construct measures of top-of-book liquidity and order-flow imbalance and map these variables into the model's propagation parameters. Top-of-book depth captures the immediate liquidity buffer available to absorb incoming order flow and provides a high-frequency proxy for short-horizon liquidity conditions.

Our empirical analysis reveals three main findings. First, liquidity deterioration precedes extreme price movements during flash-style events. During the May 2010 flash crash and the August 2015 ETF event, top-of-book depth collapses sharply in the minutes preceding the largest price movements. This pattern indicates that liquidity deterioration occurs before price dislocations rather than merely as a consequence of them.

Second, propagation-based fragility behaves differently from traditional volatility measures. During the March 2020 COVID-19 crisis, realized volatility reached historically high levels. Yet our propagation-based fragility measure remained close to its long-run baseline because liquidity providers continued to replenish depth between price movements. In contrast, propagation fragility rises sharply during flash-style dislocations that occur at moderate realized volatility. This separation indicates that liquidity-driven propagation represents a distinct form of market stress that is not captured by volatility alone.

Third, conditional on identical trading pressure, liquidity conditions strongly affect the probability of extreme outcomes. Observations characterized by high propagation fragility exhibit substantially higher forward exceedance rates than observations with comparable order-flow imbalance but stronger liquidity conditions. This pattern shows that liquidity fragility captures forward-looking information about extreme price movements beyond what is contained in order flow or volatility.

Our results are also closely related to the growing empirical literature on flash crashes and high-frequency market disruptions. Studies of the May 6, 2010 flash crash and related episodes emphasize the role of rapid liquidity withdrawal and order-book fragility in amplifying price movements (e.g., Kirilenko et al., 2017; Menkveld and Yueshen, 2019). These studies document that ex-

treme intraday price dislocations often occur when liquidity providers temporarily step away from the market, causing trading pressure to propagate rapidly through prices. Our framework provides a complementary perspective by formalizing this mechanism through a propagation-based model of price adjustment. Rather than focusing on the specific institutional triggers of individual flash events, the Depth-Driven Cascade model captures the general microstructure channel through which temporary liquidity deterioration can amplify otherwise moderate trading pressure.

These findings contribute to several strands of the literature. First, the paper contributes to the literature on market microstructure and price impact. Previous work shows that liquidity affects the magnitude of price responses to order flow (Kyle and Obizhaeva, 2016; Bouchaud, 2011). Our framework complements this perspective by focusing on the temporal propagation of price adjustments. Rather than modifying the magnitude of price impact alone, deteriorating liquidity alters how trading pressure propagates across successive price adjustments.

Second, the paper relates to theories of endogenous amplification in financial markets. Prior work shows that large price movements can emerge from nonlinear trading dynamics or granular shocks (Gabaix et al., 2003; Gabaix, 2011). Our mechanism differs in that extreme returns arise from changes in the market's response function rather than from the size of shocks themselves. In this sense, propagation amplification represents a microstructure channel through which liquidity conditions influence the distribution of returns.

Related empirical and theoretical work studies the joint determination of order flow and prices. Deuskar and Johnson (2011) model a system in which trading flows affect prices and prices feed back into flows, showing that order flow accounts for a substantial fraction of observed price variation. These results speak primarily to the magnitude of price responses to trading pressure. By contrast, they remain largely silent on the timing of adjustment—that is, on how rapidly prices incorporate trading pressure under different liquidity conditions. Our analysis complements this literature by showing that, conditional on comparable trading pressure, liquidity conditions govern primarily how quickly prices adjust rather than how far they ultimately move.

Our mechanism is also closely related to theories of endogenous liquidity withdrawal. In models of intermediary constraints and strategic liquidity provision, reductions in the risk-bearing capacity of some intermediaries induce other liquidity providers to withdraw as well, generating synchronized declines in market depth and amplifying price movements (Vayanos, 2004; Vives, 2014; Cespa and Vives, 2015). In these frameworks, liquidity supply exhibits strategic complementarities: when some market participants reduce their willingness to provide liquidity, others optimally follow, producing rapid collapses in market depth. Our empirical results are consistent with this mechanism. We show that when top-of-book liquidity deteriorates sharply, trading pressure propagates through successive price adjustments before being absorbed, generating large price movements even when underlying order-flow shocks are not unusually large.

This interpretation suggests a simple two-layer structure for extreme price dislocations. At a first layer, intermediary constraints and strategic complementarities in liquidity provision can generate synchronized withdrawals of market depth, as predicted by theory. At a second layer, once liquidity has withdrawn, trading pressure is no longer absorbed locally but propagates through sequences of rapid price adjustments. In this sense, our propagation depth measure can be interpreted as a reduced-form summary of the microstructure consequences of endogenous liquidity withdrawal. Third, the paper contributes to the literature on extreme market events and financial stability. Existing monitoring frameworks typically rely on volatility-based indicators to detect stressed market conditions. Our results suggest that volatility-based indicators may fail to detect microstructure-driven fragility because propagation risk arises from liquidity deterioration rather than from large shocks.

More broadly, the framework suggests a reinterpretation of extreme intraday price movements. Rather than viewing these events solely as responses to unusually large shocks, they may reflect temporary breakdowns in the market's capacity to dissipate trading pressure. When liquidity collapses, propagation accelerates and price dynamics depart from diffusion. When liquidity recovers, diffusion-like behavior re-emerges.

The remainder of the paper proceeds as follows. Section 2 describes the economic mechanism underlying the model and explains how liquidity conditions govern the propagation of trading pressure through the order book. Section 3 develops the propagation-based framework for short-horizon price dynamics. Section 4 formalizes the liquidity-conditioned propagation model using a quantum-walk representation and highlights its classical diffusion limit. Section 5 introduces the data and constructs the microstructure variables—such as order-book depth and order-flow imbalance—that discipline the model. Section 6 examines the distributional implications of liquidity-driven propagation. Section 7 presents event-level evidence from major stress episodes across several markets. Section 8 compares the propagation framework with diffusion, jump–diffusion, Hawkes, and regime-switching benchmarks. Section 9 evaluates the framework as a monitoring tool for market fragility using rare-event metrics such as precision–recall curves and Recall@K. Section 10 explores whether machine-learning methods can improve real-time inference of propagation regimes. Section 11 concludes with implications for market monitoring and price formation.

2. The Mechanism in Economic Terms

This section describes the economic mechanism underlying the Depth-Driven Cascade (DDC) framework. The central idea is that short-horizon price dynamics depend not only on the magnitude of trading pressure but also on the resilience of market liquidity available to absorb that pressure. When liquidity is abundant, trading pressure is absorbed locally and prices adjust gradually. When liquidity deteriorates, identical trading pressure can propagate across multiple price

levels, producing rapid and extreme price movements. This interpretation is consistent with theoretical models of strategic liquidity provision, in which liquidity supply can collapse endogenously when risk-bearing capacity deteriorates (Vayanos, 2004; Vives, 2014; Cespa and Vives, 2015).

2.1. Model Inputs

The framework relies on two observable microstructure variables extracted from the limit order book at each point in time.

1. **Top-of-Book Depth (D_t).** Top-of-book depth is measured as the total quantity available at the best bid and best ask. It captures the market's immediate capacity to absorb market orders without moving prices. Higher values of D_t indicate a more resilient order book with greater local liquidity.

2. **Order-Flow Imbalance (OIB_t).** Order-flow imbalance measures the directional tilt of liquidity at the best quotes. It is defined as the normalized difference between bid-side and ask-side depth,

$$OIB_t = \frac{D_t^{\text{bid}} - D_t^{\text{ask}}}{D_t^{\text{bid}} + D_t^{\text{ask}}},$$

   where D_t^{bid} and D_t^{ask} denote the quantities available at the best bid and best ask, respectively. The statistic ranges from -1 (all liquidity on the ask side) to $+1$ (all liquidity on the bid side). It serves as an empirical proxy for short-horizon directional trading pressure.

Prior research shows that order flow plays a central role in price formation. Early work demonstrates that signed trade flow contains information about future price movements and contributes significantly to short-horizon price dynamics (Hasbrouck, 1991; Evans and Lyons, 2002). Subsequent studies document strong relationships between order imbalance, liquidity, and market returns (Chordia et al., 2002). These studies typically measure trading pressure using signed trade imbalance constructed from transaction data. Because the mechanism studied in this paper concerns the market's ability to absorb trading pressure through available liquidity, we instead measure imbalance using limit-order-book variables that capture the state of liquidity prior to trade execution. This perspective emphasizes the liquidity conditions governing how trading pressure propagates through prices rather than the magnitude of realized order flow alone.

2.2. Model Output

Given these inputs, the framework produces a conditional distribution of short-horizon price changes. When liquidity is abundant and order-flow pressure is balanced, the implied return distribution

is narrow and approximately Gaussian, consistent with standard diffusion-based models. When liquidity deteriorates or directional pressure becomes pronounced, the distribution widens and becomes increasingly skewed and heavy-tailed. Importantly, these changes arise directly from observable order-book conditions rather than from latent regime variables or exogenous shock processes.

2.3. Propagation Depth

The link between microstructure inputs and return distributions operates through a single economic variable: propagation depth N_t. This variable represents the effective number of internal price-adjustment steps occurring within a fixed calendar interval. When the order book is deep, incoming orders are absorbed at the best quotes and liquidity replenishes rapidly. Propagation depth therefore remains close to one and price dynamics resemble a classical random walk. When liquidity collapses, incoming orders walk the book, executing against limit orders at progressively worse prices. Each price level consumed corresponds to an additional internal adjustment step. Propagation depth therefore increases, allowing directional pressure to propagate across multiple price levels within the same calendar interval.

The key implication is that extreme price movements need not originate from unusually large shocks. Instead, identical trading pressure can generate much larger price adjustments when market liquidity is depleted.

The propagation depth parameter N_t can be interpreted as a microstructure-based measure of the effective propagation horizon of trading pressure. In limit-order-book markets, incoming orders are typically absorbed locally when liquidity is abundant, producing small price adjustments. When liquidity deteriorates, however, identical trading pressure can traverse multiple price levels before being absorbed. In this sense, N_t measures the number of price levels across which trading pressure propagates within a given interval. Rather than representing an exogenous modeling parameter, N_t summarizes the state of the order book and the market's capacity to dissipate trading pressure.

2.4. Distributional Implications

Figure 1 illustrates how return distributions vary with propagation depth while holding order-flow pressure constant.

- When $N_t = 1$ (abundant liquidity), the return distribution is narrow and approximately Gaussian, consistent with classical diffusion models.
- When $N_t = 10$ (moderate liquidity stress), the distribution widens and develops heavier tails.
- When $N_t = 50$ (severe liquidity deterioration), the distribution becomes wide and strongly

skewed, placing substantial probability on price movements that would represent multi-standard-deviation events under diffusive dynamics.

In each case the underlying trading pressure is identical. What differs is the amount of liquidity available to absorb that pressure and therefore the number of price levels across which it propagates. Diffusion, jump–diffusion, and Hawkes models typically generate large price movements through unusually large shocks or elevated event intensity. In contrast, the propagation framework generates extreme outcomes endogenously through liquidity-driven amplification.

2.5. Interpretation of the Quantum-Walk Representation

The model does not claim that financial markets obey quantum mechanics. The quantum walk serves purely as a mathematical representation of the propagation mechanism. It provides a compact way to model state-dependent directional persistence across successive price-adjustment steps while preserving a classical diffusion limit when $N_t = 1$.

The framework therefore characterizes the conditional distribution of returns given observable order-book conditions. Its primary role is diagnostic rather than predictive: it identifies market states in which liquidity fragility increases the likelihood of large price movements.

This interpretation is closely related to the market invariance framework of Kyle and Obizhaeva (2016, 2023), which emphasizes that trading costs, volatility, and price impact scale systematically with trading activity and liquidity. In that literature, liquidity primarily governs the magnitude of price responses to order flow. The propagation framework complements this perspective by focusing on the timing and persistence of price adjustment: when liquidity collapses, trading pressure propagates through successive price levels more rapidly, effectively increasing the number of internal price-adjustment steps within a given calendar interval.

When liquidity collapses, the propagation horizon expands and effective market time accelerates, producing rapid sequences of price adjustments within a short calendar interval.

3. A Liquidity-Conditioned Propagation Theory of Price Dynamics

3.1. Economic Mechanism

Modern electronic markets operate through limit-order-book mechanisms in which prices adjust as incoming market orders interact with standing limit orders. Under normal market conditions, liquidity at the best quotes absorbs trading pressure locally and price adjustments occur gradually. During periods of liquidity stress, however, this absorption mechanism can fail. When depth at the best quotes deteriorates sharply, incoming orders are less likely to be absorbed locally and trading pressure propagates through rapid sequences of price adjustments.

This observation suggests that short-horizon price dynamics depend not only on the magnitude of trading pressure but also on the market's ability to absorb that pressure through available liquidity. Let OIB_t denote order-flow imbalance during a short interval and let D_t denote top-of-book depth, defined as the quantity available at the best bid and best ask. A standard reduced-form price impact relation can be written as

$$r_t = \lambda(D_t) OIB_t, \tag{1}$$

where $\lambda(D_t)$ is an inverse-liquidity function capturing the sensitivity of prices to trading pressure.

While this representation captures the empirical fact that thinner markets produce larger price responses to a given imbalance, it implicitly assumes that price adjustment occurs instantaneously. In practice, prices adjust through sequences of quote revisions and executions. When liquidity is abundant, most trades execute near the best quotes and the adjustment sequence is short. When liquidity deteriorates, trading pressure can propagate through multiple successive price adjustments before liquidity replenishes.

To capture this mechanism, we introduce the concept of a *propagation horizon*. Let N_t denote the number of internal price-adjustment steps occurring within a fixed calendar interval. Economically, N_t measures how far trading pressure propagates through the price formation process before being absorbed.

When liquidity is deep, N_t is close to one and price changes resemble classical diffusive dynamics. When liquidity deteriorates, N_t increases endogenously because trading pressure persists across multiple adjustment steps.

Returns are therefore generated by the accumulation of local directional adjustments:

$$r_t = \sum_{k=1}^{N_t} \delta_{t,k}, \tag{2}$$

where $\delta_{t,k}$ denotes the incremental price adjustment associated with the k-th propagation step.

Figure 1 illustrates the distinction between standard diffusive dynamics and propagation-driven dynamics. Under classical diffusion, short-horizon returns remain approximately Gaussian. Under propagation dynamics, directional persistence across internal adjustment steps generates wider, asymmetric, and potentially multi-modal return distributions.

This formulation highlights a key distinction between two mechanisms that can generate extreme returns:

- **Shock scaling:** the magnitude of shocks increases.
- **Propagation amplification:** the number of adjustment steps increases.

Standard diffusion and jump models emphasize the first mechanism. Our framework focuses on the second.

A concrete example is provided by the May 6, 2010 flash crash in CME E-mini S&P 500 futures, shown in Figure 3. The speed and magnitude of the decline illustrate how extreme price movements can emerge over very short horizons when normal diffusive adjustment breaks down.

3.2. Liquidity and Propagation Depth

The propagation horizon N_t depends on the resilience of the order book. In limit-order-book markets, resilience reflects both the level of available liquidity and the speed with which new liquidity replenishes after executions. When depth is abundant, incoming orders are absorbed quickly and propagation remains local. When liquidity deteriorates, identical trading pressure can persist across multiple price adjustments before being absorbed.

We therefore model propagation depth as a function of observable liquidity conditions:

$$N_t = f(D_t, \Delta D_t), \tag{3}$$

where D_t denotes current depth and ΔD_t captures recent changes in depth.

Propagation depth is therefore large when the order book is both thin and deteriorating. These conditions correspond precisely to episodes in which trading pressure propagates through successive price adjustments rather than being absorbed locally.

Figure 4 shows this mechanism during the May 6, 2010 flash crash. Top-of-book depth collapses sharply during the dislocation, indicating that trading pressure is no longer absorbed locally. Figures 5 further illustrate how the distribution of depth shifts from a normal regime with substantial liquidity to a crash regime in which depth is frequently close to zero.

3.3. Implications for Return Distributions

The propagation mechanism has immediate implications for the distribution of returns. Because returns are generated by the accumulation of N_t adjustment steps, variation in propagation depth affects both the dispersion and the shape of the return distribution.

[Diffusion in Liquid Markets] When liquidity is abundant and order-flow imbalance is approximately balanced, the propagation horizon satisfies $N_t \approx 1$. In this case, returns converge to standard diffusive dynamics and short-horizon price changes are approximately Gaussian after volatility normalization.

[Tail Amplification under Liquidity Collapse] When order-book depth deteriorates sharply, the propagation horizon rises endogenously ($N_t \gg 1$). The conditional return distribution becomes wider and more heavy-tailed even when the underlying order-flow shock is not unusually large.

[Extreme Returns without Extreme Shocks] Conditional on comparable order-flow imbalance, markets with lower liquidity and larger propagation horizon are more likely to generate extreme short-horizon returns.

These propositions formalize the intuition that extreme price movements can arise even when trading pressure remains within historical ranges. What changes during stress episodes is the mechanism through which trading pressure propagates through prices. Figure 5 illustrate this mechanism by comparing the distribution of top-of-book depth in normal and crash regimes. During the crash, the depth distribution shifts dramatically toward zero, reflecting severe liquidity depletion.

4. Formal Framework: The Depth-Driven Cascade Model

4.1. Economic Summary

The Depth-Driven Cascade (DDC) model separates three distinct elements of short-horizon price formation:

1. Trading pressure, measured by order-flow imbalance.
2. Liquidity conditions, measured by top-of-book depth.
3. Propagation persistence, measured by the number of internal adjustment steps N_t within a calendar interval.

Liquidity conditions determine the propagation horizon N_t, and N_t determines the shape of the conditional return distribution. When liquidity is abundant, propagation remains local and price dynamics resemble classical diffusion. When liquidity deteriorates, trading pressure propagates through multiple internal adjustment steps, amplifying price movements.

The empirical manifestation of this mechanism can be seen during flash-style market disruptions. Figure 6 illustrates the time evolution of the liquidity collapse during the May 6, 2010 flash crash, showing how the top-of-book depth structure deteriorates rapidly before recovering as market conditions stabilize.

4.2. Quantum-Inspired Propagation Kernel

To implement this mechanism we require a representation that allows directional trading pressure to persist across successive adjustment steps while preserving a diffusion limit in liquid markets. We implement this mechanism using a quantum-inspired propagation kernel.

The quantum walk is used purely as a mathematical device. It does not imply quantum behavior in financial markets. Rather, the quantum-walk structure provides a compact way to represent directional persistence and mixing across propagation steps.

The directional state of the market is represented by a two-component vector

$$\psi_t = \begin{pmatrix} a_t \\ b_t \end{pmatrix}, \tag{4}$$

where a_t and b_t measure the amplitudes of upward and downward directional pressure.

The state evolves through repeated application of a mixing operator $U(\theta_t, \phi_t)$ governing directional persistence and asymmetry:

$$\psi_{t+1} = U(\theta_t, \phi_t)\psi_t. \tag{5}$$

Within each calendar interval, the operator is applied N_t times:

$$\psi^{(t+1)} = U^{N_t}(\theta_t, \phi_t)\psi_t. \tag{6}$$

The resulting distribution of price changes is obtained from the squared amplitudes of the final state vector.

Figure 2 illustrates the propagation kernel used in the Depth-Driven Cascade model and the distributions it generates under different propagation depths. As propagation depth increases, directional persistence across adjustment steps produces wider and more asymmetric return distributions. This behavior captures the key mechanism of the model: when liquidity conditions deteriorate, price dynamics depart from the diffusion benchmark because directional pressure persists across multiple internal adjustment steps.

4.3. Parameter Mapping from the Order Book

All model parameters are deterministic functions of observable microstructure variables.

