Homogeneity, Systemic Risk, and the Regulation of Automated Trading in India, the US, and the EU

Siwon Lee*

Abstract

Automated trading maintains a major presence in execution across advanced economies and is spreading through emerging markets. We ask when the homogeneity of automated trading becomes systemic and how strongly a regulator must respond. In a continuous-time model of one trading day, a homogeneous sector trades a common signal at heterogeneous speeds, and the price impact of its staggered orders turns the superposition into a feedback, so the amplification is derived from the frequency structure rather than assumed. Systemic risk is the chance that the dislocation grows disproportionately far, judged against the same sector trading on diverse signals. Fragility is governed by penetration, the scale of the sector relative to residual depth, not by the adoption share, and penetration and homogeneity are complements that carry the dislocation to a threshold. We model a position limit, a margin, and a circuit breaker read as a call auction that removes the amplifier, and the intensity needed to stay below the threshold rises with penetration, so a market that imports a mature jurisdiction's rules inherits a constraint calibrated to the wrong scale, while between markets of like depth the stricter regulator holds the higher threshold. We calibrate the model to the United States, the European Union, and India, and the results make an economic case for preemptive and jurisdiction-specific regulatory calibration in emerging markets.

Keywords: Algorithmic trading, Artificial intelligence, Systemic risk, Market microstructure, Signal correlation, Market penetration, Financial regulation, Emerging markets.

JEL Classification: G01, G12, G14, G15, G18, G23, D82.

*Ph.D. Candidate in Economics, Washington University in St.Louis. Email: l.siwon@wustl.edu

1 Introduction

Automated trading powered by artificial intelligence or by classical algorithms has increased over the past two decades, and it now accounts for more than 70 percent of equity volume in the United States (Hendershott, Jones, and Menkveld 2011; Securities and Exchange Commission 2020). The same pattern has reached the emerging market of India, where in the most recent fiscal year automated systems made up close to 54 percent of equity cash turnover and 70 percent of equity derivatives turnover (National Stock Exchange of India 2025). This shift changes the informational structure of the market, the institutions that regulate it, and the participants who trade in it. For the regulator it raises a new problem, because the tools built to regulate human traders may not reach a market that machines increasingly run. When multiple automated participants share information and rely on highly correlated algorithms, their forecasts converge on a common direction. Although they operate at different frequencies and do not act at the same instant, they execute trades in the same way. Consequently, their orders arrive in succession and move the price one way. The impact of these orders builds through the trading day instead of clearing at once, and this accumulation is a source of systemic fragility rather than a feature of orderly price discovery. This paper addresses two fundamental questions. The first is how a regulator should contain this systemic risk. First, how should a regulator contain this specific form of systemic risk? Second, what is the exact mechanism through which the homogeneity of automated forecasts generates this risk when trading is distributed across multiple frequencies?

Regulation is the first of these questions, and recent work has brought its difficulty into focus. Reinforcement-learning algorithms can sustain supra-competitive profits through independent convergence without explicit agreement or communication. Since a constraint aimed at containing one regulatory harm can unintentionally worsen another while reducing market efficiency, designing an effective rule demands a precise understanding of how these algorithms interact with the market microstructure (Dou, Goldstein, and Ji 2025). Turning to artificial intelligence for regulatory enforcement offers no simple resolution. While AI monitors individual systems effectively, it cannot supply the system-wide causal judgment required for macroprudential prevention (Danielsson, Macrae, and Uthemann

2022). Jurisdictions have nonetheless begun to act, and they differ in how far the regulator reaches into algorithmic trading before an order is filled, with that reach rising from the broker, to the firm, to the individual strategy.[1] We model regulation through the instruments a regulator actually controls, namely circuit breakers, margins, and position limits, and we let these instruments act on the market from the outside rather than on the internal workings of the algorithms.

The same three markets differ as much in the limits their exchanges place on price itself. In the United States, a market-wide circuit breaker halts all trading when the broad index falls 7, 13, or 20 percent from the prior close. Separately, a limit-up and limit-down band pauses any single stock moving outside a specified range around its recent price, without a fixed daily cap. India pairs a market-wide breaker at 10, 15, and 20 percent of the index with hard daily price bands of 2 to 20 percent on individual stocks, beyond which trading is constrained until the limit framework is reset. The European Union sets no market-wide breaker and instead requires each venue to interrupt a stock, often into an auction, once its price would leave a dynamic or static range, the most decentralized of the three and the lightest in how far it constrains price. These structural limits are the empirical counterpart of the circuit-breaker band in our model, and they form a second axis alongside the regulation of algorithms. Along the two dimensions our model emphasizes, the oversight of algorithms and the constraint on price movement, India is the strictest of the three, while the United States and the European Union trade places, the United States lighter on the design of algorithms and the European Union lighter on the movement of prices. A market strict on both, as India is, is a different object from one loose on each, and the model is built to price that difference.

The variable that this regulation acts on has a long established literature. A dynamic model of insider trading in sequential auctions examines the informational content of prices, the liquidity of the market, and the value of private information (Kyle 1985). In

[1] In the United States the Market Access Rule places financial and regulatory checks at the broker that provides access, while no authority reviews a strategy in advance (Securities and Exchange Commission 2010). In the European Union a firm that trades algorithmically must notify its supervisor and validate its systems through an annual self-assessment under MiFID II, and a firm that trades at high frequency must be authorized as an investment firm (European Union 2014; European Commission 2017). In India algorithms must be approved by the exchange and given a unique identifier before they can trade, a requirement applied to institutional algorithms from 2012 and extended in 2025 to retail strategies above an order-frequency threshold (Securities and Exchange Board of India 2012, 2025).

its continuous-trading limit prices follow a Brownian motion and the depth of the market is constant over time. The constant depth is a feature of that equilibrium, not a focus of the analysis. A separate and empirical literature measures how forced trading moves prices. It studies the price pressure arising when institutional funds adjust positions in response to capital flows. This pressure is measured by the fraction of a stock's average volume originating from funds under severe capital flows (Coval and Stafford 2007). The price impact of a forced flow therefore depends on its size relative to the depth available to absorb it, and we call the size of the automated sector relative to that depth its penetration. A recent theory ties this price impact to artificial intelligence, with a systemic risk multiplier that moves with the adoption share (Meng and Chen 2026). We take penetration to be a separate quantity, not the adoption share but the size of the automated sector relative to the depth that absorbs it, shaped by the regulation a market imposes and by the depth of the market. A developing market and a developed market with the same adoption share can therefore differ in penetration, and their exposure to systemic risk follows from that difference.

A recent and active body of research studies this systemic risk. The leading account builds the amplification through a feedback coupling that operates through performative prediction, algorithmic herding, and cognitive dependency (Meng and Chen 2026). The amplification there is assumed as a feedback channel, meaning the systemic effect is inherited from that assumption rather than derived from the structure of automated trading itself. Consequently, the speed of trading and the correlated following that builds within the day play no role. The empirical record, for its part, documents the pattern without a structural account of how it arises. What is missing is a measurable statement of how the homogeneity and the speed of automated trading generate systemic risk.

A separate line of work needs no feedback channel. It shows that trading at different frequencies shapes returns on its own. Infrequent rebalancing generates return autocorrelation, a positive continuation of returns over intraday horizons, and seasonality in returns (Bogousslavsky 2016). A related finding is that the first half-hour return predicts the last half-hour return, an intraday momentum that strengthens on volatile, high-volume, and macroeconomic news days (Gao et al. 2018). These results trace to a common source. Information and trades arrive in a staggered format across participants operating at differ-

ent frequencies. This staggering produces persistent, directional movements in prices. We carry this mechanism into automated trading. Algorithms share data and signals, so they are homogeneous in the direction they take. They range from high-frequency execution to infrequent portfolio reallocation, so they are heterogeneous in speed. Their correlated trades superpose within the day rather than offsetting one another. Specifically, the fast cohort acts on the price increment built by the slow cohort, meaning the heterogeneity of speed compounds the move rather than damping it. This is the channel through which the homogeneity and the speed of automated trading become systemic risk.

We study a single trading day in continuous time. A small news shock moves the price, and an automated sector that reads a common signal trades on it in one direction. Two parameters describe the sector. Its homogeneity is how alike its members read the signal, and its penetration is its size relative to the depth that would absorb it. The sector trades at two speeds. A slow cohort reacts to the move accumulated since the open and sustains a trend, and a fast cohort reacts to the latest change and amplifies it, so the order flow regenerates the very signal that produced it and a small shock can grow into a one-sided move that the slow arrival of diverse liquidity cannot reverse. A single coupling governs the strength of this loop, rising with penetration, speed, and homogeneity. Once this coupling coefficient exceeds one, the move grows exponentially without reverting. Exhaustion of the margin buffer triggers forced liquidation, accelerating the price move in either direction. High-frequency trading simply represents the region of maximum coupling and execution speed rather than an independent risk channel.

We measure systemic risk as the probability that the price moves disproportionately far during the trading day. This disproportion is evaluated against a baseline where the same automated sector trades on diverse rather than common signals, allowing the measure to isolate the fragility created by trading similarity rather than the liquidity supplied by automation. To contain this risk, the regulator holds three instruments that target separate points in the feedback loop. These tools include a position limit to reduce the size of the fast leveraged sector, a margin requirement to delay forced liquidation, and a circuit breaker to halt trading once the price move hits a designated band. This breaker reopens the market through a call auction, which clears all frequencies together and strips the amplification from the price. We show that the regulatory intensity a market needs

rises with penetration, and that the binding instrument shifts toward the circuit breaker as penetration and speed make the move faster, which is the position an emerging market that automates quickly comes to occupy.

This paper makes two contributions, and it makes each both in closed form and in calibration. The first is a unified framework where the homogeneity, penetration, and regulation of an automated sector jointly determine a critical threshold. Within this framework, the divergence between advanced and developing markets depends entirely on where each jurisdiction sits relative to that critical level. Homogeneity carries the dislocation to a critical level at which the amplifier diverges. Market penetration measures the size of the automated sector relative to the depth that absorbs it, and this penetration lowers the critical threshold. Conversely, position limits and margins push the threshold back. Consequently, two markets with the same automated share can stand at entirely different distances from the threshold due to differences in their depth, regulation, and leverage. The account derives this from the frequency structure of correlated trading rather than assuming a feedback channel, and it treats penetration as a dynamic quantity driven by regulation rather than a fixed share of volume. Under this formulation, market fragility can rise or fall even with no change in the actual proportion of automated participants. The calibration to the United States, the European Union, and India puts the difference in numbers and shows why a developing market that imports a developed market's rules can still face the greater risk.

The second contribution is the circuit breaker, which the paper formalizes as a call auction. Because the auction clears every frequency at once, it removes the amplifier built by the continuous market. It resets the dislocation through a dynamic deflation, which is deepest precisely where the amplification is largest. As a result, the circuit breaker emerges as the sole instrument whose regulatory value rises alongside the escalating speed and penetration that threaten the market The calibration quantifies the distinct impact of the breaker between India and the United States. Our results show that thin and fast markets gain the most from such halts, highlighting the circuit breaker as the primary instrument an emerging market must rely on during automation.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 sets out the model. Section 4 derives its analytical solution and interprets the results.

Section 5 calibrates the model and discusses what the calibration shows. Section 6 compares emerging and developed markets and examines the effect of regulation. Section 7 concludes. The Appendix gives the detailed proofs.

2 Literature Review

This paper builds on theoretical and empirical work from several fields. A long line of market microstructure literature studies how informed trading sets prices and volume. This body of work analyzes these dynamics by examining price impact and adverse selection (Kyle 1985; Glosten and Milgrom 1985), evaluating the number and correlation of informed traders (Holden and Subrahmanyam 1992; Foster and Viswanathan 1996), and deriving the continuous-time equilibrium for a general value distribution (Back 1992). A related line treats information in prices, its transmission between the informed and the uninformed (Grossman and Stiglitz 1980) and its aggregation from noisy signals (Hellwig 1980; Diamond and Verrecchia 1981). We take from this literature the price impact of order flow and the correlation of signals, which becomes our homogeneity, and we depart from it in two ways. The order flow comes from a homogeneous automated sector rather than from a strategic insider, and the price that reveals the fundamental and steadies the market when traders are diverse instead carries it away when they agree.

A growing literature studies the systemic risk of homogeneous machines. Recent models focus on the harmonized risk frameworks of artificial intelligence across asset management and supervision (Danielsson, Macrae, and Uthemann 2022), as well as the feedback loops generated by correlated signals (Meng and Chen 2026). A related line studies how data and computation reshape markets. This shift manifests in the analytical turn from fundamentals toward demand and order-flow data (Farboodi and Veldkamp 2020), the speed advantages and lower adaptability of quantitative managers (Abis 2020), and the supra-competitive profits that automated traders sustain without collusion (Dou, Goldstein, and Ji 2025). On the empirical side, periodic rebalancing schedules generate distinct patterns of return autocorrelation (Bogousslavsky 2016), and returns from the first part of the trading day systematically predict those of the last (Gao et al. 2018), reflecting the intraday momentum pattern carried by our slow trading cohort. We share with this

literature the concern that homogeneous machines create fragility, and we depart from it in two ways. First, our fragility arises from a herding loop on price and volume rather than from performative prediction. Second, the scale of the quantitative sector is parameterized as an endogenous market penetration that the regulator can actively influence, rather than a fixed, immutable adoption share.