The propagation depth is mapped from order-book conditions:

$$N_{\text{steps}}^{(t)} = \mathsf{v}_1 \mathsf{v}_2 \, \mathsf{PPV}^{(t)} \left(1 - \mathsf{PP}^{(t)}\right) \Delta\mathsf{MD}^{(t)}. \tag{7}$$

Here $\mathsf{PPV}^{(t)}$ denotes the normalized height of the primary peak of the depth distribution, capturing how frequently the book approaches zero depth. $\mathsf{PP}^{(t)}$ is the location of that peak, and $\Delta\mathsf{MD}^{(t)}$ represents the change in top-of-book depth between consecutive windows.

The scale parameters v_1 and v_2 regulate the overall magnitude of propagation and the effective frequency of internal price updates.

Economically, each term captures a different dimension of liquidity fragility. $\mathsf{PPV}^{(t)}$ increases when the order book frequently approaches zero depth. $1 - \mathsf{PP}^{(t)}$ increases when the book is already depleted. $\Delta\mathsf{MD}^{(t)}$ captures the speed of liquidity deterioration.

Propagation depth therefore becomes large precisely when the order book is structurally thin, currently depleted, and rapidly deteriorating—conditions under which trading pressure propagates through multiple adjustment steps rather than being absorbed locally. The empirical shift from normal liquidity conditions to crash regimes is illustrated by the depth distributions in Figure 5.

4.4. Classical Diffusion as a Limiting Case

When $N_t \to 1$ and directional asymmetry is small (balanced order flow), the propagation operator is applied only once within each interval. Interference effects do not accumulate and successive calendar intervals produce approximately independent increments.

By the central limit theorem, returns converge to classical diffusion. This corresponds to liquid markets with balanced trading pressure—the normal operating state of a healthy order book.

The Depth-Driven Cascade model therefore nests classical diffusion as a limiting case while allowing for rapid propagation dynamics during episodes of liquidity deterioration.

5. Data and Microstructure Variables

5.1. Markets and Sample Construction

Our primary empirical analysis focuses on the CME E-mini S&P 500 futures contract, one of the most liquid financial instruments globally and a central venue for price discovery in U.S. equity markets (Hasbrouck, 2003). The E-mini futures market operates as a fully electronic central limit order book (CLOB), making it well suited for studying intraday liquidity dynamics and extreme price dislocations.

We use a long intraday sample spanning January 2010 to December 2022, covering multiple market regimes, including the May 6, 2010 flash crash, the COVID-19 crisis, and subsequent episodes of elevated volatility. Our baseline data consist of one-minute aggregated quotes and trades obtained from Refinitiv continuous front-month futures chains (RIC: `ESc1`), synchronized to Coordinated Universal Time (UTC). The long sample allows us to distinguish regime-specific behavior from generic volatility patterns.

For external validation and robustness, selected analyses are replicated on additional markets that have experienced flash-style events, including CME Bitcoin futures, the SPDR S&P 500 ETF, and U.S. Treasury futures. These supplementary results are discussed where relevant but the E-mini futures remain our primary laboratory.

Top-of-book depth provides a high-frequency proxy for the market's immediate liquidity buffer, capturing the quantity of liquidity available to absorb incoming order flow at the best quotes. While deeper layers of the order book may also influence large order execution, variation in best-quote liquidity has been shown to explain a substantial fraction of short-horizon price impact.

5.2. Institutional Features of Modern Futures Markets

Trading in E-mini futures occurs nearly continuously from Sunday evening to Friday afternoon, with a daily one-hour maintenance break. Market participants include institutional investors, high-frequency trading firms, proprietary traders, and hedgers. Orders are submitted electronically and matched according to price–time priority.

The E-mini market is characterized by extremely high trading speeds, low latency, and intense competition among liquidity providers. Technological advances—including co-location and microwave transmission—have reduced communication delays between major trading venues to near physical limits (Laughlin et al., 2014). As a result, liquidity conditions can change abruptly, especially during periods of elevated order-flow imbalance or rapid execution of large trades.

These features make the E-mini market particularly susceptible to flash-style stress episodes, in which liquidity temporarily evaporates and prices adjust discontinuously. Regulatory safeguards, such as market-wide circuit breakers and limit-up/limit-down mechanisms, were introduced or refined following major dislocations, including the 2010 flash crash. While these mechanisms can halt trading during extreme price moves, they do not directly condition on underlying liquidity or order-book fragility.

5.3. Microstructure Variables

From the intraday quote and trade data, we construct a set of standard microstructure variables designed to capture liquidity, order-flow imbalance, and trading intensity. Top-of-book depth: the sum of quantities available at the best bid and best ask – the immediate liquidity available for execution without walking the order book. Order-Flow Imbalance: the normalized difference between the bid-side and ask-side depth at the best quotes, as our empirical proxy for short-horizon trading pressure. One-minute log returns, standardized using time-of-day volatility estimates to account for intraday seasonality. Trading intensity, measured as trade count and volume per minute.

5.4. Stress and Jump Definitions

Our empirical analysis focuses on identifying stress states rather than predicting individual price changes. We define stress outcomes using return-based thresholds calibrated from the long sample. Specifically, a stress event is recorded when the absolute standardized return exceeds a high quantile of its historical distribution. Forward-looking stress indicators are constructed to evaluate whether microstructure-based measures—including QW-derived fragility metrics—contain information about subsequent stress regimes.

This framework allows us to compare quantum-walk-based indicators to classical benchmarks such as realized volatility and Hawkes-style intensity proxies in both event studies and long-horizon monitoring exercises.

The next section introduces the quantum-walk probability law and describes how market microstructure variables are mapped to its parameters.

5.5. Motivating Evidence: The May 6, 2010 Flash Crash

To motivate our analysis, we begin with a descriptive examination of the May 6, 2010 Flash Crash in the CME E-mini S&P 500 futures contract. This event represents one of the most severe intraday dislocations in modern U.S. financial markets and provides a canonical example of a rapid transition from normal to stressed microstructure conditions.

Figure 3 plots the intraday evolution of E-mini prices on May 5 and May 6, 2010. Prices evolve smoothly during the pre-crash period before exhibiting a sudden and sharp decline beginning shortly after 13:30 Central Time on May 6, followed by a rapid but incomplete recovery. The speed and magnitude of the price movement are difficult to reconcile with diffusion-based models calibrated to normal market conditions.

Figure 4 documents the contemporaneous behavior of top-of-book depth at the best bid and ask. Liquidity deteriorates sharply in the minutes surrounding the price collapse, with depth falling by an order of magnitude relative to earlier in the day. Importantly, the withdrawal of liquidity coincides with intense trading pressure, indicating that the price dislocation reflects not a lack of trading activity but a collapse in available liquidity.

Together, these figures highlight two empirical features that recur across extreme intraday events: (i) abrupt, nonlinear price movements occurring over very short horizons, and (ii) rapid deterioration of order-book liquidity at the best quotes. These features motivate the need for models that can accommodate regime shifts in the return-generating process when market depth collapses, rather than relying solely on diffusive or jump-based dynamics.

6. Distributional Properties

Section 2 provided an intuitive preview of how the return distribution varies with propagation depth. This section characterizes the distributional implications more formally. The central result is that reductions in depth increase the effective propagation horizon, while sustained order imbalance induces directional asymmetry. Together, these forces generate heavy-tailed, skewed, and state-dependent return distributions that arise endogenously from the interaction of trading pressure and liquidity.

The key distinction from classical random walks is that large price moves do not arise from unusually large shocks, but from changes in propagation persistence. When N_t is small, the distribution is narrow and diffusive. When N_t increases, probability mass shifts rapidly toward the tails, producing large price changes even over short horizons. Directional imbalance is introduced through ϕ_t without eliminating tail risk on the opposite side – consistent with the empirical obser-

vation that flash events are often followed by partial reversals.

In the context of price formation, the quantum walk is used as a reduced-form propagation kernel rather than as a physical model. The number of steps, N_{steps}, governs the effective propagation horizon of price changes, while the phase parameter ϕ governs directional bias. Together, these parameters determine the shape of the conditional return distribution.

The key distinction from classical random walks is that large price movements do not arise from unusually large shocks, but from changes in propagation persistence. When N_{steps} is small, the distribution is narrow and diffusive. When N_{steps} increases, probability mass shifts rapidly toward the tails, producing large price changes even over short horizons. Directional imbalance is introduced through ϕ without eliminating tail risk on the opposite side.

To illustrate this mechanism in isolation, we construct artificial order-book states that mimic the qualitative features observed during extreme events, such as declining market depth and increasing order-flow imbalance. These states are mapped into quantum-walk parameters and used to generate families of probability distributions, as shown in Figure 1. This exercise highlights how changes in liquidity conditions alone can reshape return distributions, independent of changes in the underlying flow of trades.

Taken together, the results in this section establish that the quantum-walk framework provides a parsimonious representation of liquidity-conditioned propagation. By focusing on conditional distributions rather than point forecasts, the analysis clarifies the economic channel through which stressed microstructure conditions give rise to extreme price movements. In the next section, we illustrate how this mechanism operates in real markets during major flash crash and flash rally episodes.

6.1. Market depth structure and liquidity states

To use market depth at the best quotes as a control variable for the quantum walk (QW) propagation mechanism, it is necessary to first characterize the structure of depth in limit-order-book markets and how this structure evolves over time. Although the full order book is high dimensional and fluctuates rapidly at millisecond horizons, a small number of summary measures capture the economically relevant state of liquidity. We therefore adopt a parsimonious representation that focuses on the aggregate quantity available at the best bid and best ask, referred to as *top-of-book depth*.

In high-frequency environments, instantaneous depth is highly volatile, particularly when large market orders are executed in rapid succession. As a result, point-in-time depth measures can be noisy and difficult to interpret. Instead, we consider time-averaged distributions of top-of-book depth computed over rolling look-back windows. These depth distributions provide a compact representation of the prevailing liquidity state and allow us to distinguish between normal and

stressed market conditions.

Figure **??** illustrates the distribution of top-of-book depth for the E-mini S&P 500 futures contract over a representative window preceding the May 6, 2010 flash crash. The distribution is centered well away from zero, indicating that liquidity at the best quotes remains ample throughout the window. The spike at zero corresponds to brief intervals with no quote updates and is a mechanical feature of high-frequency data rather than a sign of illiquidity. Overall, the shape of the distribution is characteristic of normal market conditions in this contract.

Figure **??** contrasts this pattern with the depth distribution computed over an equally long window during the flash crash itself. In this case, the distribution shifts markedly toward zero, with a substantial fraction of observations concentrated at extremely low depth levels. This change reflects rapid depletion of standing limit orders combined with insufficient replenishment in the presence of intense trading pressure. Notably, this structural shift is not well captured by simple summary statistics such as the median: while the median depth declines during the crash, it remains strictly positive and masks the frequency with which the book effectively empties at the best quotes.

To better visualize the dynamics of this structural transition, Figure 6 plots the evolution of quantiles of the top-of-book depth distribution over time. The lower panel shows the intraday price path during the flash crash, while the upper panel reports the normalized positions of the lower quartile, median, and upper quartile of the depth distribution. All three measures exhibit a pronounced collapse as the event unfolds, reaching a minimum before gradually recovering.

Two features are particularly relevant for modeling propagation dynamics. First, the collapse in depth structure precedes the minimum in prices by several minutes, indicating that liquidity deterioration occurs before the most extreme price movements. Second, the recovery of depth is slow relative to the price rebound, consistent with elevated volatility persisting after the initial dislocation. These patterns highlight that extreme price movements are closely linked to the temporal evolution of liquidity, not merely to contemporaneous trading pressure.

Taken together, these figures show that the distribution of top-of-book depth provides a powerful diagnostic of market state. Rather than relying on single-point measures, the shape and quantiles of the depth distribution reveal when the market transitions from a regime in which trading pressure is absorbed locally to one in which price changes propagate rapidly across levels. In the next subsection, we show how these liquidity states are mapped into the parameters of the quantum walk, allowing the propagation mechanism to adjust endogenously as market depth deteriorates.

7. Illustrative Extreme Event Case Studies

This section presents event-level analyses of major flash crash and flash rally episodes. The purpose is illustrative: to make the propagation mechanism transparent by showing how liquidity collapse

interacts with order-flow pressure to generate rapid price dislocations.

We examine several prominent episodes, including the May 6, 2010 Flash Crash in the CME E-mini S&P 500 futures, flash crashes in Bitcoin futures markets, and the October 15, 2014 U.S. Treasury flash rally. Despite differences in market structure and liquidity, these events share a common pattern. In the minutes preceding the dislocation, order-book depth deteriorates sharply while trading pressure becomes increasingly one-sided. During the event window, prices adjust rapidly and nonlinearly, exhibiting dynamics that are difficult to reconcile with diffusion-based or jump-diffusion benchmarks calibrated to normal market conditions.

Within the quantum-walk framework, these episodes correspond to abrupt shifts in propagation persistence rather than to the arrival of unusually large shocks. Liquidity withdrawal increases the effective propagation horizon, while order imbalance induces directional bias, together producing fast and extreme price movements. The case studies illustrate that this mechanism operates symmetrically for crashes and rallies, with the sign of the dislocation determined by the direction of imbalance.

By grounding the analysis in well-documented historical episodes, the case studies serve two complementary roles. First, they provide intuition for how liquidity-driven propagation reshapes conditional return distributions during periods of stress. Second, they motivate the long-sample analysis that follows, in which we assess whether the same mechanisms identified in extreme events are detectable in real time and informative for monitoring stressed market regimes.

7.1. Distributional comparison metric

To evaluate how closely simulated price dynamics match observed market behavior, we compare the full conditional distributions of price changes rather than point forecasts. Our primary evaluation metric is the Wasserstein distance, also known as the Kantorovich, Mallows, or Earth Mover's distance. This metric defines a distance on the space of probability distributions and has become widely used in statistics and machine learning to compare distributions with potentially different shapes, skewness, and tail behavior.

Formally, the Wasserstein distance between two probability measures μ and ν on $\mathbb{R}^d$ is defined as

$$W_p(\mu,\nu) = \inf_{\substack{X\sim\mu \\ Y\sim\nu}} \left(\mathbb{E}\left[\|X-Y\|^p\right]\right)^{1/p}, \quad p \geq 1, \tag{8}$$

where the infimum is taken over all pairs of random variables X and Y with marginal distributions μ and ν, respectively.

Intuitively, the Wasserstein distance measures the minimal cost of transforming one distribution into another, making it well suited for comparing simulated and empirical return distributions that may differ in dispersion, skewness, or tail thickness. Unlike moment-based measures or likelihood

comparisons, it remains well defined even when distributions are non-Gaussian or heavy-tailed.

In our implementation, we compute the first-order Wasserstein distance using the `scipy` library, assigning equal weights to all observations. For each experiment, the simulated distribution is compared to the empirical distribution over the same event window or rolling evaluation window. Lower values of the Wasserstein distance indicate a closer match between simulated and observed price dynamics.

We emphasize that the Wasserstein distance is used solely as a diagnostic tool to compare conditional distributions and not as an objective function for model estimation. We choose the Wasserstein distance over alternative goodness-of-fit measures for three reasons. First, unlike Kolmogorov–Smirnov–type statistics, the Wasserstein distance is sensitive to differences across the entire distribution, including the tails, which are central to the study of extreme price dislocations. KS tests place equal weight on all quantiles and can therefore understate economically meaningful discrepancies in tail behavior.

Second, likelihood-based comparisons are ill-suited to our setting because the competing models imply fundamentally different distributional families, some of which do not admit tractable likelihoods or require strong parametric assumptions about jump intensities or tail indices. Wasserstein distance allows direct comparison of empirical and simulated distributions without imposing a common parametric structure.

Third, Wasserstein distance has a natural economic interpretation as the minimal cost of transporting probability mass from one distribution to another. In the context of price dynamics, this interpretation aligns closely with the notion of propagation error: a lower distance indicates that simulated price changes are closer, in magnitude and direction, to those observed in the data. For these reasons, the Wasserstein distance provides a transparent and robust diagnostic metric for evaluating distributional fit in stressed market conditions.

7.2. The May 6, 2010 Flash Crash

We begin with the May 6, 2010 Flash Crash in the CME E-mini S&P 500 futures contract, which serves as a canonical example of an extreme intraday dislocation in a highly liquid market. As documented in Section 3, the event is characterized by a rapid collapse in top-of-book depth followed by a sharp price decline and a gradual recovery.

Within the quantum-walk framework, this episode corresponds to a sudden increase in propagation persistence driven by deteriorating liquidity conditions. As the depth distribution shifts toward zero, the effective number of propagation steps rises sharply, allowing trading pressure to transmit rapidly across price levels. Directional imbalance during the event induces asymmetry in the propagation kernel, generating a pronounced downside skew while preserving substantial tail risk.

A complete simulation of a crash episode requires identifying the point at which the direction of propagation reverses. Empirically, we observe a consistent feature across flash crash events: an abnormal increase followed by a sharp decline in the positions of the depth distribution quantiles (lower quartile, median, and upper quartile), as shown in Figure 6. This pattern coincides closely with the price turning point and reflects a transient attempt at liquidity replenishment that fails under continued trading pressure. We use this feature as an indicator of the turning point within the simulation, at which the direction of propagation is reversed.

Figure 7 compares simulations from the depth-driven cascade model and a classical Geometric Brownian Motion (GBM) benchmark to the observed E-mini price dynamics during the event window. Both models are calibrated using comparable information on magnitude and direction of price changes. The GBM simulation captures neither the rapid acceleration into the crash nor the persistence of volatility during the recovery phase. By contrast, the model simulation reproduces the distinct pre-crash, crash, and post-crash regimes by allowing propagation persistence to vary endogenously with liquidity conditions.

This comparison highlights a central distinction between the two approaches. In the GBM framework, extreme price movements require unusually large shocks. In the propagation framework, similar outcomes arise from changes in propagation induced by liquidity withdrawal, even when the underlying trading pressure remains within historical ranges.

7.3. Flash crash episodes in Bitcoin futures

We next examine flash crash episodes in Bitcoin futures markets, which provide a contrasting environment characterized by lower baseline liquidity and higher volatility. These events allow us to assess whether the liquidity-driven propagation mechanism identified in the E-mini market operates similarly in less centralized and less liquid markets.

To illustrate the mechanism, we focus on the April 23, 2023 Bitcoin futures crash, during which prices declined by more than 8% within approximately 90 minutes before partially recovering. As in the E-mini case, the event is preceded by a pronounced deterioration in top-of-book depth and an increase in directional trading pressure.

Figure 8 compares the depth-driven cascade model simulation to a Geometric Brownian Motion benchmark. The observed Bitcoin price path is shown in red, while simulated paths are shown in black, with grey lines indicating individual simulation runs. As before, the model parameters are driven primarily by the evolution of top-of-book depth and order imbalance.

Despite the higher noise and lower liquidity of the Bitcoin market, the same qualitative mechanism is evident. Prior to the crash, propagation persistence increases as depth deteriorates. During the crash, rapid propagation generates large price movements without requiring unusually large shocks. The turning point is again associated with a transient spike and subsequent collapse in

the depth distribution positions, which we use to reverse the direction of propagation within the simulation.

The GBM benchmark fails to reproduce these regime transitions, instead producing price paths that scale smoothly with volatility and do not adapt to changes in liquidity conditions. Consistent with this visual evidence, the Wasserstein distance between simulated and observed prices is substantially smaller for the model than for GBM. Panel B of Table 3 reports average distances across 100 simulation runs, indicating superior fit of the model during the crash window.

These results demonstrate that the propagation-based mechanism operates across asset classes and liquidity environments. While the timing of turning points relative to price minima varies across markets—occurring earlier in less active markets such as Bitcoin—the underlying link between liquidity collapse and rapid price propagation remains robust.

7.4. The October 15, 2014 U.S. Treasury Flash Rally

We next examine the flash rally in U.S. Treasury markets on October 15, 2014, focusing on the CME 10-year Treasury note futures. On that day, yields on the 10-year Treasury fell by approximately 16 basis points within minutes before rebounding sharply, producing one of the most abrupt intraday price increases in modern fixed-income markets. Unlike equity and cryptocurrency flash crashes, this episode was characterized by intense buying pressure rather than selling, yet it occurred under similarly fragile liquidity conditions.