A further literature studies the market quality of automated trading and the channels through which forced flows move prices. Automated trading improves liquidity in equity markets (Hendershott, Jones, and Menkveld 2011) and improves price efficiency in foreign exchange markets (Chaboud et al. 2014). High-frequency trading aids price discovery and efficiency (Brogaard, Hendershott, and Riordan 2014), and automated trading improves liquidity and informational efficiency but increases short-term volatility (Boehmer, Fong, and Wu 2021). At the same time, the latency advantage of speed invites over-investment (Biais, Foucault, and Moinas 2015). The high-frequency intermediary loses on inventory but earns the spread (Menkveld 2013). These benefits in calm markets are why we keep the automated sector in place and set its homogeneity to zero to define the benchmark rather than removing the sector. The amplification under stress is the mechanism the model formalizes, as seen in the Flash Crash where high-frequency traders amplified the one-sided move rather than absorbing it (Kirilenko et al. 2017).

When dislocations are large, arbitrage can fail to correct them and can itself bear too much or too little risk (Shleifer and Vishny 1997; Gromb and Vayanos 2002). The transmission of forced flows is the channel the model builds on most directly. Liquidity can dry up suddenly because margins are destabilizing, rising when fundamental volatility rises and funding providers are imperfectly informed, so that market liquidity moves with volatility and with funding conditions (Brunnermeier and Pedersen 2009). Forced sales by distressed mutual funds, and inflow-driven purchases alike, push transaction prices away from fundamental value when they are widespread relative to available counterparties, since short-run demand curves for stocks are less than perfectly elastic (Coval and Stafford 2007). Hedge funds sell when their returns are poor (Ben-David, Franzoni, and Moussawi 2012), and an asset is fragile when concentrated ownership or correlated or volatile liquidity shocks among its owners make its demand susceptible to non-fundamental shifts (Greenwood and Thesmar 2011). These supply the forced-liquidation layer of the model,

where a margin breach turns a homogeneous position into correlated selling whose price impact is set by penetration against the residual depth that absorbs it.

A related literature studies the crowding of quantitative strategies. Crowding and leverage create externalities, so a rise in the number of sophisticated traders need not bring market efficiency (Stein 2009). As strategies become known, the traders that follow them crowd the same signals (McLean and Pontiff 2016). The quant event of August 2007 shows empirically that a market-wide deleveraging and a sudden withdrawal of market-making risk capital can produce significant losses (Khandani and Lo 2011). We read this crowding as homogeneity, and we view the resulting deleveraging and withdrawal of risk capital as the forced liquidation and thinning of heterogeneous depth that our model places at its center.

A final body of work concerns the regulation of automated trading, the variable our penetration responds to. The remedies studied are varied. A market-design response replaces continuous trading with frequent batch auctions to remove the mechanical arbitrage rent that continuous markets create (Budish, Cramton, and Shim 2015). Alternatively, capital constraints may be required when leverage drives the systemic problem (Stein 2009). Furthermore, because the benefits of automated trading are stronger in large stocks than in small ones, a uniform rule inevitably reaches unevenly across a market (Boehmer, Fong, and Wu 2021). The circuit breaker has a theory of its own. A trading halt can help a market reach the Walrasian outcome that continuous market and limit orders cannot provide, as neither allows traders to condition on the full information set revealed by a shock (Greenwald and Stein 1991). However, this benefit may not carry across linked markets. While a trading halt in the market where liquidity is concentrated reduces price variability and liquidity locally, order flow often shifts to satellite markets and increases their volatility. (Subrahmanyam 1994). Our reading of the breaker as an auction that admits the diverse liquidity the price dislocation had outrun is the favorable case made precise. Jones (2013) provides a comprehensive review of the theory, evidence, and regulatory questions raised by high-frequency trading. In terms of implementation, turning to artificial intelligence for the oversight itself faces strict structural limits. This technology serves micro-level supervision well but cannot establish the causal judgments required for macroprudential control (Danielsson, Macrae, and Uthemann 2022). Conse-

quently, actual regulatory jurisdictions differ sharply in how far the authority reaches into the system before an order is filled. This regulatory reach from the broker in the United States, to the firm in the European Union, to the individual strategy in India (Securities and Exchange Commission 2010; European Union 2014; European Commission 2017; Securities and Exchange Board of India 2012, 2025). No existing framework represents this gradient as a variable the authority can set. We model regulation through the primary instruments an authority actually controls, which include the circuit breaker, the margin, and the position limit. Each of these tools acts on the market from the outside. Under this setup, we address three critical questions. We determine how strong these constraints must be. Additionally, the paper demonstrates how these regulatory instruments must tighten as market penetration grows. Finally, we identify which specific tool the regulator must prioritize as penetration and speed accelerate the market move.

3 The Model

Trading takes place over a single day, from the open at time zero to the close at time T. Let $P(t)$ be the price and $P(0) = p_0$ a fixed constant. The open itself operates as a call auction that clears overnight orders and positions at a single price. This clearing mechanism is identical to the one the circuit breaker imposes within the trading day. Through this initial auction, the prior day's informational errors are settled, allowing the trading day to begin in an undislocated state.

$$y(t) = P(t) - p_0, \tag{1}$$

the accumulated intraday price impact at time t. A positive y is a price above the open and a negative y a price below it. The only randomness within the day is a small news shock carried by a standard Brownian motion W, and the question is when that shock grows into a disproportionate move.

The price moves for three reasons. News arrives as the shock. Automated order flow moves the price through its market impact. Heterogeneous liquidity arrives on the other side and pulls the price back toward the open. Writing $Q(t)$ for the cumulative net

automated order flow in the common direction,

$$dy(t) = \sigma\, dW(t) + \lambda\, dQ(t) - \kappa\, y(t)\, dt. \tag{2}$$

We follow Kyle (1985) in reading λ as the price impact per unit of order flow, the inverse of market depth. The coefficient κ is the rate at which slow and diverse liquidity absorbs a dislocation, the heterogeneous depth whose thinning leaves the market fragile. The coefficient σ scales the shock. We specify the automated order flow in (2) in two stages. The homogeneous feedback below runs through the day and is the baseline the paper solves first, and a second force joins it once the dislocation grows large, the forced liquidation a leveraged sector meets at a margin, which we add after the baseline is in hand.

The automated sector trades on a common signal. Trained on overlapping data and built on similar architectures, its members read that signal alike and trade the same way. The homogeneity $\rho \in [0, 1]$ is the degree to which they read it identically, the correlation of their forecast errors. At $\rho = 0$, their orders cancel in aggregate and no common pressure forms, while at ρ near one the sector moves as one block, the case the paper is about. The sensitivity χ turns a given signal into an order. The mass θ represents the absolute size of the automated sector. Its penetration measures this size relative to the residual depth available to absorb it, defined mathematically as $\lambda\theta$. Because the price impact parameter λ is the inverse of market depth, a thin market inherently experiences deeper penetration for any given sector size. Penetration is distinct from the adoption share, and two markets with the same automated share differ in penetration when their residual depth differs. It is $\lambda\theta$, not the mass θ alone, that enters the coupling below and that the propositions mean by penetration.

The sector is not uniform in speed. One part trades at high frequency, revising many times within the day, and one part trades slowly, reallocating as the day's move accumulates. We model two cohorts. A fast cohort of mass θ_f trades at frequency v_f and a slow cohort of mass θ_s at frequency v_s, with v_f far above v_s and $\theta = \theta_f + \theta_s$. The fast cohort is the high-frequency system that acts the instant its signal updates, and the slow cohort is the strategy that trades once an alert has accumulated. A continuous band of frequencies lies between them, and the two cohorts stand for its poles, the fast end that amplifies and

the slow end that sets the trend.

The homogeneity ρ, the sensitivity χ, the masses θ_f and θ_s, the frequencies v_f and v_s, the impact λ, and the absorption κ are independent primitives. The model does not tie a sector's synchronized behavior to its speed, its size, or the depth it trades against. Because of this independence, a comparative static can move one parameter while holding the rest constant. The propositions can therefore read the independent effects of homogeneity, speed, and penetration one at a time.

What each cohort reads is determined by its observation frequency. A trading system measures the price against the time window elapsed since its last update. Consequently, this observation window is short for the fast cohort and long for the slow cohort. The slow cohort reads the accumulated move from the open, the level y, the direction set early in the day that carries toward the close. The fast cohort reads the contemporaneous change, the increment dy over the instant. Each trades in the common direction with an intensity set by penetration, frequency, and homogeneity, so we take the order flow, before any forced liquidation, to be

$$dQ(t) = \theta_s v_s \chi \rho \, y(t) \, dt + \theta_f v_f \chi \rho \, dy(t). \tag{3}$$

The masses θ_s and θ_f represent the components of the dynamic mass $\theta(t)$ defined in (8), maintaining fixed levels between circuit-breaker halts. Because this mass shifts within the day only during a halt, the coupling mechanism treats these components as constants over each interval between halts. The slow term is a flow proportional to the level, and the fast term is proportional to the contemporaneous increment. The loop sits here. When a price rise is reinforced by accompanying volume, the sector interprets it as a repetition of the same signal and trades the same way again. Consequently, this market move regenerates the initial signal that produced it. The balance between the trend and its amplification is not chosen but set by the composition of the sector across speeds, a high-frequency market weighting the amplifier and a slower one the trend. The split between the two responses is not a separate assumption but a consequence of the difference in speed, since the level and the increment are the moves the slow and the fast cohorts read over their own windows, and the common signal fixes the direction while the homogeneity sets its

strength. The masses and the frequencies are free primitives, set in the calibration rather than estimated, so the empirical record of intraday momentum (Gao et al. 2018) and of high-frequency direction (Brogaard, Hendershott, and Riordan 2014) enters not as a measurement of them but as a check that the two responses take the form the difference in speed implies.

Substituting (3) into (2) gives $dy = \sigma\, dW + \alpha\, y\, dt + \beta\, dy - \kappa\, y\, dt$, with the slow and fast couplings α and β of (5). The fast term βdy carries the identical price increment that stands on the left side of the equation. This alignment occurs because the fast cohort trades on the contemporaneous change that its own order flow helps generate. Collecting these incremental terms onto the left side yields $(1-\beta)dy = (\alpha-\kappa)ydt + \sigma dW$. Dividing both sides by $1-\beta$ establishes the linear dynamics for the dislocation move.

$$dy(t) = \frac{\alpha - \kappa}{1-\beta}\, y(t)\, dt + \frac{\sigma}{1-\beta}\, dW(t), \tag{4}$$

with

$$\alpha = \lambda\chi\rho\,\theta_s v_s, \qquad \beta = \lambda\chi\rho\,\theta_f v_f. \tag{5}$$

Because the fast cohort responds to the same increment it helps create, β enters not as a separate dynamic but as the factor $1/(1-\beta)$ that multiplies both the trend and the shock. Write $G = (\alpha-\kappa)/(1-\beta)$ for the net growth rate and $\sigma_* = \sigma/(1-\beta)$ for the effective shock. The slow cohort is the engine of the trend and the fast cohort is its amplifier.

Stability turns on two ratios. The trend coupling

$$r = \frac{\alpha}{\kappa} = \frac{\lambda\chi\rho\,\theta_s v_s}{\kappa} \tag{6}$$

compares the slow cohort's amplification with the absorption opposing it, and $G = \kappa(r-1)/(1-\beta)$. When r is below one, the dislocation reverts and settles near the open, with stationary variance $\sigma_*^2/(2|G|)$. As r approaches one, the reversion vanishes and the variance grows without bound. Once r reaches one, the dislocation grows without reverting, causing a small shock to become a disproportionate move. The fast cohort sharpens this distortion through the amplifier $1/(1-\beta)$, which rises as β approaches one. Both ratios rise with penetration, frequency, and homogeneity, which represent the three

forces at the center of this paper. Together, they form the core content of the coupling mechanism that sets this model apart from a generic feedback multiplier.

A second force enters once the move is large. The funds carry leverage, and a dislocation past the margin buffer m forces them to liquidate. Liquidation closes positions in the direction the price has moved, pushing it further the same way, and the push is symmetric, since a long book is sold on a fall and a short book is bought back on a rise. Adding it,

$$dy(t) = G\,y(t)\,dt + \sigma_*\,dW(t) + \psi\,\theta_f\,\operatorname{sign}\big(y(t)\big)\,\big(|y(t)| - m\big)_+\,dt, \tag{7}$$

where $(x)_+ = \max\{x, 0\}$. The forced sale is itself order flow, so it passes through the price impact and the fast-cohort feedback like any order, and ψ collects that sensitivity at the level of the dislocation dynamics, so the term adds directly to (4) without a further factor of $1/(1-\beta)$. Below the buffer, this forced liquidation term is zero. Above the buffer, the drift gains an outward push that steepens the price move. Because this destabilizing push enters continuously at the boundary m, the mathematical slope stiffens smoothly without a discrete jump. The push scales with the leveraged fast mass θ_f, so penetration $\lambda\theta_f$ drives the steepening as well as the amplifier.

Volume can be incorporated explicitly by scaling the sensitivity χ with a multiplier that mirrors market activity. This modification introduces nonlinearity into the feedback loop, potentially causing the price to run away in finite time. Since volume tends to be heavier during market declines than during rises, this runaway effect tilts downward. Ultimately, we maintain the linear baseline because it already captures the core mechanism, leaving this nonlinearity as an extension.

The mass θ is dynamic. Regulation moves it and the sector's growth raises it, and it is not a fixed share of volume, so we write

$$\theta(t) = \theta_{\text{base}} + g(R, t), \tag{8}$$

where θ_{base} is the level the market's structure sets and g is the effect of the regulation vector $R = (\ell, m, b)$. Penetration $\lambda\theta$ moves with the mass, since λ is fixed within the day. The instruments move θ at three places and over three horizons. A position limit of

tightness ℓ lowers the deployable mass of the leveraged fast cohort,

$$\theta_f(\ell) = (1 - \ell)\,\theta_f^0, \tag{9}$$

where θ_f^0 is the unconstrained mass and $\ell = 0$ leaves it untouched, a durable change that lowers the amplifier, the fast part of the trend, and the forced-liquidation push together. A margin requirement sets the buffer m and touches the trigger alone. A circuit breaker halts trading at the first time the dislocation reaches a band b,

$$\tau = \inf\{\, t \geq 0 : |y(t)| \geq b \,\}, \tag{10}$$

and it is the instrument that moves θ within the day, since during the halt the automated flow stops and the price freezes, so θ falls for the length of the halt, and trading then reopens through a call auction that clears every frequency at once, strips the amplifier from price formation, and reopens the dislocation at a deflated level, the removal of the amplification developed in Section 4. The coupling mechanism reads the level of θ that prevails between halts. Specifically, the position limit alters that level durably, the circuit breaker moves it only for the duration of a halt, and the baseline θ_{base} expands across days as the sector grows. Consequently, these three instruments operate on three distinct time horizons, meaning that the required policy intensity must track the slowest among them. Operationally the within-day move of θ is a regime switch at the stopping time τ rather than a time-varying coefficient. The path is solved in pieces, the feedback dynamics at the prevailing θ before the halt and the deflated reopening at τ, so the baseline θ_{base} is held fixed within the day and varied across days for the scaling trap. The margin and the position limit slow the move, and only the halt removes the amplifier outright at the reopening, so the instruments are not interchangeable. Speed bumps that lower v_f and diversity requirements that lower ρ act on the same dynamics, but the regulator does not hold them directly or cannot verify them, and the baseline sets them aside.

The price impact plays two roles, and the second is where the model departs from Kyle (1985). In the price equation, λ maps order flow into the price at a date, the static price-formation map. In the dynamics, the same impact sits inside a feedback that turns a small shock into a growing trend. Kyle (1985) builds his dynamics from the gradual

revelation of an insider's information, with a constant impact and a price that converges to fundamentals. Ours come from a homogeneous sector that reads its own price impact as a fresh signal and trades on it again, a loop the static map does not contain. What this loop produces, and a static impact cannot express, is the growth rate G, a speed.

High-frequency trading is the sharp end of the mechanism, not a separate channel. It is where penetration, frequency, and homogeneity are highest, so β is closest to one and the move is steepest, and where slow heterogeneous liquidity is most thoroughly outrun. This gives the Flash Crash a structural reading, in which automated participants amplified a one-sided move (Kirilenko et al. 2017).

The systemic risk of the day is the probability that the dislocation grows disproportionately large at some point before the close. The event is the breach itself, not the closing price, since a probability read at one instant would miss a move that crossed the threshold at midday and returned. The measure therefore follows the running maximum of the path,

$$\mathrm{SR}(t) = \Pr\Big(\max_{0 \le s \le t} |y(s)| > \bar{c} \Big), \tag{11}$$

the probability of a breach by time t. It is cumulative, and its terminal value $\mathrm{SR}(T)$ is the probability that a disproportionate move occurred during the day. When r is at or above one, the maximum almost surely exceeds the threshold and SR approaches one. When the day is stable, SR follows the first passage of the linear dynamics. Specifically, this probability increases with the coupling, the amplifier, and penetration, while decreasing as the threshold rises or the horizon shortens. The circuit breaker enters here, since stopping the path at the band b caps the running maximum and lowers SR directly, which a point-in-time probability cannot express.

The dislocation is written for one asset, but the risk it measures is the market's. The automated sector reads a signal common across the cross-section and trades the same direction in every name that loads on it, so the one-sided move the model traces is not idiosyncratic but the representative path of the many assets the sector holds at once. A breach in the representative asset is the coincidence of breaches across the names the sector trades similarly. This happens at the moment their synchronized action carries a disproportionate move through the entire book rather than through a single security. It is in this sense that the measure is systemic. We carry one asset for tractability, and the

cross-asset contagion of a literal many-asset market we leave to extension.

What counts as disproportionate is fixed against the same sector trading on diverse rather than common signals. When the sector is diverse, its orders cancel and homogeneity is zero. Under these conditions, the dislocation follows the Ornstein–Uhlenbeck process $dy = -\kappa y\,dt + \sigma\,dW$ with stationary standard deviation

$$\sigma_0 = \frac{\sigma}{\sqrt{2\kappa}}, \tag{12}$$

the move the news alone would justify. The threshold is a multiple of it,

$$\bar{c} = k\,\sigma_0, \tag{13}$$

so a move is disproportionate when it exceeds by a factor k what a diverse sector facing the same news would produce, and k is fixed in the calibration against an empirical tail. Measuring against the diverse sector rather than a market with no automation is deliberate, since automation adds depth in calm markets and raises κ, so removing it would lower the benchmark and mix the cost of homogeneity with the loss of that depth. Holding the sector fixed and setting only its homogeneity to zero isolates the fragility that synchronized behavior creates. Under homogeneity, the dislocation has standard deviation $\sigma_0/\sqrt{(1-r)(1-\beta)}$, so the threshold is crossed far more often as r and β approach one. The threshold $\bar{c}$ defines the event and the band b is the policy trigger, and a breaker that acts in time sets b below $\bar{c}$.

Reported alongside SR is the amplification the feedback produces, the variance multiplier $1/(1-\beta)$ and the analogous $1/(1-r)$, which diverge as β or r approaches one. These are the static amplification factors reported in earlier work. By placing the path and the horizon on top of them, the measure developed here allows a circuit breaker, which acts on time, to enter the calculation.

The day is now described. The price equation and the two-cohort feedback fix the growth rate G and the amplifier $1/(1-\beta)$ through penetration, frequency, and homogeneity. The forced-liquidation layer adds the steepening past the buffer. The regulation rule sets the position limit, the margin, and the band. The following solution characterizes the growth rate and its comparative statics. Furthermore, it defines the first-passage

probability under each instrument and illustrates how the binding constraint shifts toward the circuit breaker as market volatility accelerates. Finally, it determines the necessary policy intensity a market must maintain as its sector expands.

4 Solution

We solve the model in closed form and collect every proof in the appendix, so the text here carries the results and what they mean. Below the margin buffer the dynamics of the previous section are linear,

$$dy(t) = G\,y(t)\,dt + \sigma_*\,dW(t), \qquad y(0) = 0, \tag{14}$$

with

$$G = \frac{\alpha - \kappa}{1-\beta} = \frac{\kappa(r-1)}{1-\beta}, \qquad \sigma_* = \frac{\sigma}{1-\beta}, \tag{15}$$

and $\alpha = \lambda\chi\rho\,\theta_s v_s$, $\beta = \lambda\chi\rho\,\theta_f v_f$, $r = \alpha/\kappa$. Both α and β are proportional to homogeneity, so we write

$$r = \rho\,A_s, \quad \beta = \rho\,B_f, \qquad A_s \equiv \frac{\lambda\chi\theta_s v_s}{\kappa}, \quad B_f \equiv \lambda\chi\theta_f v_f. \tag{16}$$

The constant A_s collects what the slow cohort contributes to the trend and B_f what the fast cohort contributes to the amplifier, and homogeneity scales both. The model is well-defined while $\beta < 1$, that is while $\rho < 1/B_f$, since the fast feedback cannot exceed the increment it responds to. Everything that follows reads off this one linear equation.

Lemma 1 (Explicit solution)**.** *The solution of* (14) *is the Gaussian process*

$$y(t) = \sigma_* \int_0^t e^{G(t-s)}\,dW(s), \tag{17}$$

with mean zero and variance

$$V(t) \equiv \mathrm{Var}\big(y(t)\big) = \sigma_*^2\,\frac{e^{2Gt}-1}{2G}. \tag{18}$$

The dislocation is Gaussian because the feedback is linear and the only shock is Brownian, so the whole path collapses to the single number $V(t)$. Every result below is a

statement about that variance.

Lemma 2 (Terminal tail)**.** *For any threshold $\bar{c} > 0$,*

$$\Pr\left(|y(T)| > \bar{c}\right) = 2\,\Phi\left(-\frac{\bar{c}}{\sqrt{V(T)}}\right), \tag{19}$$

where Φ is the standard normal distribution function, and the probability is strictly increasing in $V(T)$.

The systemic-risk measure $\mathrm{SR}(T) = \Pr(\max_{s\le T} |y(s)| > \bar{c})$ dominates this terminal tail. A path whose close lies beyond $\bar{c}$ has carried its running maximum beyond $\bar{c}$, so the event of the lemma is contained in the event of a breach, but the reverse fails, since a path can cross $\bar{c}$ at midday and return inside it by the close, a breach the terminal tail does not count. The terminal tail therefore establishes a strict lower bound on $\mathrm{SR}(T)$. This bound remains tight when the horizon is short but becomes loose when the path has sufficient time to cross the threshold and revert. We rely on this tail probability for our closed-form solutions because it depends on the price path entirely through the single variance metric $V(T)$. Given the strict monotonicity in $V(T)$, every comparative static of systemic risk regarding homogeneity, penetration, and speed is already captured within this tail. Our subsequent numerical calibration then recovers the running maximum that this analytical bound inherently understates. The force of the lemma is that the tail rises with the variance and with nothing else, which turns every question about a probability into a question about V.

Lemma 3 (Benchmark)**.** *At $\rho = 0$ the dislocation is the absorption process $dy = -\kappa y\,dt + \sigma\,dW$, stable for all penetrations and all frequencies, with variance $V(t) = \sigma_0^2(1-e^{-2\kappa t}) \to \sigma_0^2 = \sigma^2/(2\kappa)$, and with the threshold $\bar{c} = k\,\sigma_0$ the terminal exceedance is $2\Phi(-k)$, independent of penetration θ and of the frequencies v_s, v_f.*

The benchmark fixes the meaning of disproportion. This move represents the level justified by news alone under diverse liquidity. It is the exact level that a sector with no common signal would produce. This baseline does not change with market depth, speed, or automation. Therefore, we treat $2\Phi(-k)$ as the reference against which our propositions measure the excess created by synchronized action.

Lemma 4 (Stability). *Write* $\rho_{\text{trend}} = 1/A_s$ *and* $\rho_{\text{amp}} = 1/B_f$. *On the range* $\rho < \rho_{\text{amp}}$ *where the model is well-defined, the growth rate* $G = \kappa(\rho A_s - 1)/(1 - \rho B_f)$ *has the sign of* $\rho - \rho_{\text{trend}}$, *so*

$$G < 0 \iff \rho < \rho^*, \qquad \rho^* = \min\{\rho_{\text{trend}}, \rho_{\text{amp}}\}. \tag{20}$$

The variance of Lemma 1 is finite below ρ_{trend}, *linear in time at* ρ_{trend}, *and unbounded between* ρ_{trend} *and* ρ_{amp},

$$V(t) \to \frac{\sigma_*^2}{2|G|} \;\; (G < 0), \qquad V(t) = \sigma_*^2\, t \;\; (G = 0), \qquad V(t) \sim \frac{\sigma_*^2}{2G}\, e^{2Gt} \;\; (G > 0). \tag{21}$$

In the stable regime, the stationary variance is

$$V_\infty = \frac{\sigma_*^2}{2|G|} = \frac{\sigma^2}{2\kappa\,(1 - \rho A_s)(1 - \rho B_f)} = \frac{\sigma_0^2}{(1 - r)(1 - \beta)}, \tag{22}$$

and it diverges as $\rho \uparrow \rho^*$, *through* $(1 - r) \to 0$ *when* $\rho^* = \rho_{\text{trend}}$ *and through* $(1 - \beta) \to 0$ *when* $\rho^* = \rho_{\text{amp}}$.

The critical homogeneity ρ^* is the point at which the market tips from absorbing a shock to amplifying it without bound, and the two channels reach it in different ways. When the trend binds, $\rho^* = \rho_{\text{trend}}$, the drift G crosses zero and the dislocation begins to grow on its own, the slow cohort's same-direction accumulation overtaking absorption, so the failure is a runaway in the level. When the amplifier binds at $\rho^* = \rho_{\text{amp}}$, the drift stays negative but the fast impact multiplier $1/(1-\beta)$ blows up, carrying the effective shock σ_* to infinity. Consequently, the variance diverges even while the price level mathematically reverts. In this regime, the market failure manifests as an unbounded sensitivity to random news rather than a self-sustaining directional trend. The binding channel is the smaller threshold, the one the market crosses first.

Proposition 1 (Homogeneity and penetration). *On* $0 \le \rho < \rho^*$ *the variance* $V_\infty(\rho)$ *is strictly increasing in* ρ,

$$\frac{\partial V_\infty}{\partial \rho} = \sigma_0^2\, \frac{A_s(1 - \rho B_f) + B_f(1 - \rho A_s)}{\big[(1 - \rho A_s)(1 - \rho B_f)\big]^2} > 0, \tag{23}$$

so the systemic risk $\mathrm{SR}(T)$ *is strictly increasing in* ρ, *and* $\mathrm{SR}(T) \to 1$ *as* $\rho \uparrow \rho^*$.

This proposition answers the paper's second question regarding how the homogeneity of automated forecasts becomes systemic risk. The answer is that homogeneity itself drives this transition. At $\rho = 0$, systemic risk remains strictly at the benchmark. This flat baseline holds regardless of the depth of the book, the speed of the traders, or the size of the automated sector. Systemic risk climbs above the benchmark only as the machines come to read the common signal similarly. It ultimately reaches certainty as ρ approaches ρ^*. This mechanism demonstrates that homogeneity activates the risk, while penetration and speed act as the amplifiers that fuel a runaway surge. Therefore, a regulator who only counts automated participants or fears speed in itself is focusing on factors that are silent at the benchmark. The foundational quantity that activates these amplifiers is the correlation of the forecasts. A diverse market would normally absorb this homogeneity. However, once the automated sector exhibits homogeneity, that same homogeneity transforms into a disproportionate move. At this point, the market does not merely drift toward a danger zone. Instead, it aggressively crosses a catastrophic frontier.