As in the crash events analyzed above, the rally was preceded by a marked deterioration in order-book depth at the best quotes. Although aggregate trading activity remained elevated, available depth thinned rapidly, indicating a withdrawal of liquidity provision. During this period, order flow became strongly one-sided, with aggressive buy orders dominating the book. Prices responded nonlinearly, rising rapidly across multiple price levels within a short time window.

Within the quantum-walk framework, the Treasury flash rally corresponds to the same propagation mechanism identified in crash episodes, differing only in the sign of directional imbalance. The collapse in depth increases the effective propagation horizon, while sustained buy-side imbalance induces a positive asymmetry in the walk. As a result, large upward price movements occur over short horizons without requiring unusually large informational shocks. The resulting return distribution exhibits the same non-diffusive features observed in crash events, including heavy tails and rapid propagation, but with opposite skew.

Figure 10 illustrates the QW simulation alongside the observed price path. The model reproduces the rapid acceleration and subsequent stabilization of prices more closely than a geometric Brownian motion benchmark calibrated to the same period. Because price movements in this episode were less extreme than in equity or cryptocurrency crashes, diffusion-based models perform comparatively better; nevertheless, the propagation model continues to deliver a lower

Wasserstein distance relative to GBM, as reported in Panel B of Table 3.

Taken together, this episode highlights an important symmetry: flash crashes and flash rallies arise from the same liquidity-driven propagation mechanism, with the direction of the event determined by order-flow imbalance rather than by fundamentally different dynamics. This symmetry reinforces the interpretation of the quantum walk as a reduced-form propagation kernel that adapts endogenously to stressed microstructure conditions, irrespective of whether pressure is on the buy or sell side.

7.5. The August 25, 2015 ETF Flash Crash and the Role of Circuit Breakers

We next examine the flash crash in U.S. equity exchange-traded funds (ETFs) on August 25, 2015, focusing on the SPDR S&P 500 ETF (SPY). On that morning, equity markets experienced severe opening dislocations, during which more than one-fifth of all U.S.-listed ETFs were temporarily halted. Although market-wide circuit breakers had been introduced following the 2010 Flash Crash, these mechanisms did not prevent sharp intraday price dislocations in ETF markets.

As in the futures and cryptocurrency events analyzed above, the ETF crash was preceded by a rapid deterioration in liquidity at the best quotes. Order-book depth thinned sharply at the market open, while trading pressure became strongly one-sided as market participants attempted to reprice risk in response to overnight information. In contrast to centralized futures markets, however, ETF trading occurred across multiple venues and was subject to security-level trading halts triggered by price movements rather than by liquidity conditions.

Within the propagation framework, this episode again corresponds to a sudden increase in propagation persistence induced by liquidity withdrawal. As top-of-book depth collapsed, even moderate trading pressure was transmitted rapidly across price levels, generating large price dislocations within minutes. Directional imbalance induced a negative asymmetry in the walk, producing a crash-like return distribution despite the absence of unusually large informational shocks.

Figure 9 compares the observed SPY price path to simulations generated by the quantum-walk model and by a geometric Brownian motion benchmark. Prior to the activation of trading halts, the quantum-walk model captures the rapid price adjustment more closely than GBM, reflecting its ability to adapt to changing liquidity conditions. Once circuit breakers are triggered, however, the mapping between microstructure variables and price dynamics changes abruptly. At that point, the simulation necessarily diverges from observed prices, not because the propagation mechanism fails, but because trading is mechanically interrupted.

This episode highlights an important limitation of price-based circuit breakers. By conditioning halts on realized price changes rather than on liquidity conditions, such mechanisms may activate only after propagation has already accelerated. The ETF flash crash illustrates that liquidity-conditioned fragility can build prior to large price moves and that halts may therefore be reactive

rather than preventive. From this perspective, microstructure-based monitoring—such as signals derived from order-book depth and imbalance—may provide complementary information to price-based safeguards. Taken together, the ETF episode reinforces the interpretation of flash events as liquidity-driven propagation phenomena. It also underscores the relevance of modeling frameworks, such as the QW approach, that condition directly on microstructure states rather than on prices alone.

The case studies in this section identify the propagation mechanism in individual extreme events. A natural question is whether the same mechanism is detectable across the full sample in real time, rather than in hindsight. Section 9 addresses this directly. The rolling-window analysis shows that the liquidity-conditioned propagation regimes identified visually in Figures 7–11 emerge endogenously in the 2010–2022 sample, can be detected from observable order-book data, and provide early-warning information not contained in standard volatility or intensity signals. Specifically, the selective activation finding in Table 7 is the formal counterpart of what Figure 7 shows for the 2010 crash: depth collapses precede extreme price moves, and the model's fragility measure rises sharply before the largest dislocations.

We examine whether alternative mechanisms—self-excitation, exogenous jumps, or latent regime transitions—can replicate the same liquidity-conditioned propagation dynamics documented in earlier sections. We benchmark the propagation framework against commonly used alternatives. Jump–diffusion and Lévy-class models are calibrated on rolling windows and evaluated for their ability to reproduce heavy tails in return distributions. Hawkes-style benchmarks model self-exciting trade arrivals and return exceedances, capturing endogenous clustering and volatility persistence. While these models successfully generate heavy tails and cascades, they do not condition the speed of price propagation or recovery dynamics directly on contemporaneous liquidity states.

8. Benchmark Models

The joint behavior of prices and liquidity illustrated in Figures 1 and 2 suggests a regime in which extreme outcomes arise from rapid propagation under deteriorating depth rather than from exogenous jumps or purely intensity-based clustering, motivating a comparison between diffusion, Hawkes-type self-excitation, Lévy jump models, and the quantum-walk propagation framework.

8.1. Jump–diffusion and Lévy-class return models

A standard approach to modeling extreme returns augments Brownian diffusion with discontinuous jumps. Let p_t denote the log price. In continuous time, a jump–diffusion model takes the form

$$dp_t = \mu\,dt + \sigma\,dW_t + J_t\,dN_t, \tag{9}$$

where W_t is a standard Brownian motion, N_t is a Poisson process with intensity λ, and J_t denotes the jump size.

In discrete time, returns satisfy

$$r_t = \mu\Delta t + \sigma\sqrt{\Delta t}\,\varepsilon_t + \sum_{k=1}^{N_t} J_{t,k}, \tag{10}$$

with $\varepsilon_t \sim \mathcal{N}(0,1)$. Lévy-class models generalize this specification by allowing for infinite-activity jump processes characterized by a Lévy measure $\nu(dx)$, yielding heavy-tailed return distributions (Cont and Tankov, 2003).

A central empirical reference for jump-based return models is Lee (2012), who studies the role of information flow in generating price jumps at intraday frequencies. Using high-frequency data, the author shows that jumps are strongly associated with bursts of trading activity and information arrival, and that jump risk is priced in the cross section of asset returns. Importantly, jumps in this framework are interpreted as discrete responses to informational shocks, with their timing and magnitude governed by latent news processes rather than by contemporaneous liquidity conditions. While this approach successfully links jumps to information flow, it treats the market's propagation mechanism as passive: once a jump occurs, prices adjust instantaneously, abstracting from how depth and order-book resilience shape the speed and persistence of adjustment.

This interpretation builds on the seminal jump–diffusion model of Merton (1976), in which asset prices follow a continuous diffusion punctuated by Poisson-distributed jumps with exogenously specified intensity and size distribution. These models provide a flexible statistical description of return distributions and are widely used in option pricing and risk management.

However, a common feature of both jump–diffusion and Lévy models is that extreme price movements arise from unusually large shocks rather than from changes in the market's ability to absorb and propagate trading pressure. Jump arrivals are specified independently of order-book depth, and propagation speed is invariant to liquidity conditions. As a result, these models do not capture the microstructure-driven amplification observed during flash-style dislocations, where modest trading pressure can generate large price movements when depth collapses.

In contrast, the quantum-walk framework shifts the focus from shock magnitude to propagation persistence. Extreme outcomes arise endogenously when liquidity deteriorates, increasing the effective propagation horizon even in the absence of large informational shocks. This distinction is central for understanding intraday dislocations in modern limit-order-book markets.

8.2. Self-exciting Hawkes and LOB-based benchmarks

Self-exciting point processes capture endogenous clustering in trades or extreme returns. Let N_t denote a counting process for trades or exceedances. A Hawkes process specifies the conditional

intensity

$$\lambda_t = \mu + \int_0^t \phi(t-s)\, dN_s, \tag{11}$$

where $\phi(\cdot)$ is a non-negative decay kernel governing self-excitation.

In discrete time, this implies

$$\lambda_t = \mu + \sum_{s<t} \alpha e^{-\beta(t-s)} \mathbf{1}\{r_s \in \mathscr{E}\}, \tag{12}$$

where $\mathscr{E}$ denotes an exceedance region. Prices are often modeled as

$$r_t = \sigma_t \varepsilon_t, \quad \sigma_t^2 = \sigma_0^2 + \gamma \lambda_t, \tag{13}$$

generating volatility clustering and cascades (Bacry et al., 2015).

Hawkes models capture endogenous amplification in activity but do not condition price propagation on the availability of liquidity. Depth enters, if at all, only indirectly through correlations with order flow.

Self-exciting Hawkes processes provide a prominent framework for modeling endogenous clustering in high-frequency financial data. A large literature demonstrates that Hawkes processes successfully capture the self-reinforcing nature of order flow, trade arrivals, and extreme price movements (Bacry et al., 2015). In these models, past events increase the likelihood of future events, generating cascades and volatility clustering even in the absence of exogenous news.

Filimonov and Sornette (2012) apply Hawkes processes to study flash crash dynamics, showing that extreme market events can emerge endogenously from feedback loops in trading activity. Their results highlight how market impact and order flow self-excitation can amplify initial disturbances into large dislocations. Related work by Aït-Sahalia et al. (2015) emphasizes the role of Hawkes processes in modeling high-frequency volatility and jump clustering. In these frameworks, price changes are driven by event intensities that evolve endogenously, and volatility emerges as a function of the intensity of order flow or price jumps.

While Hawkes models capture important features of endogenous amplification, their propagation mechanism operates through event intensity rather than through the availability of liquidity. Depth and order-book resilience enter, if at all, indirectly through their correlation with trade arrivals. As a result, Hawkes models do not explicitly condition the speed of price adjustment or recovery dynamics on contemporaneous market depth.

In contrast, the quantum-walk framework conditions propagation persistence directly on observable liquidity states. Extreme price movements arise when depth collapses, increasing the effective propagation horizon, rather than solely from elevated trading intensity. This distinction is particularly relevant in flash-style dislocations, where prices can move rapidly even when the

arrival rate of trades does not increase proportionally.

8.3. Markov-switching regime models

Markov-switching models provide a widely used framework for capturing regime-dependent dynamics in asset prices. In this class of models, returns evolve according to a latent state variable that follows a finite-state Markov process, allowing parameters such as volatility, drift, or jump intensity to differ across regimes.

The seminal contribution of Hamilton (1989) introduces Markov-switching autoregressive models to capture business-cycle regimes. In financial applications, these models have been used extensively to study volatility regimes, crash risk, and time-varying risk premia. For example, Ang and Timmermann (2012) show that regime shifts help explain changes in volatility dynamics and asset returns across tranquil and crisis periods.

In a typical two-state specification, returns satisfy

$$r_t \mid S_t = k \sim \mathcal{N}(\mu_k, \sigma_k^2),$$

where the latent regime $S_t \in \{1,2\}$ evolves according to fixed transition probabilities

$$\mathbb{P}(S_t = j \mid S_{t-1} = i) = P_{ij}.$$

Extended versions allow for regime-dependent jump intensities or higher moments, improving the fit of return distributions during turbulent periods.

While Markov-switching models capture persistent shifts in volatility and tail risk, regime transitions are governed by latent probabilities that are not directly linked to observable market microstructure. As a result, regime changes are typically identified with a delay and may not align closely with rapid liquidity deterioration observed during flash-style dislocations. Moreover, within a given regime, price propagation remains diffusive or jump-driven, abstracting from how depth and order-book resilience govern the speed and asymmetry of adjustment.

In contrast, the quantum-walk framework dispenses with latent regime variables and conditions propagation dynamics directly on observable liquidity states. Regime-like behavior emerges endogenously as market depth collapses or recovers, allowing for rapid transitions that are tightly coupled to real-time microstructure conditions.

Taken together, these benchmarks differ in how extreme outcomes arise: through exogenous shocks, intensity feedback, latent regime transitions, or liquidity-conditioned propagation. The empirical analysis below assesses which mechanism best aligns with observed microstructure dynamics during stress.

8.4. Evaluation principle

All benchmark models are estimated on rolling windows and evaluated using identical distributional metrics. The comparison focuses on conditional return distributions during stressed periods, emphasizing tail behavior, asymmetry, and recovery dynamics rather than point forecasts. This design allows us to assess which modeling framework best aligns with observed liquidity-driven propagation in real time.

Table 5 highlights that the quantum-walk framework differs fundamentally from existing approaches by modeling extreme price movements as a consequence of liquidity-conditioned propagation rather than as realizations of large shocks, clustered arrivals, or latent regime shifts.

8.5. Relation to time-change and market invariance models

The proposed framework is closely related to classical time-change models of asset prices. In particular, Clark (1973) models returns as a Brownian motion subordinated by stochastic trading activity, while subsequent work interprets volatility fluctuations as arising from variation in transaction intensity or market activity. In these models, price dynamics remain diffusive, but the effective speed of time varies endogenously.

Our approach generalizes this insight by allowing the time-change itself to be driven by limit order book liquidity. The market-time intensity N_t plays a role analogous to the subordinator in Clark (1973)'s model. When depth is abundant, directional pressure is absorbed quickly and the effective number of propagation steps remains small. When depth collapses, trading pressure propagates across multiple price levels and N_t increases, corresponding to an acceleration of effective market time.

The framework is also consistent with the market invariance hypothesis of Kyle and Obizhaeva (2016), which emphasizes that trading costs and price impact scale with market activity and liquidity. In our setting, liquidity governs not only the magnitude of price fluctuations but the number of effective price update steps, providing a structural interpretation of impact scaling at intraday horizons.

Unlike classical time-change models, however, the quantum walk formulation allows for transient deviations from diffusive behavior when liquidity collapses and directional pressure becomes unstable. These deviations are state-dependent and episodic, rather than permanent features of the price process. As a result, the model nests time-changed Brownian motion as a limiting case while capturing the endogenous breakdowns of classical propagation observed during extreme events.

9. Monitoring Performance Evaluation

While event studies illustrate how liquidity-driven propagation operates during extreme episodes, an important question is whether the framework provides useful information for real-time monitoring of market fragility. We therefore evaluate the Depth-Driven Cascade (DDC) model as a monitoring tool using both distributional comparisons and rare-event metrics designed for tail-risk detection, including precision–recall curves and Recall@K. These metrics assess how effectively the model identifies periods in which the probability of extreme price dislocations is elevated relative to benchmark models.

9.1. Liquidity-Conditioned Fragility Signal

At each time interval we compute a fragility indicator derived from the propagation depth parameter N_t implied by the state of the order book. Higher values of N_t correspond to states in which trading pressure is more likely to propagate across multiple price levels, increasing the probability of extreme short-horizon price movements. The resulting fragility measure provides a forward-looking indicator of stressed microstructure states.

A central requirement for any monitoring signal is *selective activation*. A useful indicator should remain quiescent during normal trading conditions and activate sharply only when market microstructure resembles states observed around dislocations—specifically, when order-book depth deteriorates and directional trading pressure becomes unstable. This criterion distinguishes fragility measures that reflect breakdowns in price propagation from generic volatility or activity proxies, which often remain elevated for extended periods and generate excessive false alarms.

Accordingly, our evaluation focuses on two complementary dimensions of performance. First, we compare the full conditional distributions implied by each model to realized returns using distributional metrics that emphasize tail behavior and asymmetry. Second, we evaluate the usefulness of model-implied tail-risk measures as early-warning signals for subsequent stress using precision–recall and ranking-based criteria suited to rare events.

9.2. Definition of Stress Events

We consider two related stress definitions. For descriptive analysis, a contemporaneous stress event is recorded when the absolute standardized return exceeds the 99th percentile of its historical distribution:

$$|r_t| > q_{0.99}.$$

For monitoring evaluation, the relevant object is a forward-looking stress label. Let H denote

the prediction horizon and let τ denote a high return threshold. We define

$$y_t = 1\left\{\max_{1\le h\le H} |r_{t+h}| > \tau\right\}.$$

This forward label captures whether a stress event occurs within the next H minutes and is therefore the appropriate target for evaluating early-warning performance.

9.3. Empirical Setup and Evaluation Protocol

Let r_t denote one-minute log returns computed from mid-quotes, optionally standardized by a time-varying volatility proxy to remove predictable intraday seasonality. We implement a rolling evaluation procedure. At each minute t we (i) estimate benchmark model parameters on a trailing window of length W minutes, and (ii) evaluate out-of-sample distributional and monitoring performance over a forward horizon of length H minutes. We emphasize distributional comparisons, consistent with Section 6 and the event-window analysis, and avoid likelihood-based comparisons across fundamentally different distributional families.

Distributional evaluation. For each model m and time t we construct a predictive distribution for returns over the next H minutes. When the model admits a closed-form predictive density, we sample from it; otherwise we simulate from the implied mechanism. We then compare the model-implied distribution to the realized forward returns $\{r_{t+1},\dots,r_{t+H}\}$ using the first-order Wasserstein distance

$$W_1\left(\widehat{F}_t^{(m)}, \widehat{F}_t^{(\text{real})}\right), \tag{14}$$

where $\widehat{F}_t^{(m)}$ is the empirical CDF of simulated returns from model m for the forward horizon and $\widehat{F}_t^{(\text{real})}$ is the empirical CDF of realized returns over the same horizon. Lower values indicate closer distributional fit, including in the tails.

Monitoring evaluation. We also evaluate models as early-warning indicators of subsequent stress. Let $s_t^{(m)}$ denote the model-implied score at time t, interpreted as a measure of forward tail risk, such as $\mathbb{P}^{(m)}(|r_{t+1}| > \tau \mid \mathscr{F}_t)$ or the corresponding forward-horizon exceedance probability. Because extreme events are rare, ROC-based measures can be misleading; we therefore report precision–recall performance and operational ranking criteria such as Recall@K.

9.4. Models: DDC versus Benchmark Mechanisms

Section 8 highlights that competing models generate extreme outcomes through different mechanisms: exogenous jumps (jump–diffusion and Lévy-class specifications), intensity feedback (Hawkes),

latent regime transitions (Markov-switching), or liquidity-conditioned propagation (DDC). In the full-sample monitoring setting we compare:

1. **DDC propagation (main model).** DDC parameters are inferred from observable microstructure states. Liquidity governs the effective propagation horizon via N_t, while directional pressure induces asymmetry through the phase or bias parameter. The DDC model is implemented using a quantum-walk propagation kernel. In normal regimes the model collapses to a diffusion-like limit, while in stressed liquidity states it generates non-Gaussian, asymmetric, heavy-tailed conditional laws.

2. **Gaussian diffusion baseline.** A rolling Gaussian benchmark calibrated by rolling mean and variance. This baseline captures time-varying dispersion but imposes diffusive propagation and thin tails.

3. **Heavy-tail return law (Lévy proxy).** A rolling Student-t specification fitted on each window. This benchmark captures excess kurtosis in returns but does not condition tail thickness on contemporaneous liquidity states.

4. **Jump–diffusion (Merton-type).** A rolling Merton jump–diffusion calibrated on each window. This benchmark captures discontinuities as exogenous jump arrivals with Poisson intensity and Gaussian jump sizes, but does not link propagation speed to depth states.

5. **Self-exciting intensity (Hawkes-style).** A Hawkes-type benchmark fit to exceedance arrivals (or trade-intensity bursts) and mapped to a forward exceedance probability. This captures endogenous clustering in stress arrivals, but its mechanism operates through intensity rather than liquidity-conditioned propagation.

6. **Latent regimes (Markov-switching).** A Markov-switching volatility benchmark with regime-dependent dispersion and filtered regime probabilities. This captures persistent shifts in volatility regimes, but regime changes are latent and may not align with rapid depth collapses.

9.5. Precision–Recall Evaluation

Because stress events are rare relative to normal trading periods, precision–recall metrics provide a more informative evaluation than traditional ROC curves. Precision measures the fraction of predicted stress periods that correspond to actual events, while recall measures the fraction of realized events that are correctly identified.

We compute precision–recall curves for the DDC fragility signal and compare them with several benchmark models including diffusion, jump–diffusion, Hawkes processes, and Markov-switching

volatility models. The empirical results show that the DDC fragility signal concentrates a significantly larger fraction of extreme events within the highest-risk deciles than diffusion, jump–diffusion, Hawkes, and regime-switching benchmarks.

Importantly, the signal activates selectively during microstructure-driven dislocations while remaining stable during macro-driven volatility regimes such as the March 2020 COVID-19 crisis.

9.6. Recall@K Monitoring Performance

To evaluate the operational usefulness of the fragility signal for monitoring applications, we compute Recall@K statistics. This measure evaluates how many realized stress events occur within the top K highest fragility scores. Recall@K is particularly relevant for monitoring frameworks in which regulators or risk managers can only investigate a limited number of alerts. Higher Recall@K values indicate that the monitoring signal concentrates extreme events within a small fraction of high-risk periods.

Recall@K. Because stress events are rare, standard classification metrics such as the area under the ROC curve can be misleading. We therefore emphasize ranking-based performance measures. Let $y_t \in \{0,1\}$ indicate whether a stress event occurs over the next H minutes, and let s_t denote the model's stress score at time t. Let $\mathscr{T}_K$ be the set of indices corresponding to the K largest values of s_t. Recall@K is defined as

$$\text{Recall@}K = \frac{\sum_{t \in \mathscr{T}_K} y_t}{\sum_t y_t}. \tag{15}$$

Recall@K measures the fraction of realized stress events captured among the top K model-ranked observations and is well suited to rare-event monitoring settings with limited alert capacity.

9.7. Selective Activation and Full-Sample Evidence

Figure 13 summarizes the full-sample behavior of the inferred model parameters and the resulting DDC stress regime over 2010–2022. Stress episodes are rare and episodic (32 days out of 3,159 trading days) and cluster around well-known market dislocations such as the Flash Crash, the August 2011 turmoil, the August 2015 selloff, the February 2018 volatility episode, and the COVID-19 crisis. Outside these episodes, the stress indicator remains largely quiescent. This pattern is inconsistent with persistent regime switching and instead supports the interpretation of DDC stress as a selectively activating, microstructure-conditioned propagation regime.

A key diagnostic is whether DDC stress merely restates realized volatility. Figures 14 and 15 show that while realized volatility is higher on average during stress periods, there is substantial overlap between stress and calm observations across a wide range of volatility levels. Stress and non-stress observations coexist at similar volatility levels, indicating that DDC stress does not simply coincide with volatility spikes. Rather, it reflects a distinct state variable linked to propagation

depth and recovery dynamics that complements volatility-based monitoring measures.

Figure 16 illustrates the full-sample behavior of the DDC fragility measure relative to realized volatility. Two features stand out. First, DDC fragility activates episodically and remains near baseline during most of the sample, consistent with a selectively triggered propagation regime rather than a persistent high-risk state. Second, DDC fragility does not mechanically track volatility: the largest volatility episode in the sample (COVID-19) is not accompanied by a commensurate increase in DDC fragility, while earlier microstructure-driven dislocations (e.g., 2010–2012, 2015, 2018) exhibit pronounced DDC activation despite only moderate volatility.

Tables 6–8 provide formal statistical evidence supporting this interpretation. Rank correlations between daily realized volatility and DDC stress are negative (Spearman $\rho = -0.058$, Kendall $\tau = -0.048$), and stress activation is entirely absent in the highest realized-volatility deciles, instead concentrating in low-to-intermediate volatility regimes. Stress days also exhibit lower realized volatility on average than non-stress days. These findings reject the hypothesis that fragility stress is a high-volatility flag and instead point to a distinct market state.

To assess whether fragility stress captures information beyond volatility, we examine propagation depth conditional on realized-volatility bins. Within each volatility decile, stress days exhibit substantially higher propagation depth than calm days. After demeaning propagation depth within volatility bins, the median difference between stress and calm observations is large (136.2 versus -0.47) and highly significant (Mann–Whitney $p < 10^{-21}$). Thus, even holding volatility fixed, stress days are characterized by sharply elevated propagation persistence.

Table 10 reports Recall@K for operationally relevant alert budgets. Across all values of K, the propagation-based monitor captures a larger fraction of realized extreme events than diffusion-based, heavy-tailed, and jump–diffusion benchmarks. The improvement is particularly pronounced at small K, which is the most relevant region for supervisory and risk-management applications where only a limited number of alerts can be investigated. These results reflect the selective activation property of DDC: probability mass is concentrated on periods characterized by fragile propagation states rather than being spread broadly across high-volatility intervals.

Finally, an incremental information test using a logistic classification framework shows that adding propagation depth to realized volatility increases explanatory power substantially. The average log-likelihood improves from -0.0566 under a volatility-only specification to -0.00065 when propagation depth is included, corresponding to an LR-like statistic of approximately 353. This result confirms that DDC stress reflects a liquidity-conditioned propagation mechanism that is not subsumed by volatility-based measures.

9.8. The COVID-19 Finding: Macro Repricing versus Microstructure Fragility

One of the most striking results concerns what the fragility measure does not do. During the March 2020 COVID-19 crisis, U.S. equity markets experienced some of the largest realized-volatility readings in modern history, yet the DDC fragility measure remained near its long-sample baseline. This is not a failure of the model; it is a central feature.

The COVID-19 crash was primarily a macro-driven repricing episode. Trading activity was intense and volatility was extreme, but order-book depth generally recovered between the largest price moves. Liquidity providers repriced aggressively, yet the evidence does not indicate the same persistent cascade through a depleted order book that characterizes flash-style dislocations. Realized volatility was extreme; propagation fragility was not.

By contrast, flash-crash-type episodes—including the 2010 E-mini crash, the 2015 ETF dislocations, and several episodes in 2011–2012—are characterized by elevated propagation depth and persistent directional pressure, even when realized volatility is not at historical extremes. These are precisely the episodes in which the DDC fragility measure activates sharply.

Figure 17 and Table 7 document this distinction formally. DDC stress concentrates in realized-volatility deciles D1–D7 and is absent in deciles D8–D10, which are dominated by macro-driven episodes including COVID-19. The rank correlation between realized volatility and DDC stress is slightly negative (Spearman $\rho = -0.058$). The implication is that the monitoring framework separates two distinct forms of market stress: macro-driven volatility spikes and microstructure-driven propagation fragility.

This distinction is economically important. A monitor based only on realized volatility would be most alarmed in March 2020, precisely when propagation fragility was unremarkable. The same monitor would be less alarmed during episodes such as May 2010 or August 2015, when propagation fragility was elevated and flash dislocations occurred within minutes. The DDC framework therefore provides information that is complementary to volatility-based surveillance.

9.9. Regime-Dependent Distributional Fit

This section evaluates whether the DDC model improves the *distributional* fit of intraday returns in a regime-dependent manner. The objective is not to replace a volatility-conditioned diffusion benchmark in tranquil periods, but to test whether propagation-based dynamics improve fit in regimes where stress monitoring is most relevant.

Regime definition. We partition the sample into four liquidity–propagation states using median splits of (i) top-of-book depth and (ii) propagation intensity N_t (smoothed), yielding: Low depth / Low propagation, Low depth / High propagation, High depth / Low propagation, and High depth / High propagation. This classification isolates regimes in which liquidity is structurally thin and/or

propagation persistence is elevated.

Terminal return horizon. For each minute t, we compute forward H-minute cumulative returns

$$R_{t,t+H} \;=\; \sum_{j=1}^{H} r_{t+j}, \tag{16}$$

with baseline horizon $H = 30$ minutes.

Distributional distance and interpretation. Within each regime s, we compare the empirical distribution of $R_{t,t+H}$ to model-implied distributions via the first-order Wasserstein distance W_1. We report regime-level performance using

$$\Delta W_1(s) \;=\; W_1^{DDC}(s) - W_1^{\text{Diffusion}}(s), \tag{17}$$

so that $\Delta W_1(s) < 0$ indicates superior distributional fit of DDC relative to the diffusion benchmark.

State-conditional calibration. To ensure that differences in W_1 reflect *distributional shape* rather than trivial scale effects, DDC simulations are calibrated within each regime using a two-step procedure:

1. Variance matching. A state-specific scale factor is chosen so that the simulated terminal volatility of DDC matches the empirical terminal volatility within regime s.

2. Tail-shape matching. Two state-specific shape parameters are selected (via grid search) to match the empirical 95th and 99th percentiles of $|R_{t,t+H}|$ within regime s.

This calibration keeps the diffusion benchmark unchanged and focuses the comparison on the ability of DDC to reproduce non-Gaussian tail behavior across regimes.

Bootstrap inference and tail-focused fit. To quantify uncertainty, we compute bootstrap distributions of $\Delta W_1(s)$ using $B = 250$ resamples within each regime. In addition to full-distribution W_1, we compute a tail-focused Wasserstein distance using only extreme quantiles ($q \leq 0.05$ or $q \geq 0.95$). This tail metric targets precisely the part of the distribution relevant for stress monitoring.

Figure 17 reports boxplots of $\Delta W_1(s)$ for both the full distribution (Panel A) and the tail-focused distance (Panel B). Overall distributional distances are broadly comparable between DDC and diffusion, with DDC achieving lower W_1 in two of the four regimes and remaining close in the others.

The tail-focused comparison reveals a clearer pattern: DDC achieves lower tail Wasserstein distance in three of the four regimes, and the bootstrap distributions are predominantly shifted below zero in those regimes. These results indicate that DDC primarily improves the representation of extreme quantiles rather than the central mass, consistent with the interpretation that volatility-conditioned diffusion models capture bulk dispersion well while propagation-based dynamics matter most in the tails.

Connection to monitoring performance. This distinction is economically important because the monitoring framework in Section 9 evaluates early-warning performance using Recall@K, which is inherently tail-driven. Consequently, improvements in tail distributional fit provide structural validation for the monitoring gains reported in the Recall@K analysis: DDC does not uniformly dominate diffusion in normal conditions, but it improves the modeling of tail realizations associated with propagation-driven stress, which is precisely where supervisory monitoring is targeted.

9.10. Why the DDC Model Differs from Benchmark Mechanisms in Full-Sample Monitoring

The full-sample results clarify the distinction between the propagation-based DDC mechanism and the benchmark model classes considered in Section 8. Diffusion and heavy-tailed return laws primarily adjust dispersion or tail thickness while preserving independent-increment or i.i.d. structure at the evaluation horizon. Jump–diffusion models attribute large moves to exogenous jump arrivals whose intensity and size distribution are not conditioned on contemporaneous order-book depth. Hawkes models capture clustering in the arrival of trades or extremes through self-excitation in intensity, but operate independently of the cross-sectional state of liquidity. Markov-switching models allow for persistent regime changes but rely on latent transitions that may not align with the rapid and transient depth collapses characteristic of flash-style dislocations.

By contrast, DDC conditions the propagation kernel directly on observable liquidity and imbalance states. In normal conditions it nests a diffusion-like limit, while in fragile states it generates asymmetric, heavy-tailed conditional return distributions through increased propagation persistence. This nesting property explains both the selective activation of DDC stress and the state-conditional improvement in distributional fit: DDC behaves like standard diffusion when markets are liquid, but departs endogenously when microstructure deteriorates.

Taken together, the rolling evaluation evidence shows that liquidity-conditioned propagation regimes identified in event studies are not isolated anomalies. Such regimes arise endogenously in the full 2010–2022 sample and can be detected in real time using observable order-book states. Relative to the benchmark mechanisms in Section 8, DDC improves distributional fit precisely in stressed microstructure states and delivers superior monitoring performance under rare-event

evaluation metrics, without sacrificing performance in normal conditions.

We next examine whether machine learning can improve the real-time inference of propagation regimes by mapping high-dimensional microstructure data into DDC parameters, while preserving the structural interpretation of the model.

9.11. Incremental Information Beyond Order-Flow and Impact Benchmarks

A natural question is whether propagation fragility simply repackages information already contained in classical microstructure variables emphasized by price discovery and impact frameworks, such as signed order flow, volatility, and liquidity. To address this, we evaluate whether propagation fragility provides incremental monitoring content after controlling for standard order-flow and volatility measures.

Let y_t denote a forward stress indicator, defined as an exceedance over the next H minutes:

$$y_t = 1\left\{\max_{1\le h\le H} |r_{t+h}| > \tau\right\}, \tag{18}$$

where r_t is the one-minute return (optionally standardized) and τ is a high threshold. We compare monitoring signals using a parsimonious specification that conditions on (i) signed order flow, (ii) volatility, and (iii) liquidity, and then adds the DDC fragility score:

$$\Pr(y_t = 1 \mid \mathscr{F}_t) = g\Big(\alpha + \beta_1 \,\mathrm{OF}_t + \beta_2 \,\mathrm{Vol}_t + \beta_3 \,\mathrm{Depth}_t + \gamma \mathrm{Frag}_t^{\mathrm{DDC}}\Big), \tag{19}$$

where $g(\cdot)$ is a logistic (or probit) link, OF_t is minute-level signed order flow, Vol_t is a realized or EWMA volatility proxy, and Depth_t is top-of-book depth. This specification is not intended as a structural model of trading; rather, it is a diagnostic test of incremental information content. The key hypothesis is $\gamma > 0$: DDC fragility remains informative for future stress even after controlling for conventional order-flow and volatility measures.

To strengthen interpretation, we also conduct a state-conditional comparison. Specifically, we stratify observations into bins of comparable signed order flow (e.g., deciles of OF_t). Within each bin, we compare forward return distributions between high- and low-DDC-fragility states. If DDC captures liquidity-conditioned propagation rather than merely the level of trading pressure, then conditional on similar OF_t, high-fragility states should exhibit substantially higher tail risk and larger distributional deviations from diffusion. This "same order flow, different outcomes" diagnostic directly tests whether DDC is capturing propagation risk beyond what is implied by order flow alone.

We test whether propagation fragility provides incremental monitoring content beyond standard order-flow and impact controls. We estimate a nested logistic specification that conditions on signed order flow, a realized-volatility proxy, and top-of-book depth, and then add the DDC

fragility measure. The incremental contribution of DDC fragility is strongly supported: the likelihood-ratio statistic for adding the fragility measure to a specification that already included signed order flow, realized volatility, and top-of-book depth is 45,391 (df = 1), strongly rejecting the null of no incremental information. Operationally, including DDC fragility improves precision–recall performance: PR-AUC increases from 0.7765 in the baseline specification to 0.8011 when DDC fragility is included (ΔPR-AUC = 0.0245). These results indicate that propagation fragility contains forward-looking information about exceedance outcomes that is not subsumed by contemporaneous order flow, volatility, and depth.

Table 11 provides a state-conditional diagnostic that directly addresses whether propagation fragility merely reflects contemporaneous trading pressure. Within each signed order-flow decile, observations are split into low- and high-fragility states based on the bottom and top quintiles of the DDC propagation measure. If DDC were simply a transformation of order flow, forward outcomes should be similar within each decile.

The evidence strongly rejects this hypothesis. Across all order-flow deciles, high-fragility states are associated with substantially higher forward exceedance rates than low-fragility states. The differences are economically large, ranging from approximately 8 percentage points in intermediate deciles to over 30 percentage points in the lowest and highest order-flow deciles. Thus, even when trading pressure is held fixed, forward outcomes differ sharply depending on the propagation regime.

Distributional comparisons reinforce this conclusion. Within every order-flow bin, the upper quantiles of the forward return distribution shift upward in high-fragility states. For example, the 90th percentile of the forward maximum absolute return is consistently higher under high DDC fragility, and the quantile-distance metric indicates pronounced distributional separation across states. Nonparametric tests strongly reject equality of the forward return distributions between low- and high-fragility observations within each decile.

Taken together, the results in Table 11 demonstrate a "same order flow, different outcomes" phenomenon: within every decile of signed order flow, high-fragility states are associated with substantially higher forward exceedance rates than low-fragility states, with differences ranging from 8 to 32 percentage points. Conditional on similar trading pressure, markets in high-fragility states exhibit materially greater tail risk. This finding confirms that DDC fragility captures liquidity-conditioned propagation risk that is not subsumed by standard order-flow or impact measures.

10. Machine-learning augmentation: Proof of Concept

As a proof-of-concept extension, we augment the parameter mapping with an LSTM neural network trained to infer propagation regime from high-dimensional microstructure data. We stress that this extension is not the paper's core contribution; the structural results in Sections 7–9 do not

depend on it. Its purpose is to show that the framework's parameters can be estimated from richer data than the parsimonious mapping in Section 5 permits, and that the resulting model generalises across markets without retraining.

The LSTM is trained on the April 2020 WTI oil crash and applied without modification to the May 6, 2010 E-mini flash crash and the August 2024 Japanese equity futures crash. Table 4 reports Wasserstein distances: the quantum-walk model achieves the lowest distance in both cases relative to fat-tailed and Gaussian benchmarks. A fuller treatment of the machine-learning augmentation, including sensitivity analysis for the training event choice and architecture details, is provided in Online Appendix D.

Figure 11 compares the observed price path to simulations generated by the LSTM–QW model. In the period preceding the dislocation, the model identifies a sustained increase in propagation persistence associated with declining market depth and increasingly one-sided order flow. The resulting simulations reproduce the rapid price adjustment more closely than diffusion-based benchmarks, despite the absence of any event-specific training.

While the timing and magnitude of the signal differ across episodes, the qualitative dynamics are consistent. Liquidity withdrawal increases the effective propagation horizon, and directional imbalance governs the sign of the price movement. These results support the interpretation of flash events as manifestations of a common liquidity-driven propagation mechanism rather than as market-specific anomalies.

We evaluate the machine-learning–augmented quantum-walk framework using a flash crash in Japanese equity futures markets in August 2024. This episode provides a stringent test of external validity, as it occurs well outside the sample period used for model development and differs from U.S. markets in terms of trading hours, participant composition, and institutional structure.

As in previous application, the machine-learning layer infers the prevailing propagation regime, while price dynamics are generated by sampling from the resulting quantum-walk conditional return distributions. Figure 12 compares the observed price path to simulations generated by the LSTM–QW model. In the period preceding the dislocation, the model identifies a sustained increase in propagation persistence associated with declining market depth and increasingly one-sided order flow. The resulting simulations capture the rapid price adjustment more closely than diffusion-based benchmarks, despite the absence of any market-specific training.

While the timing and magnitude of the signal differ from U.S. events, the qualitative dynamics are the same. Liquidity withdrawal increases the effective propagation horizon, and directional imbalance governs the sign of the price movement. These findings reinforce the interpretation of flash events as manifestations of a common liquidity-driven propagation mechanism.

Table 4 reports the distributional fit of the machine-learning–augmented quantum-walk model relative to classical benchmarks using the Wasserstein distance. For both the May 6, 2010 E-

mini flash crash and the August 2024 Japanese equity futures crash, the quantum-walk specification achieves the lowest Wasserstein distance among the competing models. This indicates that, conditional on the same machine-learning–inferred propagation regime, the quantum-walk kernel provides a closer approximation to the empirical return distribution than either skewed Gaussian or fat-tailed alternatives. In the E-mini case, the improvement is economically meaningful: classical benchmarks struggle to reproduce both the magnitude and persistence of price movements observed during the crash window, while the quantum walk better captures the joint presence of large steps and directional propagation.