The proposition also carries the first half of the answer to the first question, what a regulator must hold down. Systemic risk climbs as homogeneity increases through the interaction of two distinct channels. Since penetration and speed sit inside A_s and B_f, this climb is steeper at every level of synchronized behavior in a market that the automated sector already dominates. Consequently, two markets with the exact same degree of synchronized machine action can stand at opposite ends of the fragility scale. This divergence occurs if their penetration levels differ. Such a gap remains completely invisible to any metric built on the adoption share. This is because the adoption share remains identical in both markets. Therefore, the regulator's true target is not automation in isolation. Instead, it is the joint product of synchronized action and penetration. The systemic danger is largest precisely where the benefits of automation under normal market conditions are greatest. For this reason, the policy task is not to slow down automation. Regulators must hold synchronized action and penetration jointly below the threshold at which the risk runs away.

Proposition 2 (Speed and the binding channel)**.** *The critical homogeneity* (20) *satisfies*

$$\rho_{\text{amp}} = \frac{1}{B_f} = \frac{1}{\lambda\chi\theta_f v_f}, \qquad \rho_{\text{trend}} = \frac{1}{A_s} = \frac{\kappa}{\lambda\chi\theta_s v_s}, \tag{24}$$

and the amplifier is the binding channel, $\rho^* = \rho_{\text{amp}} < \rho_{\text{trend}}$*, if and only if*

$$\theta_f v_f > \frac{\theta_s v_s}{\kappa}. \tag{25}$$

Both critical levels fall as the relevant frequency rises, so a faster market and a market whose penetration is concentrated in the fast cohort reach instability at a lower homogeneity.

This proposition carries the second question one step further, asking through which channel the homogeneity produces the risk, and it makes the channel a matter of speed. The fast frequency raises the amplifier. The slow frequency raises the trend. As a result, a market can tip in two distinct ways. The first way is a runaway in the level driven by the accumulation of the slow cohort in a single direction. The second way is an explosive surge in the variance driven by the impact multiplier of the fast cohort. The market tips through whichever mechanism reaches its critical homogeneity first. The amplifier binds and the market fails fast under one specific condition. This failure occurs once the fast leveraged mass weighted by its speed outweighs the slow mass measured against absorption. This threshold reveals the structural place of high frequency trading in the mechanism. High frequency trading acts as the sharp end of the risk rather than an independent channel. This specific vulnerability is the exact place the Flash Crash occupies.

The channel is the bridge to the first question, because which channel binds is which instrument the regulator needs. A market can fail fast through the amplifier channel. In this case, a circuit breaker can successfully reach the market, and a limit on the fast leveraged cohort can effectively pre-empt the failure. In contrast, a market can fail slowly through the accumulation of single-direction rebalancing. This slower failure throws the entire burden onto the policy instruments that lower penetration.

Therefore, the diagnosis of the specific channel is identical to the diagnosis of the cure. This distinction is also where the paper parts most sharply from accounts that compress the sector into a single adoption share. Two sectors of equal size and equal synchronized action require entirely different policy instruments. This divergence occurs when one sector concentrates in fast execution while the other concentrates in slow reallocation. For this reason, speed is not a secondary factor. Instead, speed acts as a direct determinant

of the policy.

Proposition 3 (Forced liquidation and the margin)**.** *Above the buffer, for $|y| > m$, the dislocation obeys*

$$dy = (G + \psi\theta_f)\, y\, dt - \psi\theta_f m\, \mathrm{sign}(y)\, dt + \sigma_*\, dW, \tag{26}$$

with growth rate $G_+ = G + \psi\theta_f$ and two unstable fixed points $\pm y^*$, where $y^* = \psi\theta_f m/G_+$, one above m and one below $-m$. A market that is stable below the buffer, $\rho < \rho^*$ so that $G < 0$, nonetheless escapes once a shock carries the dislocation past y^* upward or $-y^*$ downward, and the two points exist when

$$\psi\theta_f > |G| = \frac{\kappa(1 - \rho A_s)}{1 - \rho B_f}. \tag{27}$$

This proposition adds a second route to the systemic risk and, with it, a second target for the first question. The baseline analysis asks whether a market collapses entirely on its own. In contrast, the margin layer asks whether a market that would otherwise remain stable can still be pushed into a systemic collapse. The answer to this second question is yes. This disruption follows a clear sequence. Once a shock pushes the price dislocation past the unstable threshold, the leveraged sector is forced to liquidate its positions directly into the downward price movement. This forced selling rapidly overruns the market absorption that would normally pull the price back. This downward escape is the exact runaway collapse that our metric records. Therefore, a market that is stable in the baseline sense is not fundamentally safe. Instead, it is only conditionally safe. It remains secure only until a shock breaks the margin requirement. The specific shock required to break this margin becomes smaller the closer the market already sits to the homogeneity frontier. For this reason, synchronized behavior and leverage compound as two distinct sources of the exact same tail risk. They do not substitute for one another. Consequently, a calibration model that reads only synchronized behavior will severely understate the true risk of a leveraged market.

For the regulator, this mechanism identifies a leverage channel that remains distinct from the homogeneity channel. It simultaneously pairs two specific instruments to ad-

dress this channel. The first instrument relies on the scale of forced selling. This forced selling scales directly with the leveraged fast mass, which is the exact same mass that drives the amplifier. Therefore, a position limit shrinks this specific mass. By doing so, it successfully lowers both the forced liquidation route and the amplifier together. In contrast, the second instrument is the margin requirement. This instrument raises the loss buffer and shifts the trigger threshold without touching the overall size of the sector. Thus, one instrument works directly on how much mass crosses the threshold, while the other instrument determines where that threshold actually sits. A regulator holding both instruments can effectively reach the leverage channel from two different sides. Consequently, the buffer a market requires should be read directly off the scale of its leveraged sector, rather than being set arbitrarily by convention.

Proposition 4 (Circuit breaker)**.** *Let the breaker freeze the price at the band $b < \bar{c}$ and reopen trading through a call auction. In the auction every cohort clears at a single price with no distinction of frequency, so both feedback channels are absent from price formation and the price forms under the benchmark relation $dy = -\kappa y\,dt + \sigma\,dW$. Measured in the benchmark dispersion that fixes the threshold $\bar{c}$, the auction resets the dislocation from b to*

$$y^{+} = b\sqrt{(1-\rho A_s)(1-\rho B_f)}, \tag{28}$$

where the deflation $D = \sqrt{(1-\rho A_s)(1-\rho B_f)} \in (0,1]$ is the ratio of the benchmark dispersion to the amplified dispersion the continuous market sustains. The reset deepens, $D \to 0$, as homogeneity ρ, the slow channel A_s, and the fast channel B_f rise toward the stability boundary ρ^*, so that the systemic risk after the reset, governed by a first passage that now starts from $y^{+} < b$, falls furthest exactly where the amplification is largest.

This proposition gives the first question its central instrument. The position limit and the margin slow the price movement by lowering its overall growth rate. In contrast, the circuit breaker does not merely slow the movement. Instead, it completely dissolves the underlying mechanism that makes the movement explosive. The structural amplification lives in the staggered superposition of cohorts trading at different frequencies. Within this structure, the fast machines read the exact price increment that the slow machines build. A call auction collapses every trading frequency into a single discrete instant. This

time compression prevents the sequential superposition from forming in the first place. Because of this structural disruption, both feedback channels vanish entirely from the reopening price. Consequently, the diverse liquidity returns to price the book on equal footing. This is the same liquidity that the prior liquidation run or liquidity drain had previously outrun. The returning liquidity now acts precisely as if the automated sector had traded on completely diverse signals.

The deflation $D = \sqrt{(1-\rho A_s)(1-\rho B_f)}$ is exactly the gap between the amplified and the benchmark dispersion, the same $(1-\rho A_s)(1-\rho B_f)$ that sets the stationary variance in Lemma 4. The step that fixes the level of the reset is an assumption, that the auction prices the accumulated dislocation as the benchmark regime would. With both feedback channels absent, the price forms smoothly under that alternative regime. The threshold $\bar{c}$ is itself measured in the benchmark dispersion. Therefore, the consistent reopening level is simply the pre-halt dislocation re-expressed in those specific units. This re-expression represents the deflation. Specifically, it forms the formal content of an auction. This mechanism allows the diverse liquidity to price the book on equal footing. This is the same liquidity that the prior liquidity drain had previously outrun. The breaker therefore returns the most value where the structural amplification is largest. This containment occurs near the stability boundary in a market that is fast, highly penetrated, and homogeneous. Under these extreme conditions, the dislocation D falls to zero as ρ approaches ρ^*. This is the closed-form counterpart of the shift in the binding instrument. Where the rate instruments only shrink the feedback, the auction removes it outright at the reset, and its value grows with the very penetration and speed that make the move fast. A market that has crossed into the fast, high penetration region identifies with Proposition 2. This specific environment requires a regulator to prioritize the breaker. The more a jurisdiction deploys the rate instruments alone, the further behind its regulatory framework falls.

The band b is the empirical circuit breaker, the seven, thirteen, and twenty percent of the United States and the ten, fifteen, and twenty of India. The proposition demonstrates that these thresholds matter most exactly where the automated sector is deep and fast. This environment is precisely where the structural limits are in fact tightest. Therefore, the cross-country pattern of circuit breakers is consistent with the model's prescription.

It is not at odds with it. Specifically, consider the European approach. This framework substitutes venue level interruptions for a single market wide band. The model reads this approach as the variable b dispersed across multiple venues. Consequently, the proposition interprets the European design as a softer constraint on the exact same structural object. The result that closes the section turns this into a statement about how strong the constraint must be.

Proposition 5 (Required intensity and the scaling trap)**.** *Fix a tolerance $\bar{s}$ for systemic risk. Holding* $\mathrm{SR}(T) \leq \bar{s}$ *requires the homogeneity-weighted penetration to stay below the critical level,*

$$\rho \max\{A_s, B_f\} \leq 1 - \eta(\bar{s}), \tag{29}$$

for a margin $\eta(\bar{s}) > 0$ increasing in the strictness of $\bar{s}$. A position limit of tightness ℓ replaces B_f by $(1-\ell)B_f$, so the required ℓ rises as the baseline mass θ_f^0 grows, and a market whose penetration grows fast must tighten the limit over time merely to stay in place.

This proposition specifies the appropriate intensity of regulation. Rather than providing a static baseline number, the model formulates this required strength as a direct function of market penetration. Holding systemic risk below an acceptable tolerance requires the homogeneity weighted penetration to stay safely below the critical level. This safety gap defines the exact distance from the boundary that the regulator must maintain. Crucially, penetration grows naturally as the automated sector expands. For this reason, a regulatory rule that holds perfectly today becomes far too loose tomorrow. The position limit must tighten as the baseline mass rises. This constant adjustment is necessary merely to hold the overall market risk in place. Therefore, the required regulatory intensity is not a level set once by policy makers. Instead, it represents a dynamic path the regulator must continuously walk. A regulatory framework that reads its rules off a static snapshot is already falling behind.

This is the scaling trap, and it is the paper's answer to the comparison of an emerging and a developed market. A developing market often imports the regulatory intensity of a mature market. By doing so, it inherits a constraint calibrated strictly to that mature market's penetration level. This reference point is the wrong scale for an emerging market.

Its residual depth is significantly thinner, and its automated sector is entering at a much faster pace. Consequently, the exact same rule that safely holds a deep and slowly growing market below the critical threshold allows a thin and fast growing market to cross it. The true trap is not that an emerging market regulates too little in an absolute sense. In fact, it may even impose tighter structural limits than a mature market does. Instead, the real danger is that the emerging market regulates according to another market's irrelevant reference point. The proposition identifies the correct reference point. That baseline is the specific penetration of the market itself. Market penetration is highly dynamic and constantly moved by regulation. Therefore, two markets carrying the exact same fraction of automated activity require entirely different rules. This conclusion is the precise result that our introduction promised. It is also the critical insight that a metric built solely on the adoption share cannot deliver.

Taken together the results trace one line. The dislocation is Gaussian, the benchmark fixes the meaning of disproportion, and stability holds below ρ^*. From these propositions, the systemic risk of automated markets emerges with complete clarity. Risk rises strictly with homogeneity and penetration, diverging entirely at a critical threshold. A fast, high-frequency market reaches this instability sooner and breaks specifically through the amplification channel. Once triggered, the margin layer opens a secondary route to a crash via forced liquidation. Against these forces, the circuit breaker uniquely resets the dislocation through a call auction that neutralizes the amplification, proving most valuable precisely where the sector is fastest and most heavily penetrated. Finally, the required regulatory intensity must scale directly with this penetration, trapping any fast-automating market that mistakenly borrows a mature jurisdiction's static rules. The calibration that follows fixes plausible values for the parameters and traces these objects numerically.

5 Calibration

We calibrate the model to map out the comparative statics. During this process, we fix plausible values for the primitives. We then vary the homogeneity ρ that the paper places at the center of the analysis. Finally, we read the systemic risk of each market against its

own specific benchmark.

Because the threshold $\bar{c} = k\sigma_0$ is set directly in the benchmark dispersion σ_0, the absolute scale of the news washes out completely. Consequently, the systemic risk depends on a given market only through two specific dimensionless groups. These groups are defined as $A_s = \lambda\chi\theta_s v_s/\kappa$ and $B_f = \lambda\chi\theta_f v_f$. Specifically, the risk is determined by these two groups together with the homogeneity ρ and the threshold k. The group A_s captures the slow channel while B_f represents the fast channel. A market therefore enters the analysis through these two values and through the regulatory layer where the margin buffer m and the breaker band b are measured in the same dispersion units. The underlying primitives still retain their economic interpretation because a large A_s reflects the combination of thin market depth and an active automated sector. Ultimately the calibration explicitly fixes the pair (A_s, B_f).

The normalization does not assume the markets share a dispersion. Each market is evaluated against its own benchmark. Because the threshold $\bar{c}$ equals the market-specific baseline dispersion σ_0 multiplied by the common factor k, the absolute scale of the news shock σ cancels out within each jurisdiction. Consequently, the common metric across markets is the disproportion factor k rather than the absolute level of volatility. The absorption κ does not cancel out. This resistance remains because it directly sets both the trend group A_s and the speed of mean reversion. Therefore, this parameter is fixed separately for each individual market. The resulting comparison evaluates the exact amplification that homogeneity induces within each market against its own diverse signal baseline. It does not evaluate absolute price movements. Specifically, this precise relative evaluation is the exact comparison that our benchmark was originally built to make.

We set the primitives ordinally rather than empirically. Three core dimensions establish a clear hierarchy for the parameters across the jurisdictions. These dimensions are the depth of a market, the diversity of its liquidity, and the geographical reach of its regulation. Consequently, we fix representative magnitudes that remain strictly consistent with this structural hierarchy. Table 1 collects them. The United States represents a deep market with light algorithmic oversight. This jurisdiction features the most favorable price impact λ and absorption rate κ along with the largest sensitivity χ, fast mass θ_f, and trading frequency v_f. The European Union is deep like the United States but the

lightest on the movement of price. Therefore, its breaker band b is the widest. Conversely, India represents a thin market with strict oversight. Its price impact λ is high and its absorption κ is low. Meanwhile, rigorous approval and frequency limits suppress its sensitivity χ, θ_f, and v_f. Finally, a fourth row places the rules of the United States directly onto the market of India. This configuration carries the thin depth and low absorption of India alongside the sensitivity, fast mass, frequency, and band of the United States. This hybrid scenario represents the precise numerical form of the scaling trap.

Table 1: Illustrative calibration. The slow frequency is fixed at $v_s = 1$. The margin buffer m and the breaker band b are in units of the benchmark dispersion σ_0. The last three columns are implied, with $A_s = \lambda\chi\theta_s v_s/\kappa$, $B_f = \lambda\chi\theta_f v_f$, and $\rho^* = \min\{1/A_s, 1/B_f\}$.

Market	λ	κ	χ	θ_s	θ_f	v_f	m	b	A_s	B_f	ρ^*
United States	1.00	1.00	1.00	0.40	0.46	3.00	2.5	2.2	0.40	1.38	0.72
European Union	1.05	1.00	0.90	0.42	0.42	2.90	2.5	2.8	0.40	1.15	0.87
India	2.80	0.50	0.75	0.36	0.30	2.80	2.8	1.8	1.51	1.76	0.57
India, imported rules	2.80	0.50	1.00	0.36	0.46	3.00	2.5	2.2	2.02	3.86	0.26

Each instrument enters through one primitive, and the values in Table 1 follow from the regulations set out in the introduction. The authorization of algorithms before trade lowers both the sensitivity χ and the frequency v_f. Specifically, consider the United States. Its Market Access Rule places its checks strictly at the broker and reviews no strategy in advance. Consequently, the United States carries the largest χ and v_f. The European Union requires formal notification alongside an annual self assessment. It also explicitly authorizes high frequency firms. Therefore, the European Union carries middle values for these parameters. Conversely, India approves each individual algorithm and assigns it a unique identifier. This strict oversight now reaches retail strategies above a specific order frequency threshold. As a result, India carries the smallest values for these parameters. Furthermore, the position limit lowers the fast mass θ_f while the margin requirement raises the buffer m. India remains the strictest jurisdiction on both regulatory fronts. Therefore, it carries the lowest θ_f and the highest m. The depth of the market and the diversity of its liquidity set the price impact λ and absorption κ. Regulation cannot reach these intrinsic structural features. For this reason, India as a thin emerging market carries the highest λ and the lowest κ. Meanwhile, the two developed markets carry the more favorable values. The masses are anchored to empirical automated trading shares. This

volume exceeds 70% of equity volume in the United States. In India, automation accounts for close to 54% of cash turnover and 70% of derivatives turnover. These empirical baseline proportions order the sector size before position limits split and trim the fast cohort.

The breaker band b represents the per-asset price limit. This corresponds to the dynamic price bands in the United States. It also matches the hard daily bands applied in India. In the European Union, this structure aligns with exchange mandated trading pauses. We focus exclusively on this single-asset constraint because it directly binds the asset dislocation modeled here. Aggregate markets breakers operate as a coarser regulatory layer at higher multiples, such as the 7%, 13%, and 20% halts in the United States or the 10%, 15%, and 20% stages in India. These macro interventions remain outside the scope of our model. India carries the tightest breaker band b. The United States carries a middle value for this parameter. In contrast, the European Union carries the widest band. This jurisdiction leaves the specific limits to each individual venue and sets no rules at the aggregate market level. Homogeneity stays outside this map, since it is the quantity we vary, and the independence of the primitives established in Section 3 is what allows it to move alone.

Two computations produce the results. The closed forms of Section 4 give the stationary dispersion, the critical homogeneity ρ^*, the growth rate G, and the auction deflation D at once. The systemic risk itself represents a first passage of the running maximum. However, the explosive regime along with the forced liquidation and breaker layers leave this trajectory without a closed form solution. Therefore, we estimate this risk using Monte Carlo simulations. Specifically, we simulate the price dislocation in benchmark dispersion units across one trading day. During this process, we record the exact fraction of paths whose largest excursion exceeds $\bar{c}$. The breaker enters the simulation exactly as Proposition 4 states. This mechanism operates as a freeze followed by a call auction. When the dislocation first reaches the band b, trading freezes immediately. It then reopens at the exact theoretical level $\text{sign}(y)\, b\, D$. This specific reset eliminates discrete overshooting distortions. Here, $D = \sqrt{(1-\rho A_s)(1-\rho B_f)}$ represents the deflation factor derived in Proposition 4. After this theoretical adjustment, the feedback loops resume. The model allows only a single halt over the course of one trading day. The band b lies below the threshold $\bar{c}$, so the auction acts before a move counts as systemic. In the stable

regime the simulated terminal variance matches the closed form to several digits, which anchors the estimates.

We now vary ρ over a fine grid from 0.01 to 0.99.

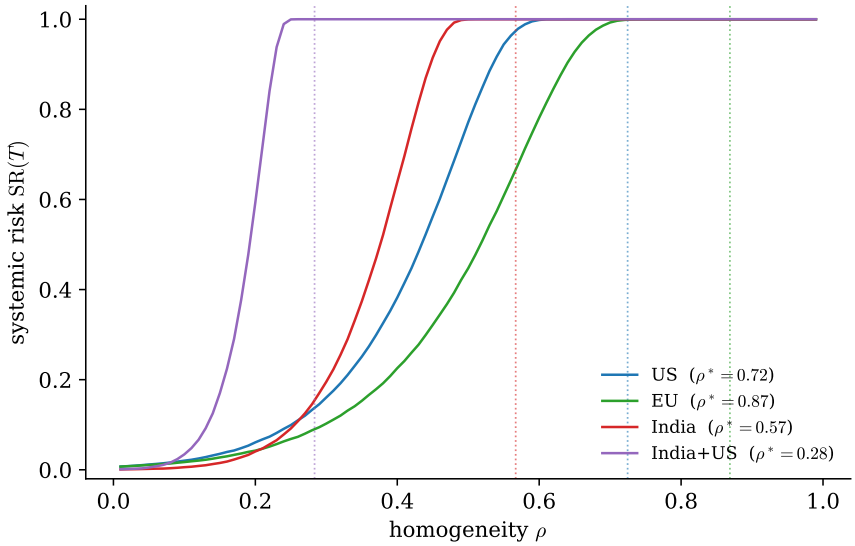


Figure 1: Systemic risk against homogeneity without the breaker. The dotted line marks each market's critical homogeneity ρ^*.

The systemic risk of each market climbs with homogeneity. This risk reaches absolute certainty at that market's own critical level where the amplifier diverges. Crucially, the precise ordering of those critical levels carries two fundamental lessons. Among the developed pair, the European Union holds the higher threshold. Tighter algorithmic rules effectively shrink the amplifier. This suppression creates additional safety margins before systemic collapse. Consequently, this mechanism demonstrates that stricter regulation successfully lowers systemic risk under equivalent market conditions. However, this dynamic shifts across the structural divide. Despite imposing the strictest regulations, India crosses its critical threshold earliest because its thin residual depth dominates the regulatory constraints. Consequently, a hypothetical market combining India's thin depth with the lighter United States regulatory framework would fail first. The maximum level of homogeneity a market can safely sustain is determined by its structure and its regulation together. Crucially, regulation yields the greatest protection when it is precisely calibrated to the specific market structure it governs.

Figure 2 isolates these critical points and sets the two developed markets apart. The

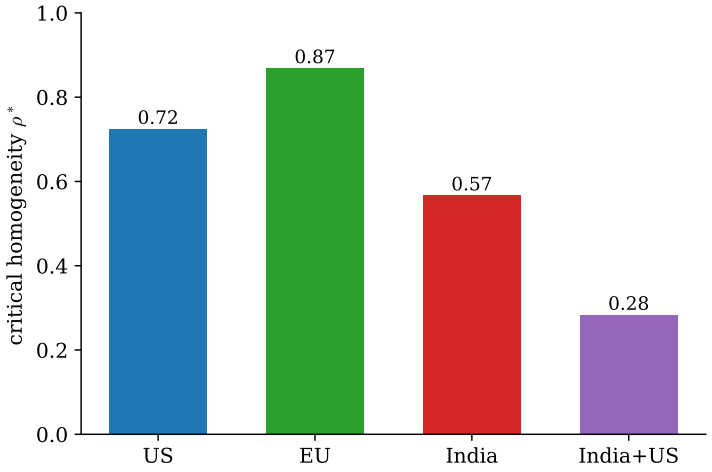


Figure 2: Critical homogeneity $\rho^* = \min\{1/A_s, 1/B_f\}$ across the markets, the safety margins each holds before the amplifier diverges.

European Union holds the largest safety margins while the United States holds less. The substantial gap between these two developed markets is entirely the work of regulation. Specifically, the stricter oversight of algorithms in the European Union holds its amplifier down. In contrast, India holds less safety margins than either developed market. Its thin structural depth completely overrides the strict rules it imposes. Finally, the hybrid market that carries the thin depth of India under the flexible rules of the United States holds the least safety margins of all. Therefore, this precise ordering tracks depth and regulation together. Section 6 reads this result as the fundamental difference between an emerging and a developed market. It also highlights the critical difference a regulator can make within each market type.

The same ordering governs the reach of the breaker, which connects directly to the critical points just shown. The deflation parameter measures the depth of the auction reset. This deflation falls toward zero as homogeneity approaches the critical boundary. Crucially, at any given level of homogeneity, this reset is the deepest for the thinnest market, India. As shown in Figure 2, this is the exact jurisdiction that holds the tightest safety margins. Therefore, the breaker cuts the price dislocation most aggressively where the critical point is nearest. This mechanism ensures that the instrument does its heaviest work in the specific market that needs it most. Ultimately, the total depth of the theoretical reset is directly tied to the inherent safety margins of the market.

Figure 4 turns this reach into the systemic risk itself. The breaker lowers the risk

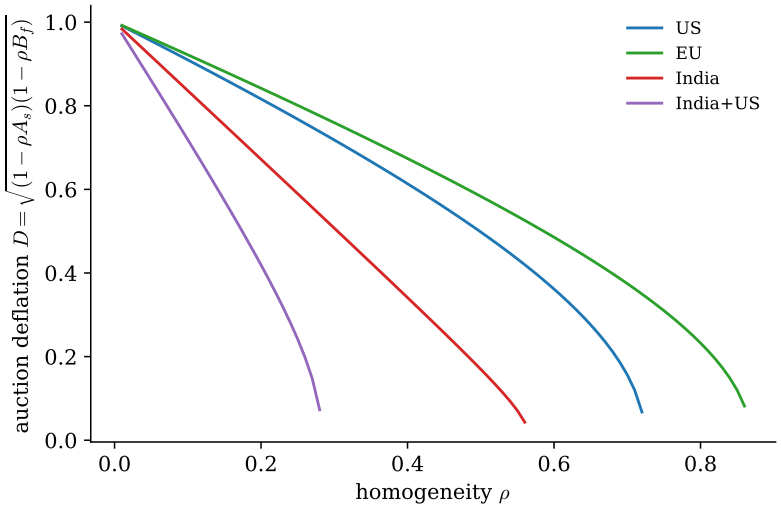


Figure 3: The auction deflation $D = \sqrt{(1-\rho A_s)(1-\rho B_f)}$, the depth of the breaker reset.