The absolute magnitude of the Wasserstein distance differs substantially across the two events and should not be interpreted as a measure of cross-market severity or model quality. In particular, the Japanese flash crash exhibits much larger Wasserstein distances than the U.S. episode because price changes in that event span a considerably wider range and occur under a higher-volatility regime. Wasserstein distance is scale dependent: when the support of the return distribution expands, the cost of transporting probability mass between distributions increases mechanically. For this reason, Wasserstein distances are comparable within an event across models, but not across events with different price scales, volatility levels, or contract specifications. Importantly, within each event the relative ranking of models is stable, with the quantum-walk framework consistently delivering the closest distributional fit despite large differences in absolute scale.

11. Discussion and Conclusion

This paper documents new empirical regularities in the timing and propagation of extreme price movements and proposes a depth-driven cascade model — implemented using a quantum-walk mathematical framework — as a parsimonious representation of liquidity-driven stress dynamics. The central finding is that extreme market events are characterized not only by large price changes, but by rapid and persistent propagation of trading pressure when liquidity becomes constrained.

The key distinction between macro-driven volatility episodes and microstructure-driven stress events is documented formally in Section 9.2. The COVID-19 crisis generated extreme realized volatility but unremarkable propagation fragility, because large exogenous shocks drove rapid price discovery in a market that retained functional liquidity between moves. Flash crash episodes generated elevated propagation fragility at moderate realized volatility, because the mechanism of price adjustment itself was disrupted by liquidity withdrawal. These two types of stress call for different monitoring tools and potentially different policy responses.

Implications for market monitoring and design: existing safeguards — price-based circuit breakers and volatility interruptions — are primarily triggered by large price movements and may therefore respond with a delay to emerging liquidity stress. The propagation framework highlights that fragility can arise from rapid propagation under constrained liquidity even when real-

ized volatility remains moderate. Monitoring propagation-based indicators derived from order-book states could provide exchanges and regulators with earlier signals of microstructure stress, enabling more targeted interventions.

References

Aït-Sahalia, Y., J. Cacho-Diaz, and R. J. Laeven (2015). Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics 117*(3), 585–606.

Ang, A. and A. Timmermann (2012). Regime changes and financial markets. *Annu. Rev. Financ. Econ. 4*(1), 313–337.

Bacry, E., I. Mastromatteo, and J.-F. Muzy (2015). Hawkes processes in finance. *Market Microstructure and Liquidity 1*(01), 1550005.

Bouchaud, J.-P. (2011). The endogenous dynamics of markets: price impact, feedback loops and instabilities. *Lessons from the 2008 Crisis*, 57–156.

Cespa, G. and X. Vives (2015). The beauty contest and short-term trading. *The Journal of Finance 70*(5), 2099–2154.

Chordia, T., R. Roll, and A. Subrahmanyam (2002). Order imbalance, liquidity, and market returns. *Journal of Financial Economics 65*(1), 111–130.

Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica: journal of the Econometric Society 41*, 135–155.

Cont, R. and P. Tankov (2003). *Financial modelling with jump processes*. Chapman and Hall/CRC.

Deuskar, P. and T. C. Johnson (2011). Market liquidity and flow-driven risk. *The Review of Financial Studies 24*(3), 721–753.

Evans, M. D. D. and R. K. Lyons (2002). Order flow and exchange rate dynamics. *Review of Financial Studies 15*(2), 461–491.

Filimonov, V. and D. Sornette (2012). Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics 85*(5), 056108.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Gabaix, X., P. Gopikrishnan, V. Plerou, and H. E. Stanley (2003). A theory of power-law distributions in financial market fluctuations. *Nature 423*(6937), 267–270.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the econometric society*, 357–384.

Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance 46*(1), 179–207.

Hasbrouck, J. (2003). Intraday price formation in us equity index markets. *The Journal of Finance 58*(6), 2375–2400.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance 72*(3), 967–998.

Kyle, A. S. and A. A. Obizhaeva (2016). Market microstructure invariance: Empirical hypotheses. *Econometrica 84*(4), 1345–1404.

Kyle, A. S. and A. A. Obizhaeva (2023). Large bets and stock market crashes. *Review of Finance*, rfad008.

Laughlin, G., A. Aguirre, and J. Grundfest (2014). Information transmission between financial markets in chicago and new york. *Financial Review 49*(2), 283–312.

Lee, S. S. (2012). Jumps and information flow in financial markets. *The Review of Financial Studies 25*(2), 439–479.

Menkveld, A. J. and B. Z. Yueshen (2019). The flash crash: A cautionary tale about highly fragmented markets. *Management Science 65*(10), 4470–4488.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of financial economics 3*(1-2), 125–144.

Vayanos, D. (2004). Flight to quality, flight to liquidity, and the pricing of risk.

Vives, X. (2014). Strategic complementarity, fragility, and regulation. *The Review of Financial Studies 27*(12), 3547–3592.

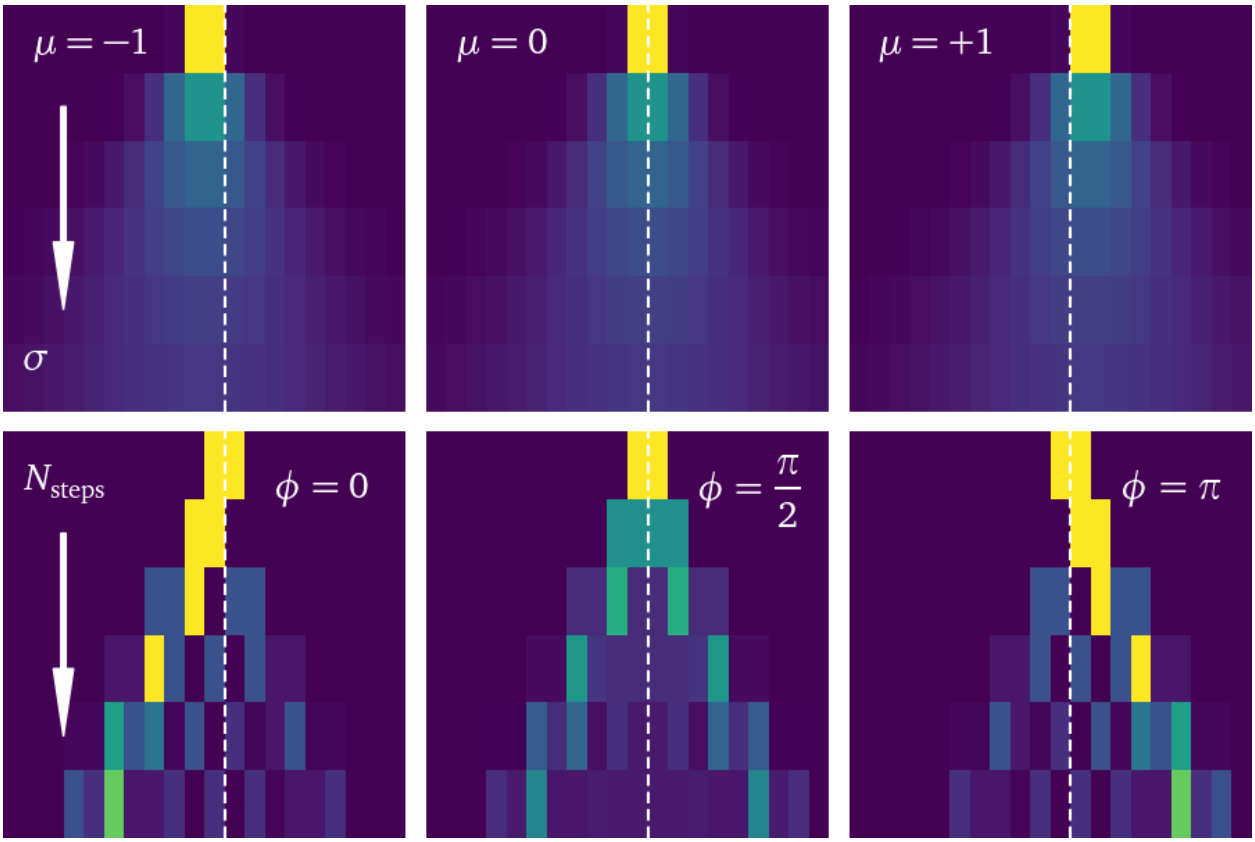


Figure 1: Probability laws generated by classical and quantum walks. The figure illustrates the probability distributions implied by a classical random walk and a quantum walk. Classical walks exhibit diffusive Gaussian scaling, while quantum walks generate non-Gaussian, asymmetric, and multi-modal distributions through path interference. These features allow the quantum-walk law to flexibly approximate return distributions observed during liquidity stress.

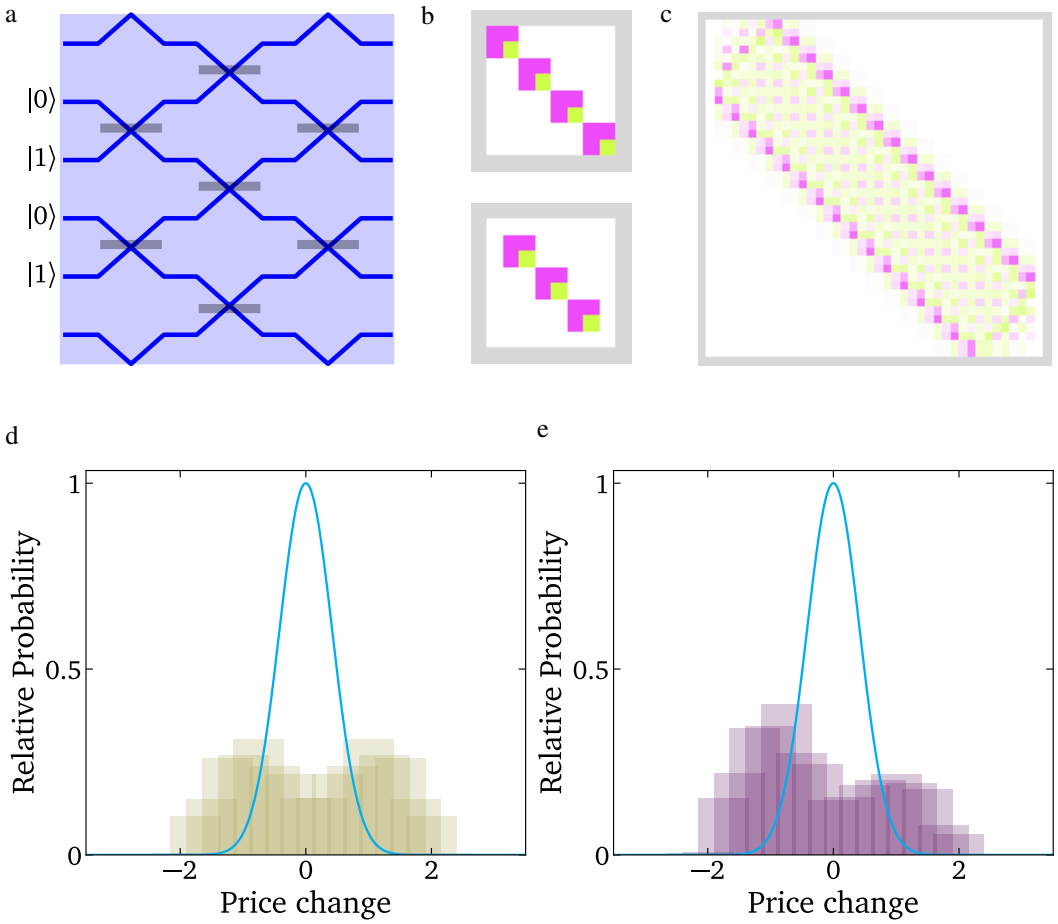


Figure 2: Quantum walk distributions: **a** Photonic circuit used for a QW simulation. The beam splitters (grey bars) are described by the Hadamard operator. **b** Matrices, U_1 and U_2 for alternate layers of the circuit. Purple and yellow indicate $\pm\frac{1}{\sqrt{2}}$, respectively. **c** Matrix for $N_{\text{qubits}} = 20$ and $N_{steps} = 6$ repeats of U_1U_2. **d** and **e** Output probability distributions for $N_{\text{qubits}} = 20$ and $N_{\text{steps}} = 4$ for two input states, eqn (**??**), with relative phase $\phi = \pi/2$ and $\phi = \pi/3$, respectively. We perform a 5-point moving average on the output. The distribution in **e** is the same as that used in Figure 3 right inset.

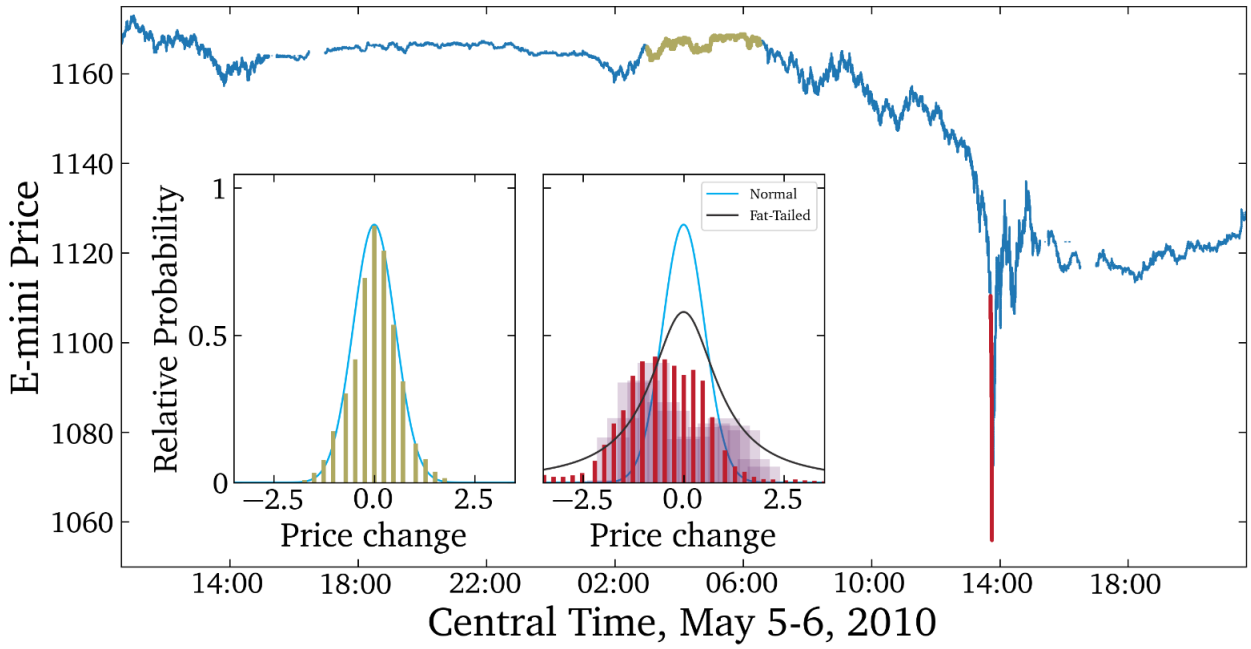


Figure 3: E-mini S&P 500 futures prices around the May 6, 2010 Flash Crash. This figure plots intraday prices for the CME E-mini S&P 500 futures contract on May 5 and May 6, 2010. Prices evolve smoothly prior to the event before exhibiting a rapid and severe decline beginning shortly after 13:30 Central Time on May 6, followed by a partial recovery. The speed and magnitude of the dislocation illustrate the breakdown of diffusive price dynamics over very short horizons.

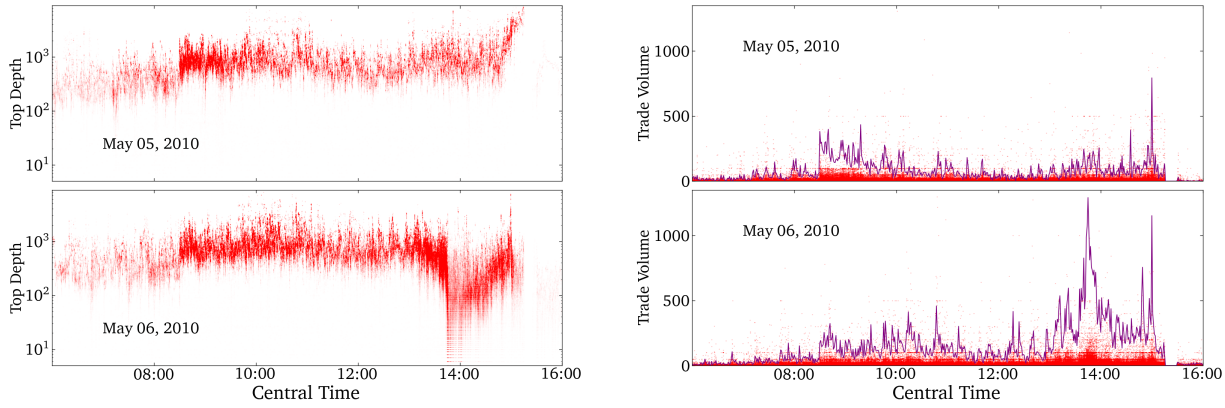


Figure 4: Top-of-book depth during the May 6, 2010 Flash Crash. This figure shows the evolution of top-of-book depth, defined as the total quantity available at the best bid and best ask, for the CME E-mini S&P 500 futures contract on May 5 and May 6, 2010. Market depth deteriorates sharply during the crash window, falling by an order of magnitude relative to earlier periods. The collapse in liquidity coincides with intense trading activity, indicating that the price dislocation is associated with liquidity withdrawal rather than an absence of trades.

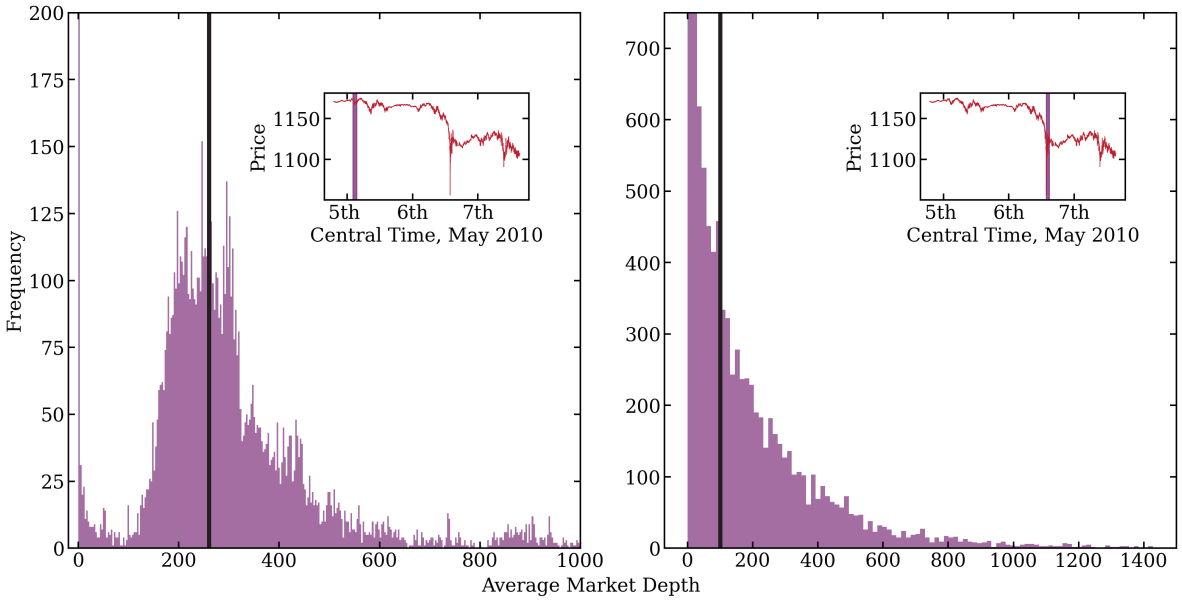


Figure 5: E-mini 2010 depth structure normal vs crash times: This figure compares the market depth structure of the E-mini during normal conditions and the May 2010 flash crash. The inset (red) displays price data for two periods: the day before the crash (May 5) and the crash day (May 6). To represent normal market conditions, we select a 5000-second window on May 5 (left panel). To illustrate crash conditions, we use the same time window on May 6 (right panel). For each period, the distribution of top-of-book depth is shown in purple, with the median depth marked by a black line. The histogram highlights the stark contrast: during normal conditions, market depth remains stable, whereas in the crash event, depth collapses to near zero, signaling a severe liquidity drain.