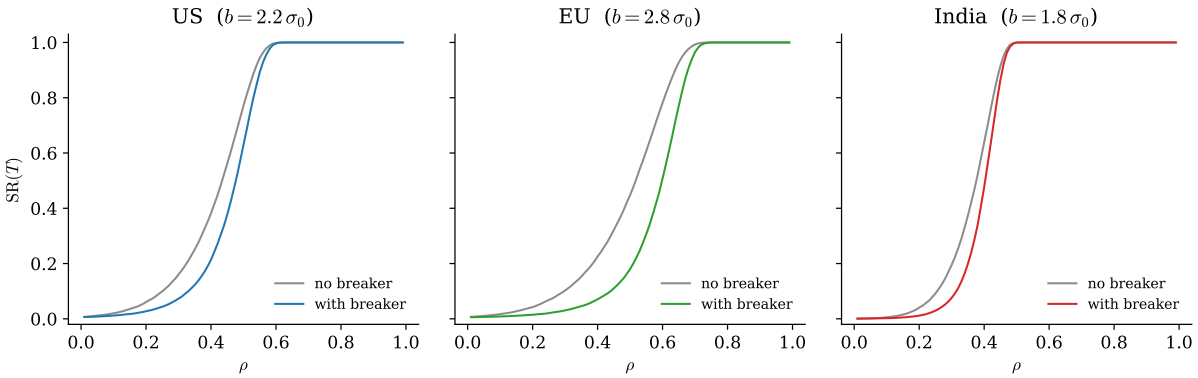


Figure 4: The circuit-breaker auction, with and without the breaker, for the three markets.

across the markets, and over the middle range of homogeneity the reduction is sizeable for each. The breaker does the most work where the price move is fast. Consequently, the widest band of the European Union helps the least. This failure occurs because the European breaker fires the latest and reopens the nearest to the critical threshold. In sharp contrast, the tightest band of India cuts the market move early, proving highly effective for that specific structure. This mechanism completes the fundamental trade-off between the two developed markets. The United States is lighter on algorithm rules. As a result, it holds lower safety margins in Figure 2. To compensate, the United States adopts a narrower breaker band than the European Union. It prioritizes this breaker mechanism, whereas the European Union relies directly on its strict algorithm rules. The auction is worth keeping in every market and worth the most where the amplification is fastest.

While the preceding analysis evaluates regulatory instruments, we now turn to the underlying structural forces including homogeneity, penetration, and speed. In the figures, homogeneity drives systemic risk upward toward the stability boundary. Conversely, penetration and speed determine the location of that boundary itself. To isolate the marginal impact of each structural variable, we vary them independently while holding the remaining parameters fixed at the Indian.

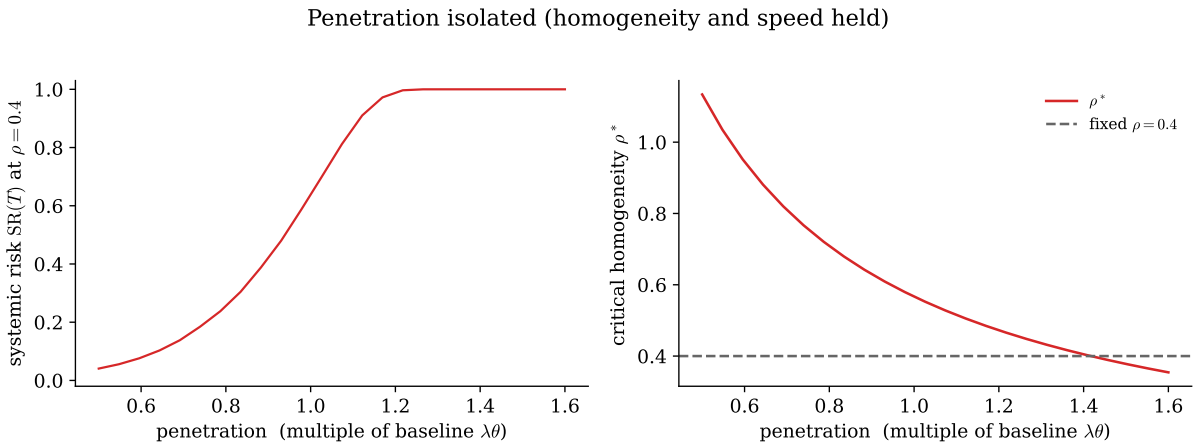


Figure 5: Penetration isolated, with homogeneity and speed held, starting from the calibration of India.

Raising the penetration lifts both parameter groups A_s and B_f in exact proportion. Here, penetration measures the size of the sector relative to the depth that absorbs it. Consequently, the systemic risk climbs while the critical homogeneity falls as the inverse of penetration. Below a specific threshold penetration, the critical homogeneity rises above one. At this stage, no degree of homogeneity can destabilize the market. This stability occurs because homogeneity represents a correlation coefficient which cannot mathematically exceed one. In this sense, penetration and homogeneity operate as direct substitutes in the creation of structural fragility. Ultimately, the systemic risk turns entirely on their mathematical product.

Raising the fast cohort's frequency relative to the slow, the ratio v_f/v_s that measures the speed heterogeneity the model places at its center, lifts only the amplifier group B_f. The trend boundary ρ_{trend} remains constant, whereas the amplifier boundary ρ_{amp} falls inversely with the speed ratio. Past their crossing point, the binding constraint permanently shifts from the trend to the amplifier. A faster market reaches its structural

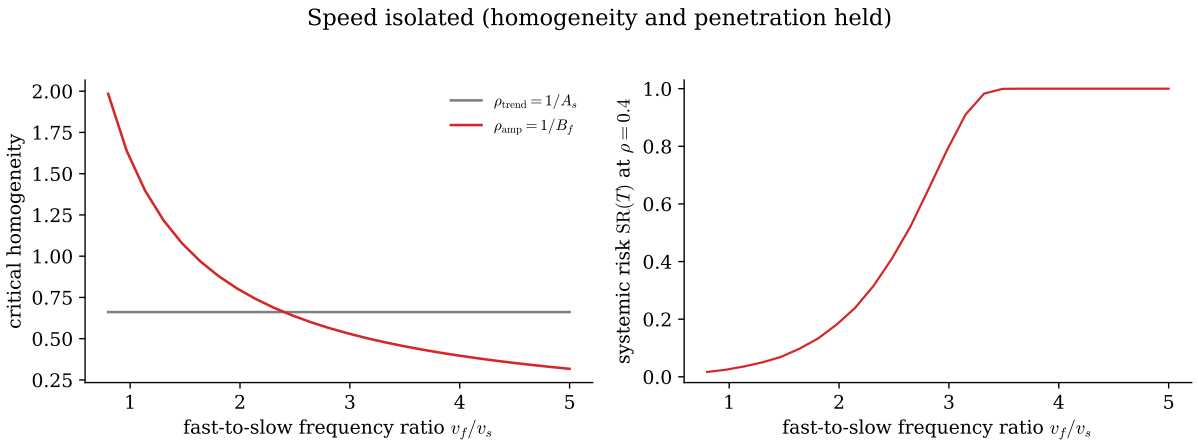


Figure 6: Speed isolated, with homogeneity and penetration held. The horizontal axis is the fast-to-slow frequency ratio v_f/v_s with the slow cohort as the unit.

breakdown through the fast channel that reads the immediate price increment. This outcome establishes that high-frequency trading operates as the sharp end of a single feedback loop rather than as a disconnected risk factor.

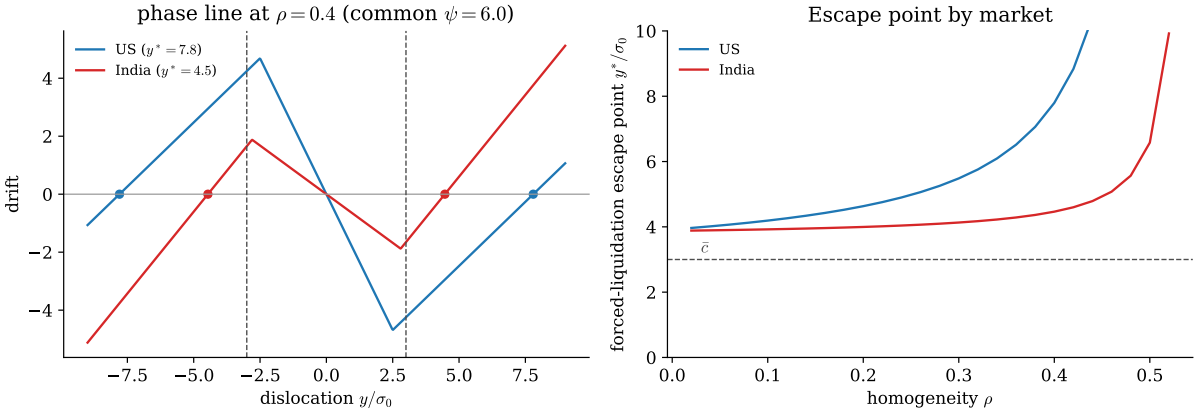


Figure 7: The forced-liquidation second route of Proposition 3, India against the United States, at the calibrated parameters with a common liquidation intensity. The left panel is the drift above and below the margin buffer, the right panel the escape point across homogeneity.

A last figure places the forced-liquidation route of Proposition 3 in the two markets. The left panel of Figure 7 draws the drift above and below the margin buffer. Below the buffer both markets revert toward the open, and the United States reverts more steeply because its deeper book carries stronger structural absorption. Above the buffer, forced selling turns the drift upward and crosses zero at the escape point. Beyond this point, the dislocation runs away into a crash. The escape point is markedly farther for the

United States than for India, so the runaway is the more remote in the deep market. The right panel traces the exact trajectory of the escape point across different levels of homogeneity. Throughout the entire spectrum, the Indian line remains strictly below the American line and sits significantly nearer to the critical threshold. Both lines rise and diverge sharply as homogeneity approaches the boundary where the amplifier takes over. At the calibrated buffers, the escape point stays safely above the threshold in each market. Consequently, this second route represents a remote tail of the distribution rather than its center. However, this risk remains a much nearer tail for the thin market of India.

A last calculation fixes a tolerance for systemic risk and finds the position limit each market needs to meet it at a common homogeneity. The thin market needs a far tighter limit than the deep one, and the market that imports the deep market's limit cannot meet the tolerance at all. This is the scaling trap in numbers, the same limit that holds a deep market below the threshold leaving a thin one above it. The cross-country section that follows reads the three jurisdictions through these objects and draws out the comparison between an emerging and a developed market.

6 Emerging and Developed Markets

The difference between an emerging market and a developed one is usually drawn as a difference in how much of the trading is automated. The model draws it elsewhere. Two markets can carry the exact same share of automation yet stand in very different risk positions. This divergence occurs because the true exposure is determined by penetration rather than the automated share itself. Here, penetration measures the precise size of the automated sector relative to the liquidity depth that absorbs it. Consequently, the penetration of a thin market is exceptionally large, whereas the same automated sector remains small in a deep market. India and the two developed markets differ first in this structural quantity, the price impact λ and the absorption κ of Table 1, and the rest of the comparison follows from it.

The first consequence is a different critical homogeneity. By Proposition 1 the systemic risk rises with homogeneity and diverges at ρ^*, and by Proposition 2 the binding boundary is the amplifier's, $\rho_{\text{amp}} = 1/B_f$, once the fast channel dominates. Because the parameter

group B_f carries the penetration, the thin market reaches that boundary at a lower homogeneity level. Figure 5 traces this exact mechanism directly through the steady fall of ρ^* with rising penetration. Additionally, Figure 1 shows India breaking down at a lower homogeneity level that the United States and the European Union can still safely tolerate. Crucially, Figure 2 synthesizes this entire gap as the explicit safety margins each market holds. These margins remain wide for the deep markets but stay critically narrow for the thin jurisdiction. The threshold a market can carry is set by its structure and its regulation together, and the same degree of homogeneity that a deep market absorbs tips a thin one over.

The second consequence is that an emerging market cannot reach safety by copying a developed market's rule. Proposition 5 makes the required intensity a function of the penetration, so a rule calibrated to a deep market is calibrated to the wrong scale when it is carried to a thin one. The fourth row of Table 1 places the flexible rules of the United States directly onto the thin Indian market. Under this hybrid scenario, India's critical homogeneity ρ^* falls to the absolute lowest level among all four studied configurations. As a direct consequence, its systemic risk in Figure 1 climbs to the highest peak, causing the market to break down first. This dramatic failure delivers a powerful policy lesson. Blind imitation of foreign frameworks completely fails to import the structural safety of a developed market. Instead, it imports a constraint that is too loose for the thin book it now governs. The market that copies is thus left more exposed than the market that wrote its own rule. The instinct to borrow the rule of a mature jurisdiction is exactly the instinct the model warns against, since the borrowed rule raises the systemic risk of the market that adopts it rather than lowering it.

Within the developed pair where depth is roughly equal, the model delivers a sharper lesson showing that stricter regulation lowers systemic risk. The European Union oversees algorithms more tightly than the United States. Figure 2 shows it holds a higher threshold because the tighter rules shrink the sensitivity, the fast mass, and the frequency that feed the amplifier. The gap in their safety margins is therefore the work of regulation alone. The United States applies lighter oversight on algorithms. Consequently, it sits lower on the baseline threshold. To make the difference up, the American framework relies heavily on its other instrument. Specifically, the United States holds a tighter circuit breaker

band than the European Union, whereas the European Union leaves its band the widest. As a result, the two jurisdictions trade places across the algorithm rules and the price limits. Neither regulatory approach dominates the other. The European Union, long the heavier hand in regulation, is on the model's reading buying real stability with it.

The lesson does not carry across the structural divide unchanged. India regulates its algorithms more tightly than either developed market. Yet, it still holds the least safety margins of the three. This vulnerability persists because India's thin structural depth raises the amplifier much faster than its strict rules can lower it. Stricter regulation successfully buys a higher baseline threshold when market depth is held fixed. However, it simply cannot lift a thin market to the safety level of a deep one. Consequently, a thin market that borrows a deep market's lighter rule suffers the greatest collapse of all. This catastrophic decline is clearly demonstrated in the imported rule column. Ultimately, regulation works effectively, but only when it is precisely matched to the specific market structure it governs.

What the thin market can do is rely on the instrument whose value rises with the very forces that threaten it. The circuit breaker lowers the systemic risk in every market, as Figure 4 shows, and by Proposition 4 its auction strips the amplifier at the reopening, so the deflation is the deepest where the amplification is the largest. Figure 3 places India's deflation curve strictly below the developed markets at every level of homogeneity. This positioning demonstrates that the breaker performs its heaviest work in the specific market that holds the narrowest safety margins. While the position limit shrinks the feedback loop and the margin requirements push back the forced liquidation route, the call auction removes the feedback mechanism entirely. Crucially, this auction format eliminates the most risk where the market move is the fastest. This high speed region represents the exact space that a fast automating emerging market naturally comes to occupy.

That region is set by speed. By Proposition 2 a higher fast-to-slow frequency lowers the amplifier boundary, and Figure 6 traces the binding constraint passing to the amplifier as the fast cohort quickens. The precise loop an emerging market must manage is this frequency driven feedback mechanism. Within this loop, the fast cohort continuously reads the price increment that its own order flow creates. The circuit breaker operates as the unique instrument that dissolves this feedback entirely by collapsing all high fre-

quencies into a single call auction. Consequently, the core task for a market like India is to locate the exact point of equilibrium. The breaker must be tight enough to catch the loop before it amplifies out of control. Concurrently, it must remain loose enough not to halt ordinary, healthy trade. This balance requires an engineered optimal band rather than a blindly borrowed framework. The rate instruments India has already tightened cannot by themselves offset a market structure that is inherently thin and fast. Therefore, the breaker is the exact mechanism where the remaining regulatory weight must fall. Ultimately, its precise calibration becomes the fundamental safety margins on which an emerging market's structural stability turns.

The margin layer sharpens the same divide from a second direction. By Proposition 3 a market stable below its buffer can still escape once forced selling above the buffer overruns the reversion, and the point at which it does so is set by the strength of that reversion. Figure 7 places the two markets side by side. The deep order book of the United States reverts strongly. Consequently, its escape point remains highly remote. In sharp contrast, the thin order book of India reverts weakly. Its forced liquidation runaway trajectory therefore sits far nearer to the critical threshold. This structural difference fundamentally alters the nature of systemic risk. The fire sale operates as an immediate, live danger for the emerging market, whereas it remains a distant concern for the developed jurisdiction. Under these conditions, the margin requirement functions as a much more delicate regulatory instrument for the thin market. This sensitivity arises because the exact same safety margins that hold a deep market well clear of the runaway leave a thin market perilously close to it. This proximity reveals a second face of the emerging market's structural difficulty. The thin jurisdiction stands much nearer to the homogeneity boundary and the forced liquidation boundary simultaneously. As a direct result, the regulation tasked with holding both boundaries at once must do significantly more work on a much thinner institutional base.

The comparison thus turns the usual picture around. The emerging market is not behind because it regulates too little, since on the oversight of algorithms and the limits on price it may regulate more than a developed market does. It is exposed because its structure is thin and its sector is entering fast, and because the rule it is most tempted to borrow is calibrated to a depth it does not have. The model names the quantity the

rule should track, the penetration of the market itself, and the instrument that carries the weight when penetration and speed are high, the circuit breaker read as an auction.

7 Conclusion

This paper has asked when the homogeneity of automated trading becomes a source of systemic risk and how a regulator should contain it. The answer rests on one continuous time model of a single trading day. Within this framework, a homogeneous automated sector trades a common signal at heterogeneous speeds. The price impact of these staggered orders subsequently turns their mathematical superposition into a destructive feedback loop that amplifies a small initial shock. Systemic risk is defined here as the exact probability that the market dislocation grows disproportionately far before the close of trade. This probability is evaluated directly against a baseline scenario where the same automated sector trades on entirely diverse signals. Consequently, what this quantitative measure captures is the fundamental fragility of homogeneity. It does not merely reflect the baseline liquidity of automation.

Three structural forces drive that fragility, and the model separates their effects. Homogeneity raises the systemic risk and carries it to certainty at a critical level, where the amplifier diverges. Penetration measures the size of the sector relative to the depth that absorbs it and enters as the multiplier on that homogeneity. The two are therefore complements, meaning that the critical homogeneity falls as the inverse of penetration. Consequently, no level of homogeneity can destabilize a market whose penetration is low enough. Speed enters the system through the specific behavior of the fast cohort. The frequency of this group relative to the slow cohort directly sets the mathematical value of the amplifier. As a direct consequence, a faster market reaches its critical breakdown threshold through this fast channel and at a lower level of homogeneity. The three act through the same coupling, and high-frequency trading is the region where that coupling is largest rather than a force apart.

The regulator holds three instruments, and each acts at a separate point. A position limit lowers the fast leveraged mass and shrinks the amplifier. A margin raises the buffer against forced liquidation, the second and bidirectional route to a crash that opens once

a move overruns the reversion. A circuit breaker freezes the price and reopens it through a call auction. Because the auction clears every frequency at once, it removes the amplifier outright and resets the dislocation by a deflation that is the deepest where the amplification is the largest. The regulatory intensity required to keep systemic risk below an acceptable tolerance level rises alongside its penetration. Consequently, the binding policy instrument shifts decisively toward the circuit breaker as rising penetration and speed accelerate market movements. A regulatory regime that reads its rules off a static snapshot inevitably falls behind an automated sector that keeps growing.

This paper provides three primary contributions to the study of modern market risk. First, it treats market penetration as a dynamic variable shifted by policy rather than a static volume share. Second, it models regulation through the specific operational instruments that regulators control to measure their direct efficacy against systemic failures. Third, rather than assuming a feedback channel a priori, the paper derives macro amplification endogenously from the frequency structure of correlated trading. This mathematical approach establishes the systemic outcome as an explicit property of execution mechanics.

The analytical framework is intentionally tractable. Consequently, its clean structural choices naturally suggest several logical directions for future extension. The calibration serves an illustrative purpose rather than an empirical estimation. Its parameters are ordered by the qualitative structure and regulation of each jurisdiction rather than fitted to historical data, making the cross-country analysis qualitative. Systemic risk is evaluated as a first-passage probability of the running maximum, which the analytical section approximates through the terminal tail and the simulation confirms.

The comparative analysis of emerging and developed jurisdictions follows directly from these principles. Two markets displaying an identical share of automated volume can occupy entirely different risk positions because the thin depth of an emerging market increases its structural penetration and lowers its critical homogeneity threshold. A jurisdiction that imports rules from a mature market inherits a policy constraint calibrated for a depth it does not possess, causing the borrowed standard to increase rather than decrease systemic risk. Within the developed pair, the model adds that stricter regulation lowers systemic risk. The European Union applies tighter oversight of algorithms

to hold a higher threshold. Conversely, the United States applies a lighter approach to algorithms but relies instead on a tighter breaker, allowing the two jurisdictions to trade places across the instruments. This logic does not cross the structural divide because the emerging market enforces tighter algorithmic restrictions than either developed market yet retains the narrowest safety margin. India's thin liquidity base accelerates the amplifier much faster than its strict rules can suppress it. Consequently, this structural deficit pushes the trading venue substantially closer to both the homogeneity boundary and the forced liquidation boundary simultaneously. Regulation remains effective only when it is matched precisely to the underlying market structure.

The primary policy insight is that regulatory intensity in an automating market cannot be treated as a static choice. It must operate as a dynamic path that continuously tracks structural penetration. A rapidly automating emerging market cannot safely borrow the regulatory standards of a mature jurisdiction. As execution speeds accelerate, the primary policy tool for maintaining stability shifts decisively to the circuit breaker operated as a clearing auction. This specific format functions as the only instrument capable of eliminating the automated amplifier entirely. Other regulatory tools merely slow its progress. The requirement for proactive, jurisdiction-specific calibration is dictated by the absolute scale of the underlying market book rather than the adoption rate of its technology.

While this framework offers clear insights for cross-jurisdictional calibration, a specific limitation of our regulatory design is the assumption of a single circuit breaker per trading day, whereas real-world market-wide circuit breakers or volatility-interruption mechanisms can trigger repeatedly. We adopt this tractable bound to isolate the pure economic effect of a single call auction in eliminating the automated feedback loop. By focusing on the one time removal of the amplifier at the reset, the paper establishes a clean lower bound for regulatory efficacy. This targeted approach deliberately avoids confounding the analysis with multi period strategic execution. Furthermore, it completely eliminates the distorting effects of end of day order congestion.

The forced liquidation runway occurs exclusively beyond the systemic threshold under realistic margin buffers. It operates strictly as a tail risk mechanism rather than a central tendency. Under this condition, the linear baseline omits the severe volume nonlinearities

that would eventually tilt the runaway path downward. Each of these components can be integrated into the current analytical framework without altering its core insights.

Consequently, its clean structural choices naturally suggest several logical directions for future extension. Leaving the derivation of a closed-form first-passage bound for future research, a comprehensive extension could incorporate multiple assets, recurrent halts, and the cross-asset contagion mechanics that market-wide circuit breakers are designed to contain. Extending this framework to a fully dynamic setting with recurrent halts and endogenous machine-learning re-calibration remains a highly valuable path for future research.

Reference

[1] Abis, Simona. "Man vs. Machine: Quantitative and Discretionary Equity Management." SSRN Working Paper No. 3717371 (2020).

[2] Back, Kerry. "Insider trading in continuous time." The Review of Financial Studies 5, no. 3 (1992): 387-409.

[3] Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi. "Hedge fund stock trading in the financial crisis of 2007–2009." The Review of Financial Studies 25, no. 1 (2012): 1-54.

[4] Biais, Bruno, Thierry Foucault, and Sophie Moinas. "Equilibrium fast trading." Journal of Financial economics 116, no. 2 (2015): 292-313.

[5] Boehmer, Ekkehart, Kingsley Fong, and Juan Julie Wu. "Algorithmic trading and market quality: International evidence." Journal of Financial and Quantitative Analysis 56, no. 8 (2021): 2659-2688.

[6] Bogousslavsky, Vincent. "Infrequent rebalancing, return autocorrelation, and seasonality." The Journal of Finance 71, no. 6 (2016): 2967-3006.

[7] Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. "High-frequency trading and price discovery." The Review of Financial Studies 27, no. 8 (2014): 2267-2306.

[8] Brunnermeier, Markus K., and Lasse Heje Pedersen. "Market liquidity and funding liquidity." The review of financial studies 22, no. 6 (2009): 2201-2238.

[9] Budish, Eric, Peter Cramton, and John Shim. "The high-frequency trading arms race: Frequent batch auctions as a market design response." The Quarterly Journal of Economics 130, no. 4 (2015): 1547-1621.

[10] Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega. "Rise of the machines: Algorithmic trading in the foreign exchange market." The Journal of Finance 69, no. 5 (2014): 2045-2084.

[11] Coval, Joshua, and Erik Stafford. "Asset fire sales (and purchases) in equity markets." Journal of Financial Economics 86, no. 2 (2007): 479-512.

[12] Danielsson, Jon, Robert Macrae, and Andreas Uthemann. "Artificial intelligence and systemic risk." Journal of Banking & Finance 140 (2022): 106290.

[13] Diamond, Douglas W., and Robert E. Verrecchia. "Information aggregation in a noisy rational expectations economy." Journal of financial economics 9, no. 3 (1981): 221-235.

[14] Dou, Winston Wei, Itay Goldstein, and Yan Ji. Ai-powered trading, algorithmic collusion, and price efficiency. No. w34054. National Bureau of Economic Research, 2025.

[15] European Commission. "Commission Delegated Regulation (EU) 2017/589 of 19 July 2016 Supplementing Directive 2014/65/EU with Regard to Regulatory Technical Standards Specifying the Organisational Requirements of Investment Firms Engaged in Algorithmic Trading (RTS 6)." Official Journal of the European Union L 87 (2017): 417-448.

[16] European Union. "Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on Markets in Financial Instruments (MiFID II)." Official Journal of the European Union L 173 (2014): 349-496.

[17] Farboodi, Maryam, and Laura Veldkamp. "Long-run growth of financial data technology." American Economic Review 110, no. 8 (2020): 2485-2523.

[18] Foster, F. Douglas, and S. Viswanathan. "Strategic trading when agents forecast the forecasts of others." The Journal of Finance 51, no. 4 (1996): 1437-1478.

[19] Gao, Lei, Yufeng Han, Sophia Zhengzi Li, and Guofu Zhou. "Market intraday momentum." Journal of Financial Economics 129, no. 2 (2018): 394-414.

[20] Glosten, Lawrence R., and Paul R. Milgrom. "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." Journal of financial economics 14, no. 1 (1985): 71-100.

[21] Greenwald, Bruce C., and Jeremy C. Stein. "Transactional risk, market crashes, and the role of circuit breakers." Journal of Business (1991): 443-462.

[22] Greenwood, Robin, and David Thesmar. "Stock price fragility." Journal of Financial Economics 102, no. 3 (2011): 471-490.

[23] Gromb, Denis, and Dimitri Vayanos. "Equilibrium and welfare in markets with financially constrained arbitrageurs." Journal of financial Economics 66, no. 2-3 (2002): 361-407.

[24] Grossman, Sanford J., and Joseph E. Stiglitz. "On the impossibility of informationally efficient markets." The American economic review 70, no. 3 (1980): 393-408.

[25] Hellwig, Martin F. "On the aggregation of information in competitive markets." Journal of economic theory 22, no. 3 (1980): 477-498.