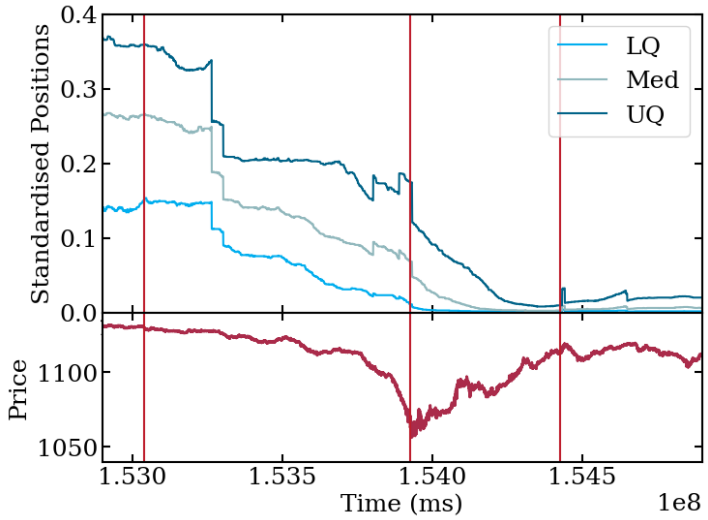


Figure 6: E-mini 2010 depth structure collapse: The lower plot (red) displays E-mini price data during the May 2010 flash crash, while the upper plot illustrates the normalized average top-of-book depth. The depth structure is further characterized by the lower quartile (LQ), median (Med), and upper quartile (UQ) positions. Red vertical lines mark the start, minimum, and recovery phases of the crash. Once the crash is triggered, market depth collapses, reaching a minimum before beginning its recovery. At the price turning point, there is a pronounced spike followed by a sharp drop in depth positions, indicating a temporary liquidity response before further dislocation.

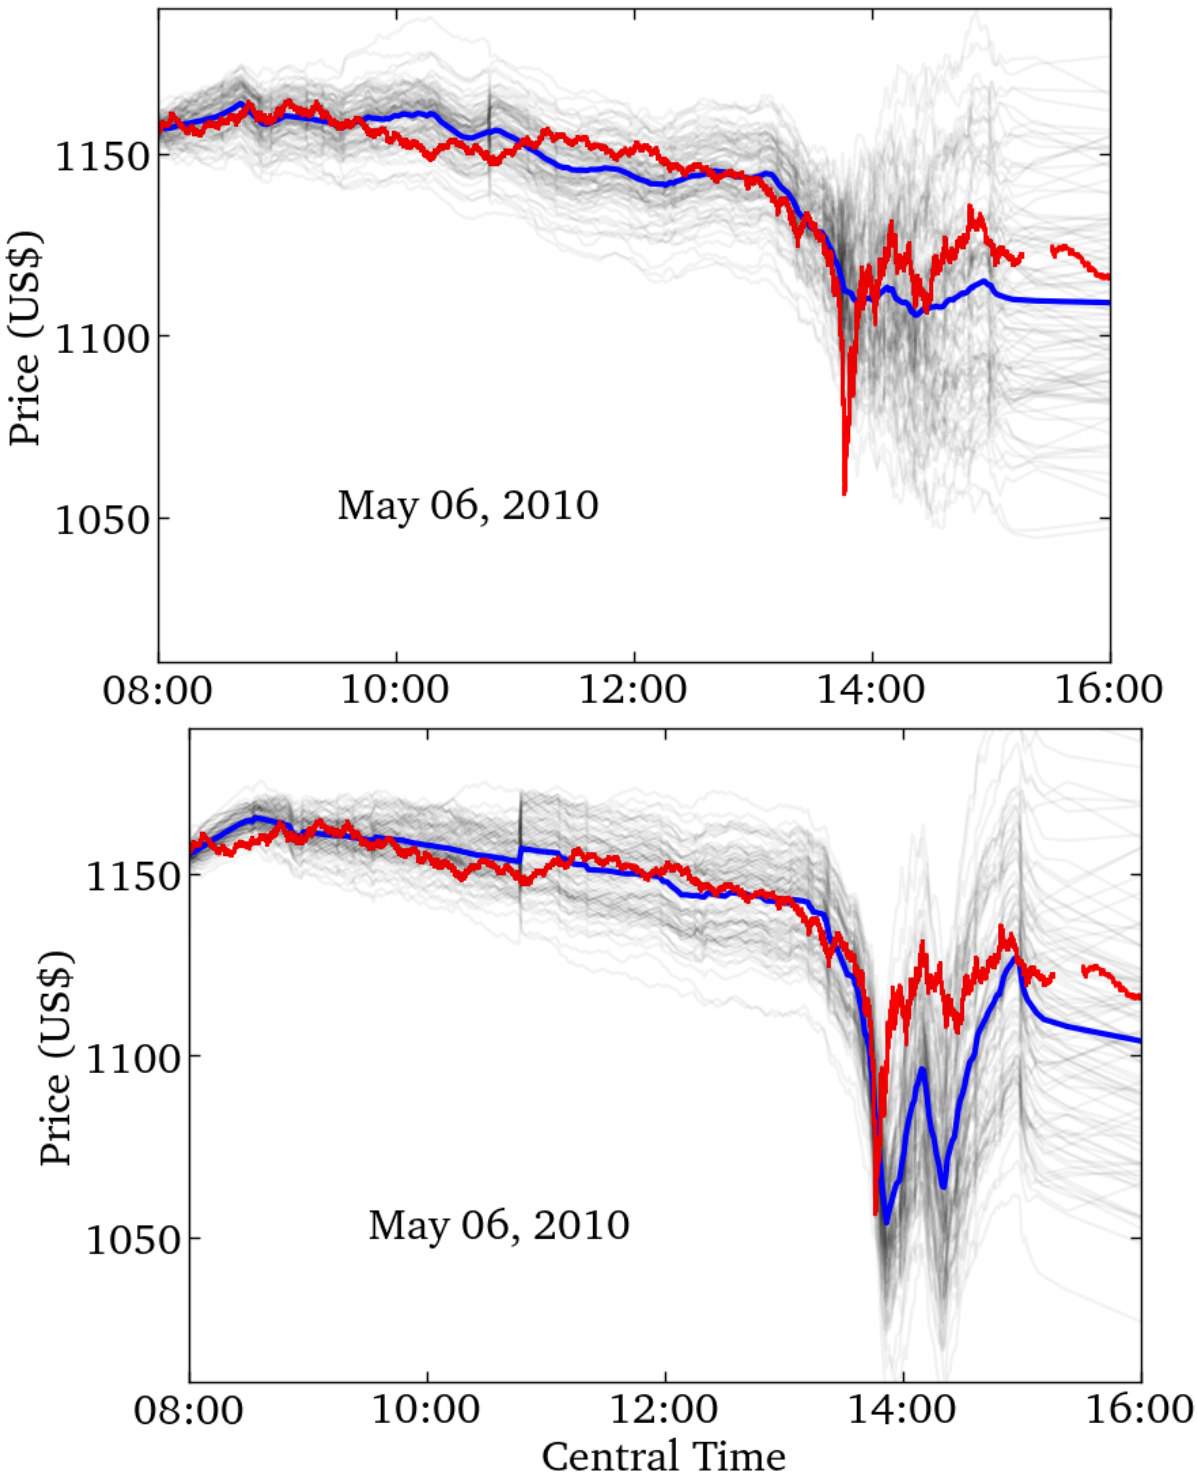


Figure 7: Conditional simulations around the 2010 flash crash. The figure compares price simulations generated by a classical random walk and a quantum walk, conditional on observed microstructure inputs. Both models track prices during normal conditions, but the classical walk fails to reproduce the abrupt transition during the crash, whereas the quantum walk adapts to the stressed regime.

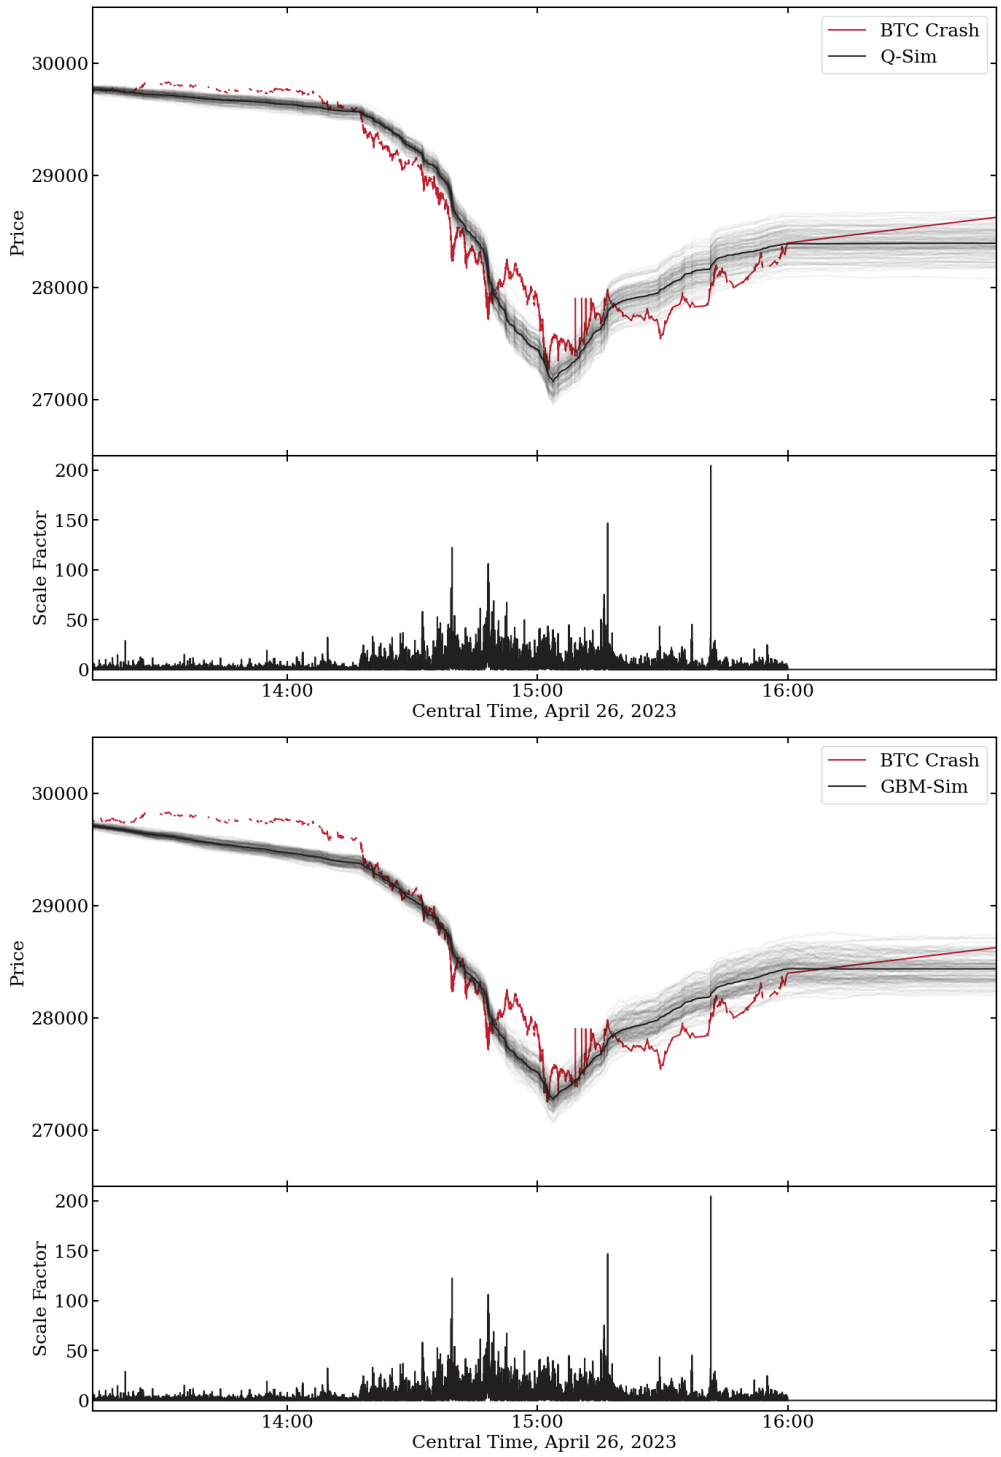


Figure 8: QW and GBM models of the bitcoin futures price dynamics on 23rd of April 2023: Top: The simulations (in black) alongside the raw price data (in red). For each model, the simulation is an average of 100 repeats (shown in grey). Bottom: The scale factor that guides the magnitude of the walk proportional to N_{steps} as detailed in section **??**.

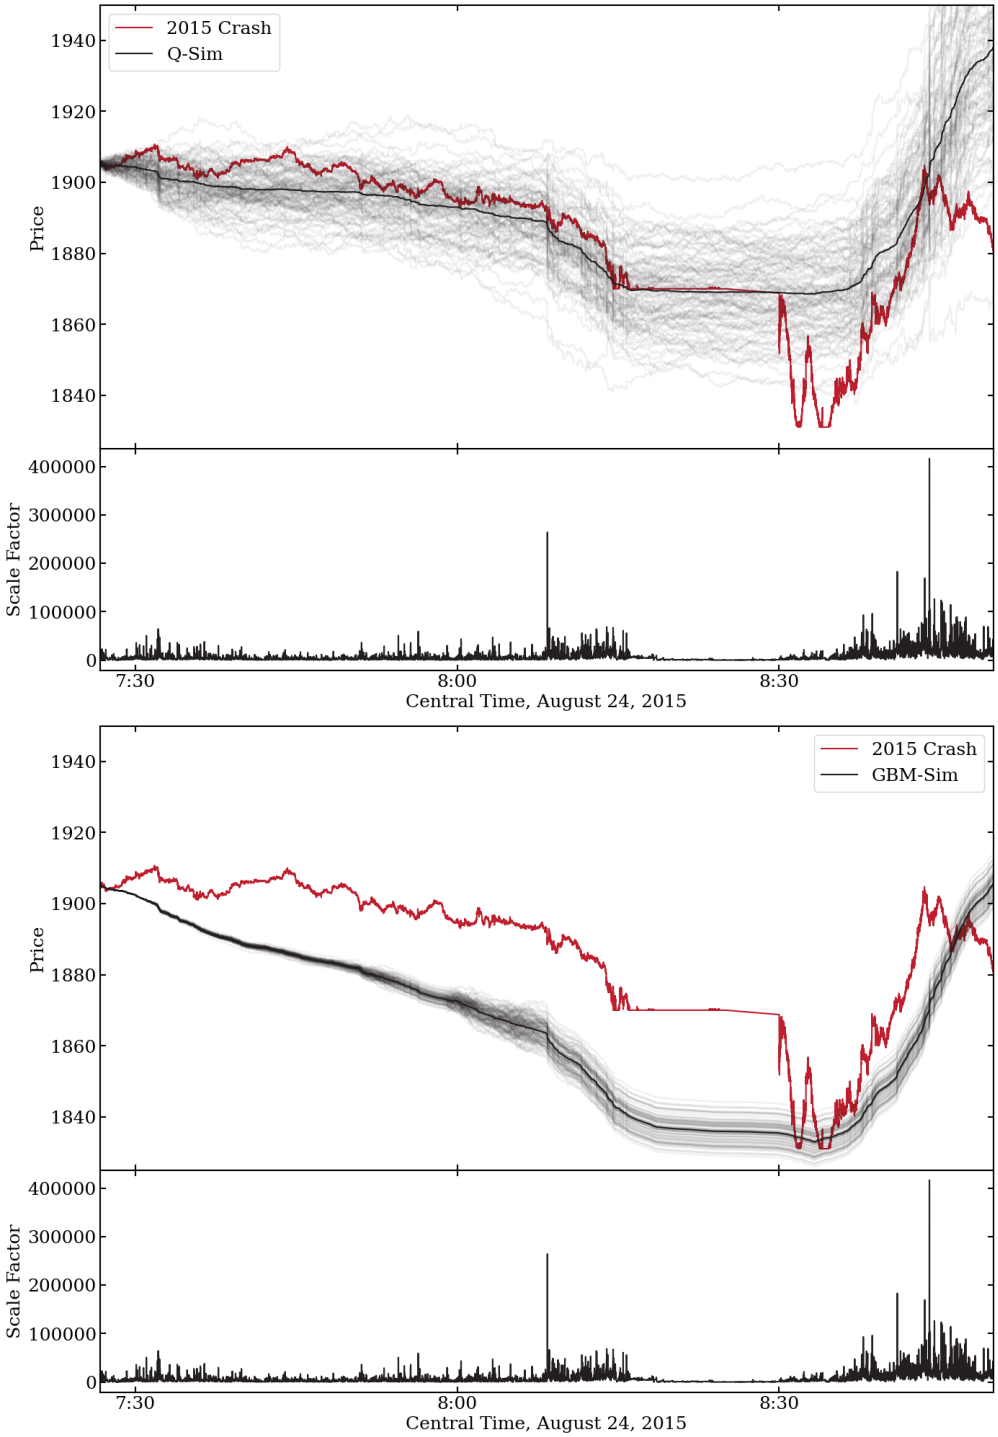


Figure 9: 2015 ETF simulation: Top Figure: The QW price simulation (in black) alongside the raw price data (in red). Bottom Figure: The GBM simulation (in black) alongside the raw price data (in red). For each Figure, the simulation is an average of 100 repeats (shown in grey); and, the scale factor that guides the magnitude of the walk. See Section **??**.

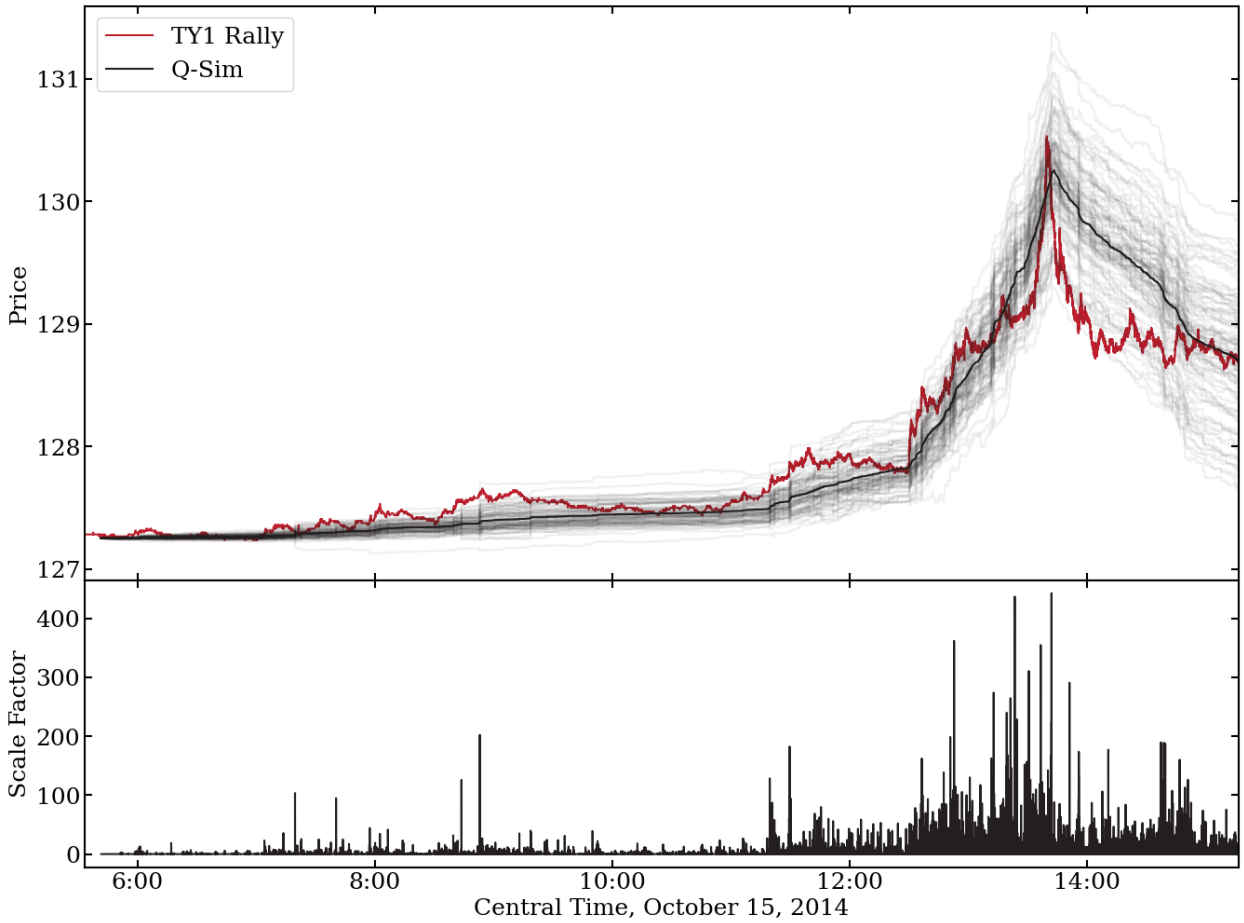


Figure 10: 2014 US Treasury flash rally simulation: The QW price simulation (in black) alongside the raw price data (in red). The QW simulation is an average of 100 repeats (shown in grey); and, the scale factor that guides the magnitude of the walk. See Sections **??** and **??**.

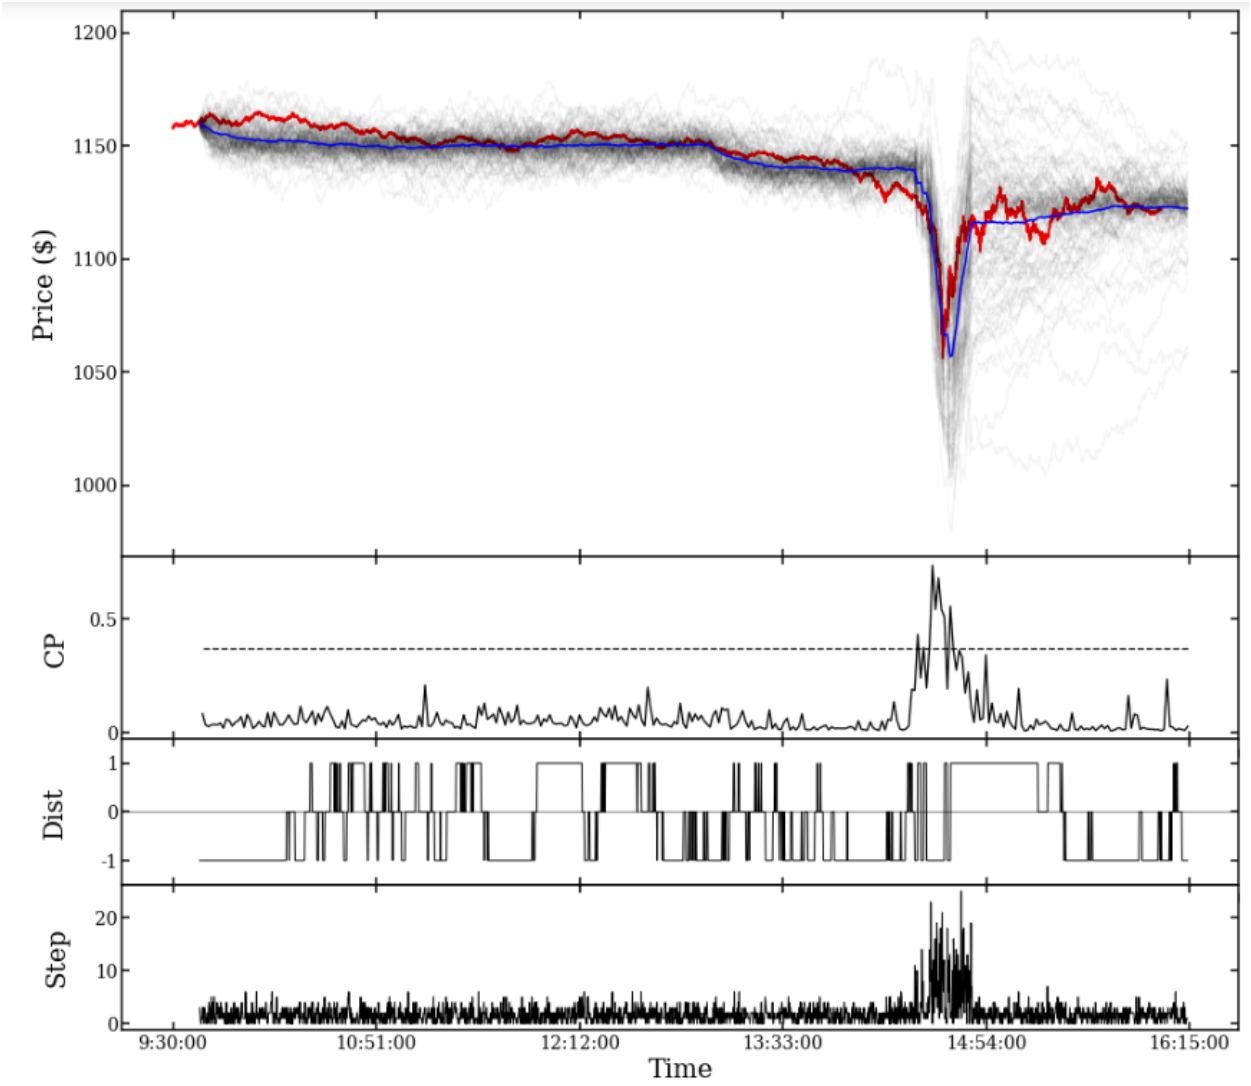


Figure 11: LSTM–QW simulation of the May 6, 2010 E-mini flash crash: From top to bottom. CME E-mini S&P 500 futures price data (red) and quantum-walk simulation (black). The simulation is generated using a LSTM model trained on the April 2020 WTI oil crash and applied to the E-mini market. Grey lines show an ensemble of repeated simulations. Below the main plot we report the inferred crash probability (CP). Also shown is the parameter Dist—derived from order-book imbalance—which governs the asymmetry parameter of the quantum walk. Finally, top-of-book depth and trading volume determine the number of propagation steps, N_{steps}, shown in the bottom panel.

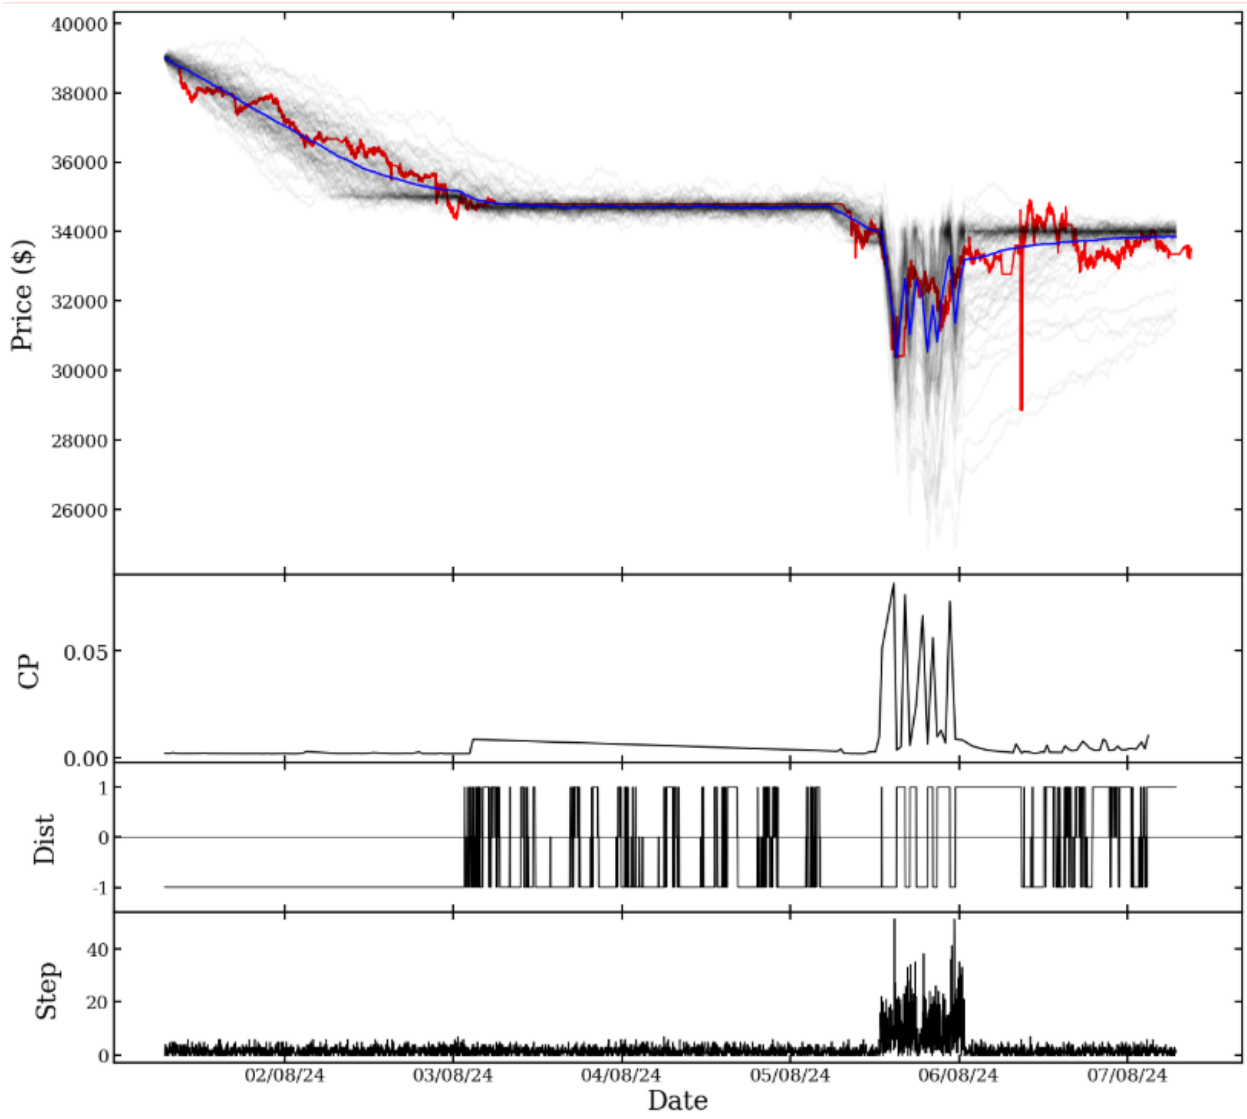


Figure 12: LSTM–QW simulation of the August 2024 Japanese equity futures flash crash: From top to bottom. Japanese equity futures price data (red) and quantum-walk simulation (black). The simulation is generated using a LSTM model trained on the April 2020 WTI oil crash and applied out-of-sample, without retraining or recalibration. Grey lines show an ensemble of repeated simulations. Below the main plot we report the inferred crash probability (CP). Also shown is the parameter Dist—derived from order-book imbalance—which governs the asymmetry parameter of the quantum walk. Finally, top-of-book depth and trading volume determine the number of propagation steps, N_{steps}, shown in the bottom panel.

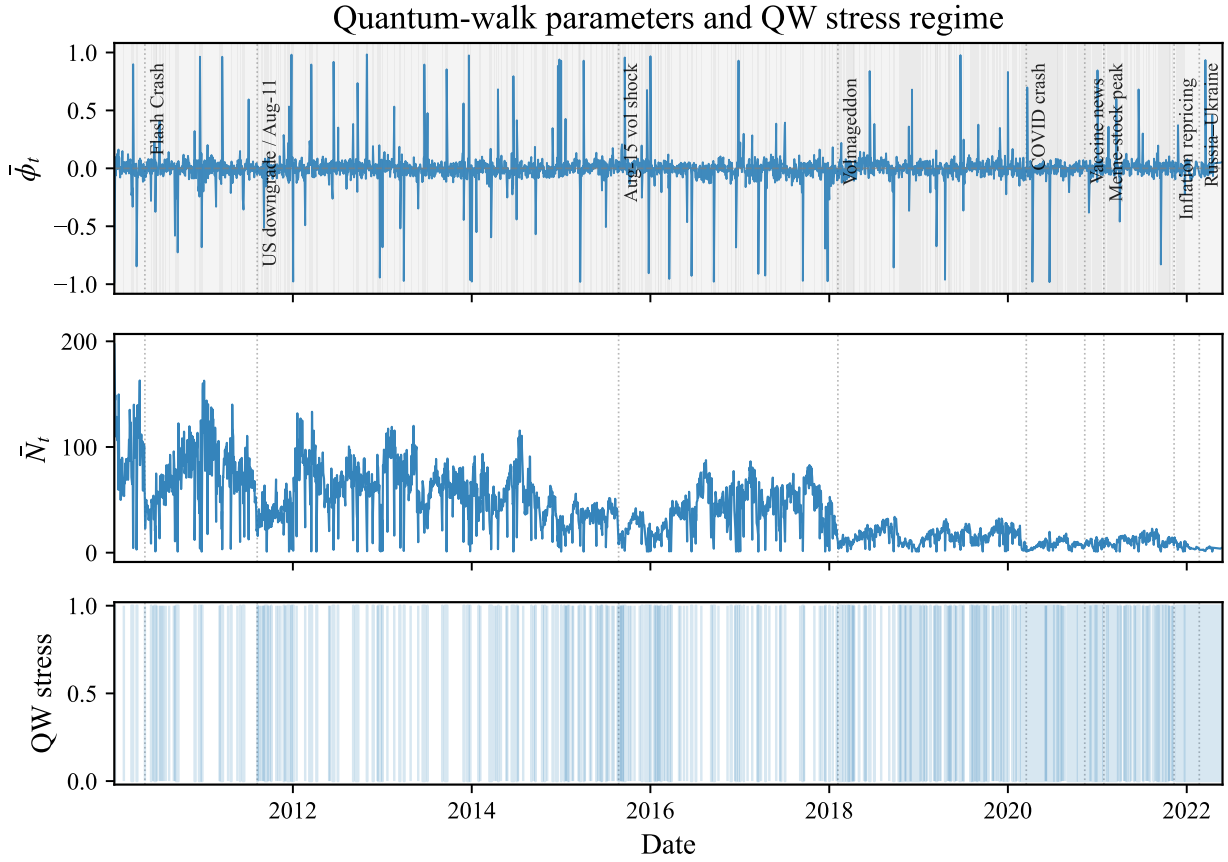


Figure 13: This figure uses the E-mini S&P 500 futures contract (E-mini) and plots Quantum-walk (QW) parameters and stress regime over the 2010–2022 sample. The inferred propagation depth $\tilde{N}_t$ and directional parameter $\tilde{\phi}_t$ vary substantially over time, while the QW stress indicator activates episodically around well-known dislocation periods and remains largely quiescent during normal conditions. This pattern supports the interpretation of QW stress as selective activation tied to liquidity-conditioned propagation and recovery dynamics rather than a generic high-volatility flag.

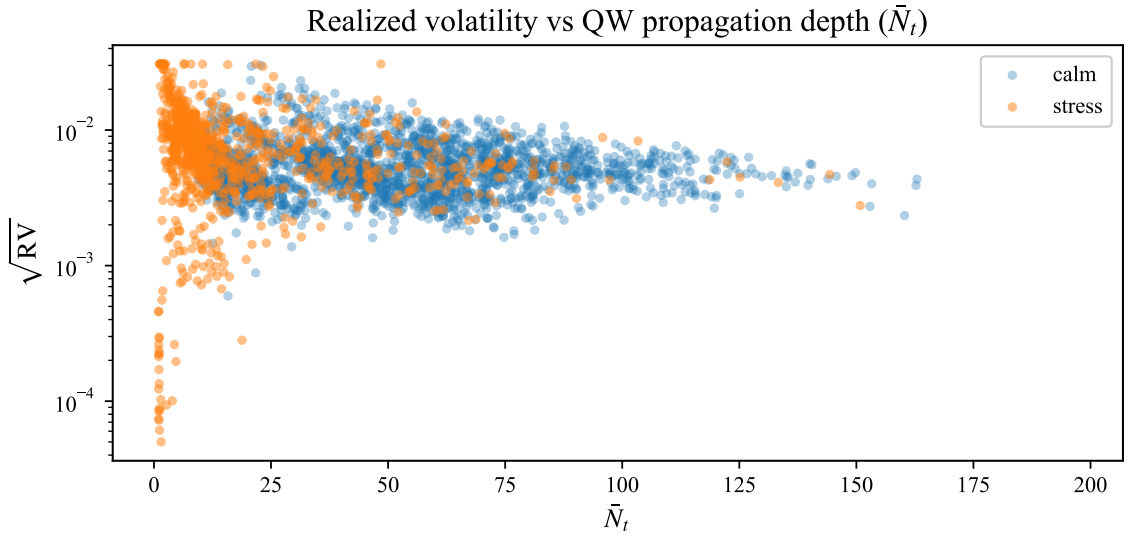


Figure 14: Realized volatility versus QW propagation depth $\tilde{N}_t$. While higher realized volatility is associated with higher propagation depth on average, the dispersion is large: stress and calm observations overlap across a wide range of volatility levels. This indicates that $\tilde{N}_t$ captures a distinct propagation-state component that is not mechanically determined by realized volatility.

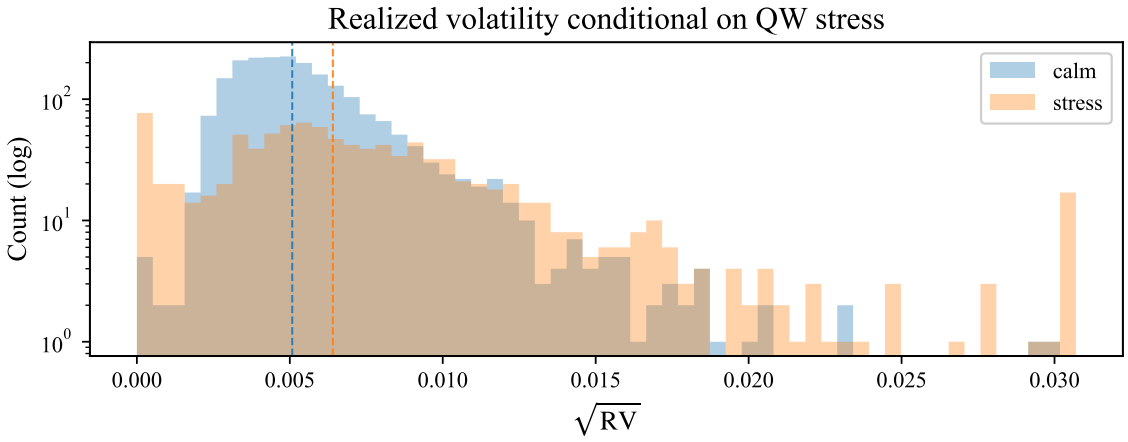


Figure 15: Realized volatility conditional on QW stress. The distribution of realized volatility shifts upward during QW-stress periods, but exhibits substantial overlap with the non-stress distribution. This overlap implies that realized volatility alone cannot identify the propagation regime, whereas QW stress isolates episodes consistent with liquidity-conditioned breakdowns of diffusive propagation.