[26] Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. "Does algorithmic trading improve liquidity?." The Journal of finance 66, no. 1 (2011): 1-33.

[27] Holden, Craig W., and Avanidhar Subrahmanyam. "Long-lived private information and imperfect competition." The Journal of Finance 47, no. 1 (1992): 247-270.

[28] Jones, Charles M. "What do we know about high-frequency trading?." Columbia Business School Research Paper 13-11 (2013).

[29] Khandani, Amir E., and Andrew W. Lo. "What happened to the quants in August 2007? Evidence from factors and transactions data." Journal of Financial Markets 14, no. 1 (2011): 1-46.

[30] Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. "The flash crash: High-frequency trading in an electronic market." The Journal of Finance 72, no. 3 (2017): 967-998.

[31] Kyle, Albert S. "Continuous auctions and insider trading." Econometrica: Journal of the Econometric Society 53, no. 6 (1985): 1315-1335.

[32] McLean, R. David, and Jeffrey Pontiff. "Does academic research destroy stock return predictability?." The Journal of Finance 71, no. 1 (2016): 5-32.

[33] Meng, Shuchen, and Xupeng Chen. "Artificial intelligence and systemic risk: A unified model of performative prediction, algorithmic herding, and cognitive dependency in financial markets." arXiv preprint arXiv:2604.03272 (2026).

[34] Menkveld, Albert J. "High frequency trading and the new market makers." Journal of financial Markets 16, no. 4 (2013): 712-740.

[35] National Stock Exchange of India (NSE). "Market Pulse." Vol. 7, no. 3, March 2025. Mumbai: National Stock Exchange of India (2025).

[36] Securities and Exchange Board of India (SEBI). "Broad Guidelines on Algorithmic Trading." Circular No. CIR/MRD/DP/09/2012. Mumbai: Securities and Exchange Board of India (2012).

[37] Securities and Exchange Board of India (SEBI). "Framework for Algorithmic Trading by Retail Individual Investors." Circular No. SEBI/HO/MRD/MRD-PoD-3/P/CIR/2025. Mumbai: Securities and Exchange Board of India (2025).

[38] Shleifer, Andrei, and Robert W. Vishny. "The limits of arbitrage." The Journal of finance 52, no. 1 (1997): 35-55.

[39] Stein, Jeremy C. "Presidential address: Sophisticated investors and market efficiency." The Journal of Finance 64, no. 4 (2009): 1517-1548.

[40] Subrahmanyam, Avanidhar. "Circuit breakers and market volatility: A theoretical perspective." The Journal of Finance 49, no. 1 (1994): 237-254.

[41] U.S. Securities and Exchange Commission (SEC). "Risk Management Controls for Brokers or Dealers with Market Access." Exchange Act Release No. 34-63241, File No. S7-03-10. Federal Register 75, no. 220 (2010): 69792-69838.

[42] U.S. Securities and Exchange Commission (SEC). "Staff Report on Algorithmic Trading in U.S. Capital Markets." Washington, D.C.: U.S. Securities and Exchange Commission (2020).

A Proofs

Proof of Lemma 1. Apply the integrating factor e^{-Gt}. By Itô's product rule, since e^{-Gt} has finite variation,

$$d\big(e^{-Gt}y(t)\big) = e^{-Gt}\,dy(t) - Ge^{-Gt}y(t)\,dt = e^{-Gt}\big(Gy\,dt + \sigma_*\,dW\big) - Ge^{-Gt}y\,dt = \sigma_*\,e^{-Gt}\,dW(t).$$

Integrate from 0 to t and use $y(0) = 0$,

$$e^{-Gt}y(t) = \sigma_* \int_0^t e^{-Gs}\,dW(s),$$

which gives (17) after multiplying by e^{Gt}. The integrand is deterministic, so $y(t)$ is Gaussian with mean zero, and its variance is the second moment

$$V(t) = \mathbb{E}\big[y(t)^2\big] = \sigma_*^2\,\mathbb{E}\left[\Big(\int_0^t e^{G(t-s)}\,dW(s)\Big)^2\right].$$

The Itô isometry takes this expectation to the deterministic integral,

$$V(t) = \sigma_*^2 \int_0^t e^{2G(t-s)}\,ds = \sigma_*^2 \int_0^t e^{2Gu}\,du = \sigma_*^2\,\frac{e^{2Gt}-1}{2G},$$

where $u = t - s$. The expression extends to $G = 0$ by the limit $V(t) = \sigma_*^2\,t$. □

Proof of Lemma 2. By Lemma 1, $y(T)$ is Gaussian with mean zero and variance $V(T)$. For a centered normal variable with standard deviation $\sqrt{V(T)}$,

$$\Pr\big(|y(T)| > \bar{c}\big) = 2\,\Pr\big(y(T) > \bar{c}\big) = 2\left(1 - \Phi\big(\bar{c}/\sqrt{V(T)}\big)\right) = 2\,\Phi\big(-\bar{c}/\sqrt{V(T)}\big).$$

Since Φ is increasing and $\bar{c} > 0$, the argument $-\bar{c}/\sqrt{V(T)}$ increases in $V(T)$, so the probability is strictly increasing in $V(T)$. □

Proof of Lemma 3. By (16), $\rho = 0$ gives $\alpha = \beta = 0$, so $r = 0$, and by (15) $G = -\kappa$ and $\sigma_* = \sigma$. Equation (14) becomes $dy = -\kappa y\,dt + \sigma\,dW$. Lemma 1 gives $V(t) = \sigma^2(1 - e^{-2\kappa t})/(2\kappa) = \sigma_0^2(1 - e^{-2\kappa t})$. Since α and β are the only places θ, v_s, and v_f enter, and both vanish, the law does not depend on them. For a day long enough that $e^{-2\kappa T}$ is negligible, $V(T) \to \sigma_0^2$, and Lemma 2 with $\bar{c} = k\sigma_0$ gives $2\Phi(-\bar{c}/\sigma_0) = 2\Phi(-k)$. □

Proof of Lemma 4. Where the model is well-defined, $1 - \beta = 1 - \rho B_f > 0$, and the sign of G is the sign of $r - 1 = \rho A_s - 1 = A_s(\rho - \rho_{\text{trend}})$ with $\rho_{\text{trend}} = 1/A_s$. Hence, $G < 0$ for $\rho < \rho_{\text{trend}}$, $G = 0$ at $\rho = \rho_{\text{trend}}$, and $G > 0$ above it, while the consistency condition caps ρ below $\rho_{\text{amp}} = 1/B_f$. The mean-reverting region is $\rho < \min\{\rho_{\text{trend}}, \rho_{\text{amp}}\}$, which is (20), and the explosive band $\rho_{\text{trend}} < \rho < \rho_{\text{amp}}$ is nonempty only when $\rho_{\text{trend}} < \rho_{\text{amp}}$. The three forms in (21) follow from (18), since $e^{2Gt} \to 0$ for $G < 0$ gives $V \to \sigma_*^2/(2|G|)$, the limit at

$G = 0$ is $\sigma_*^2 t$, and e^{2Gt} dominates for $G > 0$. Substituting $\sigma_*^2 = \sigma^2/(1-\rho B_f)^2$ and $|G| = \kappa(1-\rho A_s)/(1-\rho B_f)$ into $\sigma_*^2/(2|G|)$ gives $\sigma^2/[2\kappa(1-\rho A_s)(1-\rho B_f)] = \sigma_0^2/[(1-r)(1-\beta)]$, which is (22). As $\rho \uparrow \rho_{\text{trend}}$, $r = \rho A_s \to 1$ so $(1-r) \to 0$, and as $\rho \uparrow \rho_{\text{amp}}$, $\beta = \rho B_f \to 1$ so $(1-\beta) \to 0$, either limit sending $V_\infty \to \infty$. □

Proof of Proposition 1. Differentiate (22). With $f(\rho) = (1-\rho A_s)(1-\rho B_f)$, $V_\infty = \sigma_0^2/f(\rho)$ and $\partial V_\infty/\partial\rho = -\sigma_0^2 f'(\rho)/f(\rho)^2$. Here $f'(\rho) = -A_s(1-\rho B_f) - B_f(1-\rho A_s)$, which is negative on $\rho < \rho^*$, so $\partial V_\infty/\partial\rho > 0$ and equals (23). By Lemma 2 the terminal tail is increasing in the variance, hence in ρ, and $\text{SR}(T)$ dominates the terminal tail and rises with it. As $\rho \uparrow \rho^*$ one factor of $f(\rho)$ tends to zero, so $V_\infty \to \infty$, the terminal tail tends to $2\Phi(0^-) = 1$, and $\text{SR}(T) \to 1$. □

Proof of Proposition 2. The two expressions are the reciprocals of B_f and A_s from (16). By (20), $\rho^* = \rho_{\text{amp}}$ iff $1/B_f < 1/A_s$ iff $A_s < B_f$. Substituting, $\lambda\chi\theta_s v_s/\kappa < \lambda\chi\theta_f v_f$, which reduces to (25). Since B_f increases in v_f and in θ_f, and A_s increases in v_s and in θ_s, the reciprocals ρ_{amp} and ρ_{trend} decrease in the respective frequencies and masses. □

Proof of Proposition 3. The forced-selling term $\psi\theta_f \operatorname{sign}(y)(|y|-m)_+$ acts only above the buffer and pushes in the direction of the move, so we take the two sides separately. For $y > m$ it equals $\psi\theta_f(y-m)$ and for $y < -m$ it equals $\psi\theta_f(y+m)$, and adding each to (14) gives

$$dy = G_+\, y\, dt - \psi\theta_f m\, dt + \sigma_*\, dW \quad (y > m), \qquad dy = G_+\, y\, dt + \psi\theta_f m\, dt + \sigma_*\, dW \quad (y < -m),$$

with homogeneous growth $G_+ = G + \psi\theta_f$. The sign of the constant flips between the two sides, which is the $\operatorname{sign}(y)$ of (26), and the drift vanishes at $y = \psi\theta_f m/G_+ \equiv y^*$ above the buffer and at $-y^*$ below it.

Around either fixed point the deviation satisfies $dz = G_+ z\, dt + \sigma_*\, dW$, which is explosive when $G_+ > 0$. Since $G = -|G|$ below the stability boundary, this is the condition $\psi\theta_f > |G|$ of (27), with $|G|$ substituted from (15), and the strict inequality keeps $G_+ > 0$ so that y^* is finite. If instead $\psi\theta_f \le |G|$, then $G_+ \le 0$, no fixed point lies beyond the buffer, and the drift there points back toward it, so the dislocation does not escape, the boundary $\psi\theta_f = |G|$ being the marginal case in which $G_+ = 0$ and the restoring drift is the constant $-\psi\theta_f m \operatorname{sign}(y)$.

In the escape case $y^* = \psi\theta_f m/(\psi\theta_f - |G|) > m$ and $-y^* < -m$. At $y = m$ the drift is $-|G|m < 0$ and at $y = -m$ it is $+|G|m > 0$, so the band $-y^* < y < y^*$ reverts toward the open, and escape requires a shock to carry the dislocation above y^* or below $-y^*$, where the drift drives it away. The downward escape past $-y^*$ is the crash the systemic measure records. □

Proof of Proposition 4. In continuous trading the reduced dynamics (14) have growth $G = \kappa(\rho A_s - 1)/(1 - \rho B_f)$ and effective shock $\sigma_* = \sigma/(1 - \rho B_f)$. Below the stability boundary the dislocation has the stationary dispersion

$$V_\infty = \frac{\sigma_*^2}{2|G|} = \frac{\sigma^2}{2\kappa\,(1 - \rho A_s)(1 - \rho B_f)} = \frac{\sigma_0^2}{(1 - \rho A_s)(1 - \rho B_f)},$$

against the benchmark $\sigma_0^2 = \sigma^2/(2\kappa)$, so the amplified standard deviation exceeds the benchmark by the factor $[(1 - \rho A_s)(1 - \rho B_f)]^{-1/2}$.

The call auction sets one clearing price from the orders aggregated during the freeze, so the fast cohort has no contemporaneous increment to read and the slow cohort no staggered arrival. Both feedback terms drop from price formation, and the price forms in the benchmark regime, so a dislocation that measures b in the amplified dispersion measures

$$y^+ = b\sqrt{(1 - \rho A_s)(1 - \rho B_f)}$$

in the benchmark dispersion, which is the reset of (28).

The factor lies in $(0, 1]$ for $\rho < \rho^*$, where both $1 - \rho A_s$ and $1 - \rho B_f$ are positive, and falls to zero as $\rho \to \rho^*$. After reopening the increment exists again and both channels return, so the auction is the one-time removal of the amplification at the reset. □

Proof of Proposition 5. By Proposition 1, $\mathrm{SR}(T)$ increases in ρ from the benchmark to one across $[0, \rho^*)$, so a tolerance $\bar{s}$ above the benchmark defines a unique $\rho_{\bar{s}} < \rho^*$ with $\mathrm{SR}(T) = \bar{s}$, and $\mathrm{SR}(T) \le \bar{s}$ iff $\rho \le \rho_{\bar{s}}$. Writing $\rho_{\bar{s}} = (1 - \eta)/\max\{A_s, B_f\}$ with $\eta = 1 - \rho_{\bar{s}} \max\{A_s, B_f\} > 0$ gives (29), and a stricter $\bar{s}$ lowers $\rho_{\bar{s}}$ and raises η. The position limit sets $\theta_f = (1 - \ell)\theta_f^0$, so $B_f = \lambda\chi(1 - \ell)\theta_f^0 v_f$. To hold (29) as θ_f^0 rises, $(1 - \ell)\theta_f^0$ must stay bounded, so ℓ must rise toward one as θ_f^0 grows. □