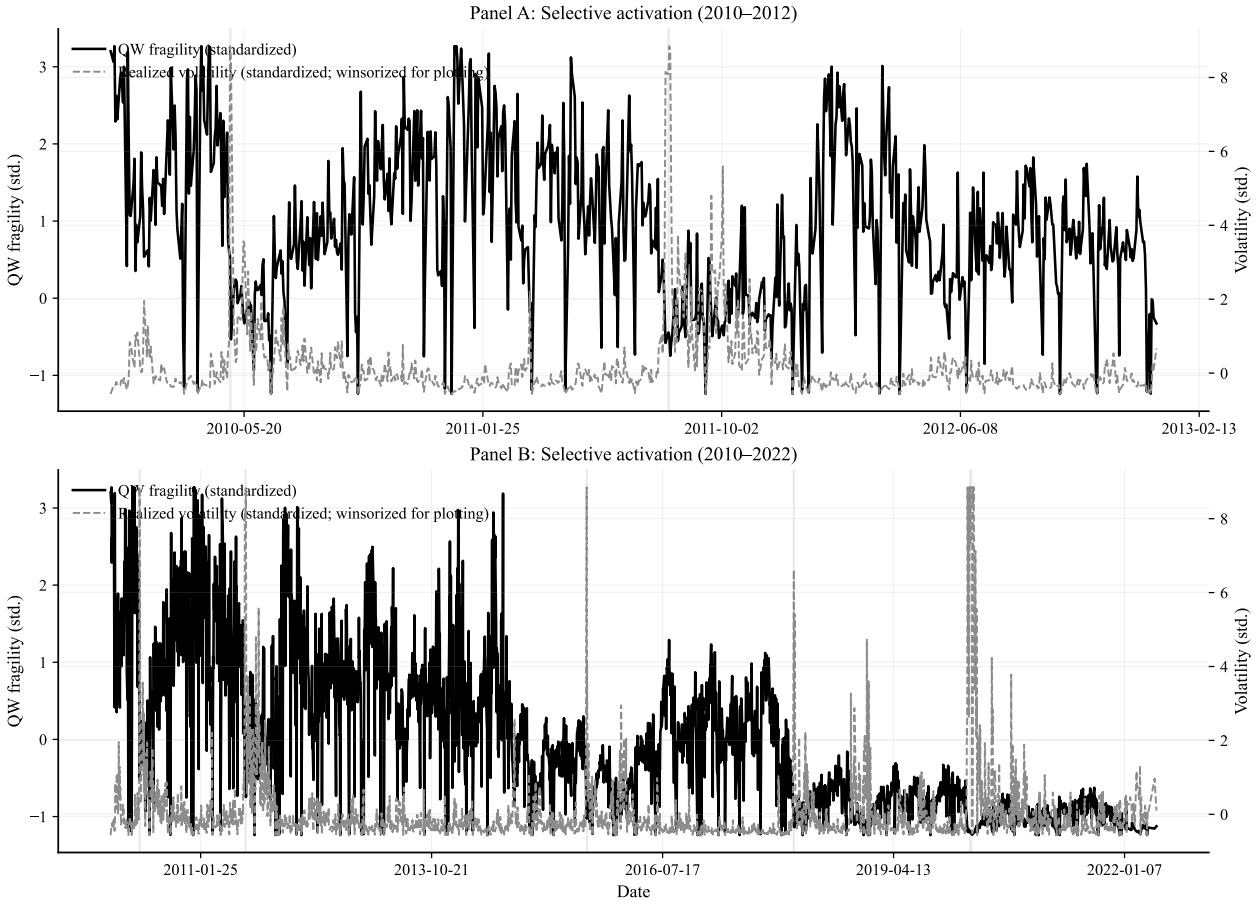


Figure 16: Selective activation of Quantum Walk fragility. The figure plots a standardized QW fragility measure (daily 90th percentile of propagation depth N_{steps}) alongside a standardized realized-volatility proxy. Panel A zooms into 2010–2012; Panel B shows the full 2010–2022 sample. Light-gray shading marks major dislocation windows. Volatility is winsorized for plotting clarity only. QW fragility activates episodically around microstructure-driven dislocations and does not mechanically track volatility, remaining muted during large macro volatility shocks such as COVID-19.

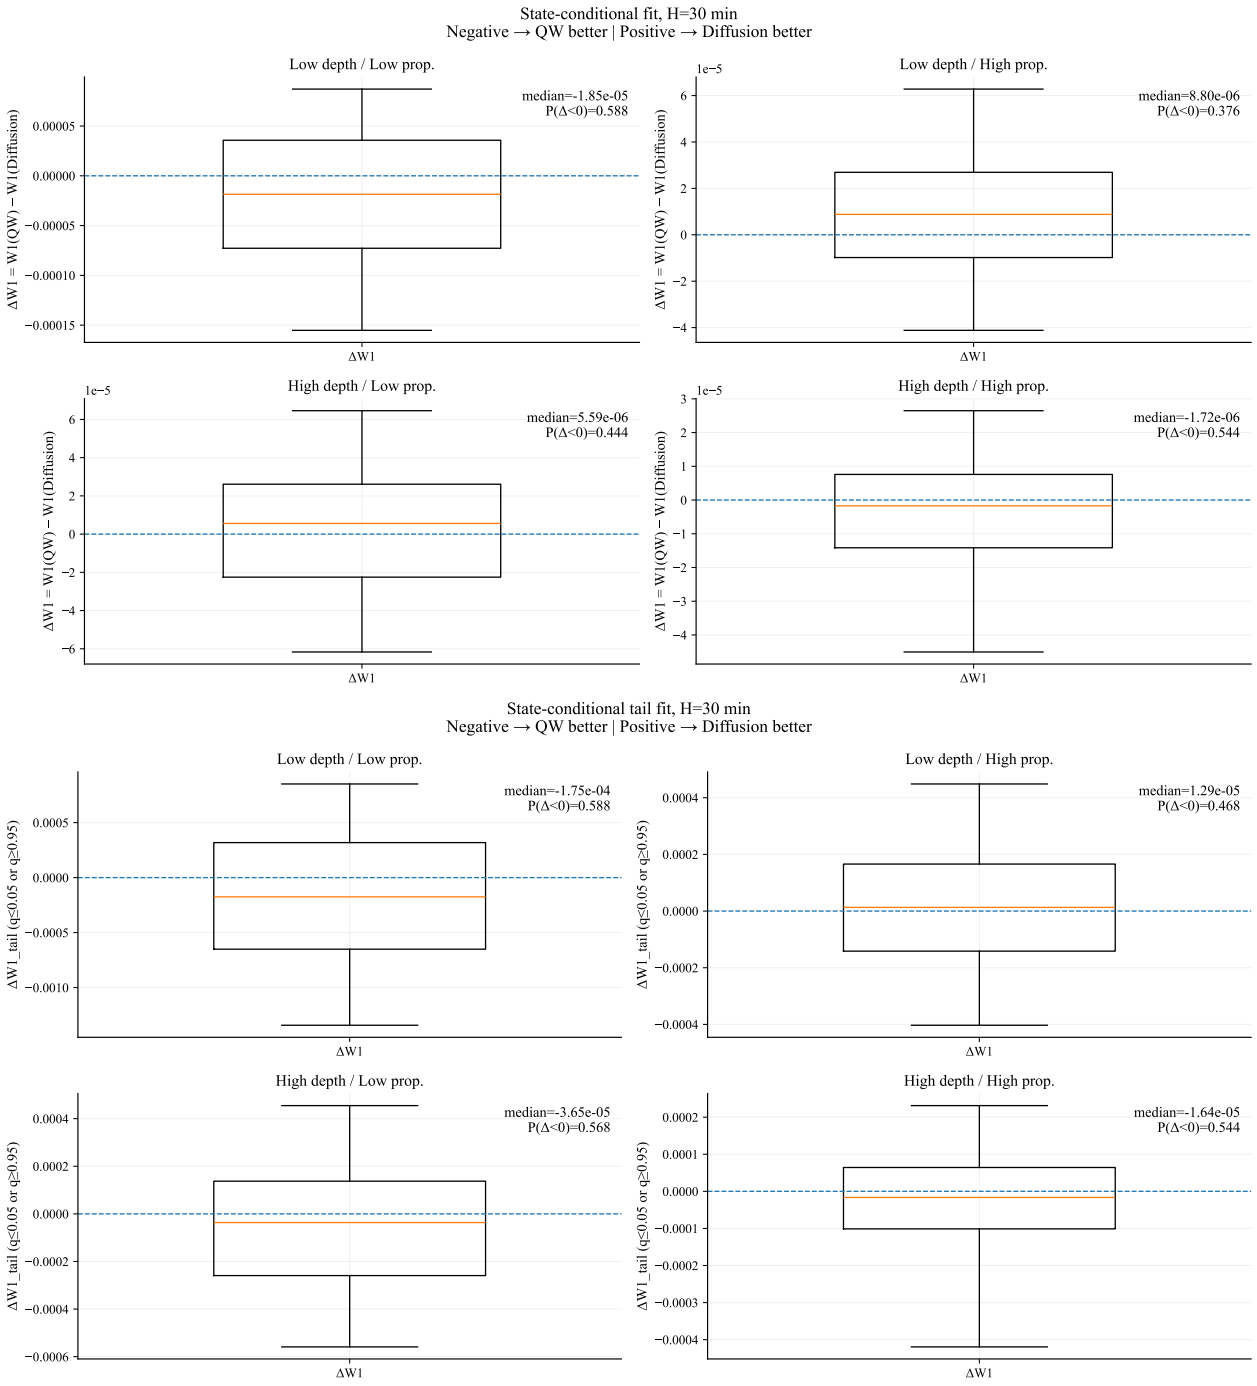


Figure 17: Regime-dependent distributional fit. Bootstrap boxplots of $\Delta W_1 = W_1^{QW} - W_1^{Diffusion}$ by liquidity–propagation regime for $H = 30$ minute terminal returns. Panel A reports full-distribution Wasserstein-1; Panel B reports tail-focused Wasserstein-1 computed using only extreme quantiles ($q \leq 0.05$ or $q \geq 0.95$). Negative values indicate superior QW fit. Each box summarizes $B = 250$ bootstrap resamples within each regime.

Table 1: Comparative features of benchmark frameworks versus liquidity-conditioned propagation

Feature	Diffusion	Jump Lévy	Hawkes	Markov switching	QW
Heavy tails / extremes	✓ †	✓	✓	✓	✓
Skewness (asymmetry)	✓ †	✓	✓ ‡	✓	✓
Volatility clustering	✓ *	✓ *	✓	✓	✓
Endogenous clustering of events	×	×	✓	✓ ‡	✓ ‡
Observable state conditioning on order book	✓ *	✓ *	✓ *	×	✓
Liquidity-conditioned propagation speed/persistence	×	×	×	×	✓
Selective activation without highest RV	×	×	× ‡	× ‡	✓

Notes. The table summarizes qualitative features across model classes. * indicates the feature can be accommodated through state-dependent parameters (e.g., time-varying volatility or covariates) but does not change the internal propagation mechanism. † requires additional structure (e.g., stochastic volatility, leverage effects, or non-Gaussian innovations). ‡ depends on the specific implementation (e.g., marked Hawkes, switching intensities, or state-dependent parameterizations). The defining feature of the quantum-walk framework is an observable, liquidity-conditioned propagation depth that governs the speed and persistence of directional adjustment within a fixed calendar interval.

Table 2: Mapping from microstructure variables to propagation model parameters

Variable	Economic interpretation	Role in model
Top-of-book depth D_t	Available liquidity	Determines propagation depth N_t
Order-book imbalance OIB_t	Directional trading pressure	Determines phase parameter ϕ_t
Returns r_t	Price change	Model output
Trade count / volume	Trading activity	Control variables

Table 3: Efficiency measure for direct market depth parameterization:

	WD	Mean of 100	Std Dev	Lower Bound	Upper Bound
Bitcoin Crash on 19th of April 2023)					
Quantum	5.11	5.35	0.40	4.02	6.05
Geometric Brownian Motion	5.91	6.75	0.41	5.11	7.06
ETF Crash on 24th of August 2015)					
Quantum	1.04	1.30	0.09	0.83	1.28
Geometric Brownian Motion	1.15	1.15	0.01	1.11	1.18
Japanese Crash on 5th of August 2024					
Quantum	0.00096	0.0017	5.03e-07	0.00172	0.00172
Geometric Brownian Motion	0.00146	0.00176	1.52e-10	0.00172	0.00176
US Treasury Rally on 15th of October 2014					
Quantum	0.026	0.030	0.001	0.022	0.031
Geometric Brownian Motion	0.033	0.034	0.001	0.031	0.034

Notes: We use the Wasserstein Distance (WD) to quantify the difference between the true and simulated distributions for each model, i.e., Quantum and Geometric Brownian Motion (GBM). We report the mean WD, the mean WD for 100 individual simulations, standard deviation, and the lower and upper bounds. Lower values of WD indicate lower dissimilarity between the true and simulated distribution. The results are reported using direct market depth parameterization for the 2023 Bitcoin flash crash, 2015 ETF flash crash, 2024 Japanese price crash, and the 2014 US Treasury flash rally.

Table 4: Machine-learning–augmented quantum-walk efficiency measure

	WD	Mean (100 runs)	Std. Dev.	Lower Bound	Upper Bound	90% CI
E-mini S&P 500 flash crash (May 6, 2010)						
Quantum walk	7.10	12.83	6.03	2.72	31.46	± 1.09
Fat-tailed benchmark	13.31	14.63	4.51	6.28	25.50	± 0.81
Normal benchmark	14.30	15.82	7.48	6.03	38.06	± 1.35
Japanese equity futures flash crash (August 5, 2024)						
Quantum walk	207.98	446.71	254.45	98.93	1415.95	± 45.80
Fat-tailed benchmark	1746.29	1757.14	290.98	1114.96	2722.96	± 52.38
Normal benchmark	826.91	877.46	199.92	598.40	1536.78	± 35.99

Notes: This table reports distributional fit for machine-learning–augmented simulations using the Wasserstein distance (WD), which quantifies the discrepancy between empirical and simulated conditional return distributions. Results are shown for the quantum-walk model and for classical benchmarks with fat-tailed and Gaussian distributions. For each event, we report the WD computed using the mean simulation, the average WD across 100 independent simulation runs, the corresponding standard deviation, lower and upper bounds, and 90% confidence intervals. Lower values of WD indicate a closer match between simulated and observed price dynamics.

Table 5: Comparison of mechanisms for extreme price movements

Model class	**Source of extremes**	**Role of liquidity**	**Propagation dynamics**
Jump–diffusion / Lévy	Exogenous shocks or large jumps	Not modeled explicitly	Instantaneous jumps; diffusive otherwise
Hawkes / self-exciting	Endogenous clustering of events	Indirect (via order flow)	Intensity-driven cascades
Markov-switching	Latent regime transitions	Implicit, not observable	Regime-dependent volatility or jumps
Quantum walk (this paper)	Liquidity-conditioned propagation	Explicit and state-dependent	Fast, asymmetric, depth-driven propagation

Table 6: Selective activation of Quantum Walk stress

Statistic	Value
Spearman(RV, QW stress)	-5.84e-02
Kendall(RV, QW stress)	-4.77e-02
Spearman(N_{steps}^{p90}, RV)	-1.47e-01
Mann–Whitney p (within RV bins)	1.12e-22
LR-like statistic (RV vs RV+N_{steps}^{p90})	3.53e+02

Notes: QW stress is defined as the top 1% of trading days ranked by the daily 90th percentile of propagation depth N_{steps}. Rank correlations assess the relationship between QW stress and realized volatility (RV). The Mann–Whitney test compares propagation depth between stress and calm days after demeaning within realized-volatility deciles. The LR-like statistic measures the incremental explanatory power of propagation depth in a logistic classification of stress relative to realized volatility alone.

Table 7: Incidence of QW stress across realized-volatility deciles

Realized-volatility decile	Stress frequency
D1 (lowest RV)	0.0158
D2	0.0095
D3	0.0190
D4	0.0253
D5	0.0095
D6	0.0095
D7	0.0127
D8	0.0000
D9	0.0000
D10 (highest RV)	0.0000

Notes: The table reports the fraction of QW stress days within each realized-volatility (RV) decile, where deciles are formed from the daily RV distribution. QW stress is defined as days in the top 1% of the propagation-based fragility measure (daily $p90$ of N_{steps}). Stress incidence is concentrated in low-to-mid RV regimes (D1–D7) and is absent in the highest volatility deciles (D8–D10). This pattern indicates that QW stress activation is not a mechanical by-product of volatility spikes; for example, the extreme RV regime associated with the March 2020 crashes falls in the upper deciles yet does not trigger QW stress, consistent with jump-driven repricing and constrained endogenous propagation.

Table 8: Propagation depth by volatility decile and stress state

RV decile	Median N_{steps}^{p90} (calm)	Median N_{steps}^{p90} (stress)
D1	28.7	197.4
D2	59.0	185.6
D3	56.3	194.3
D4	57.9	191.8
D5	56.3	195.8
D6	54.7	183.9
D7	43.6	187.5

Notes: The table reports the median daily 90th-percentile propagation depth, N_{steps}^{p90}, separately for calm and stress days within each realized-volatility (RV) decile. Across all volatility bins, stress days exhibit dramatically higher propagation depth than calm days. Notably, stress observations are absent from the highest RV deciles (D8–D10), indicating that elevated propagation is not driven by extreme volatility episodes.

Table 9: Incremental information in classifying QW stress

Model	Average log-likelihood
Realized volatility only	-0.0566
Realized volatility + N_{steps}^{p90}	-0.0006
LR-like statistic	353.48

Notes: The table reports average in-sample log-likelihoods from logistic regressions predicting QW stress. The first specification includes realized volatility only, while the second augments the model with the propagation-depth measure N_{steps}^{p90}. The LR-like statistic equals $2T(\bar{\ell}_{\text{aug}} - \bar{\ell}_{\text{RV}})$, where T is the number of daily observations, and quantifies the incremental classification power of propagation depth beyond volatility. The large improvement reflects strong separation between stress and calm states even after controlling for volatility.

Table 10: Operational monitoring performance: Recall@K (top alerts). The table reports Recall@K for the rolling 2010–2022 evaluation. Recall@K_1, K_2, and K_3 correspond to the top 1,000, 5,000, and 10,000 alerts, respectively.

Model	Recall@K_1	Recall@K_2	Recall@K_3
Quantum Walk (QW)	0.0049	0.0241	0.0480
Gaussian diffusion	0.0048	0.0229	0.0457
Student-t (heavy-tail)	0.0047	0.0221	0.0435
Merton jump–diffusion	0.0042	0.0231	0.0453
Hawkes-style intensity	0.0045	0.0226	0.0448
Markov-switching	—	—	—

Notes. Recall@K measures the fraction of realized extreme events captured among the top K alerts ranked by each model's monitoring score. Higher values indicate better concentration of true stress events among limited monitoring capacity.

Table 11: Same order flow, different outcomes.

OF decile	N_{lo}	N_{hi}	Event rate (lo)	Event rate (hi)	Diff	q_{90}^{lo}	q_{90}^{hi}	Quantile dist.
D1	162,223	162,199	0.516	0.845	+0.328	4.600	5.266	1.300
D2	162,212	162,239	0.686	0.837	+0.152	4.997	5.367	0.992
D3	162,197	162,202	0.681	0.845	+0.164	4.902	5.445	0.919
D4	162,210	162,202	0.765	0.844	+0.080	5.099	5.445	0.809
D5	171,781	171,772	0.531	0.835	+0.305	4.656	5.444	1.278
D6	152,637	152,636	0.751	0.835	+0.084	5.040	5.461	0.720
D7	162,202	162,214	0.759	0.839	+0.080	5.056	5.468	0.746
D8	162,196	162,199	0.708	0.838	+0.130	5.004	5.471	0.938
D9	162,200	162,215	0.692	0.840	+0.148	4.983	5.446	1.089
D10	162,194	162,206	0.486	0.837	+0.352	4.614	5.340	1.776

Notes. Within each signed order-flow decile (based on OF_t), observations are split into low vs. high QW fragility states (bottom 20% vs. top 20% of $N_{\text{steps},t}$ within the decile). The table reports the forward exceedance event rate and distributional diagnostics for the forward maximum absolute return over horizon H (as defined in the text). "Quantile dist." denotes the mean absolute difference between selected quantiles of the forward-max-$|r|$ distributions. Mann–Whitney tests strongly reject equality of forward outcomes across high vs. low QW fragility states within each order-flow decile (all $p \approx 0$ given the large sample).