



Research article

Optimal automation under overdispersed discrete risk: thresholds and hysteresis in a Negative Binomial model

Jinho Cha^{1,*}, Sahng-Min Han¹, Long Pham², Joseph Mollick² and Justin Yu³

¹ Department of Computer Science, Gwinnett Technical College, Georgia, USA

² Department of Decision Sciences and Economics, College of Business, Texas A&M University–Corpus Christi, Texas, USA

³ Scheller College of Business, Georgia Institute of Technology, Georgia, USA

* **Correspondence:** Email: jcha@gwinnettech.edu.

Abstract: Decision-making under clustered uncertainty requires models that accommodate discrete overdispersed risk and endogenous control mechanisms. Standard approaches often impose equi-dispersion or Gaussian assumptions and treat control as exogenous or binary, limiting structural realism. In this paper, we developed a stochastic optimization framework in which risk follows a dynamic Negative Binomial process with time-varying success probability, generating persistent overdispersion and temporal dependence. A continuous control variable mitigated non-tail and tail exposure under convex implementation and adjustment costs. Analytical results showed that overdispersion reshapes the optimal policy, generating interior solutions, dispersion-dependent thresholds, and regime-dependent hysteresis effects that do not arise under equi-dispersed specifications. Simulation evidence confirms that overdispersion amplifies tail risk and induces nonlinear responses. An empirical analysis using daily S&P 500 index data documents significant overdispersion and persistent clustering. The results provide a structural basis for integrating discrete risk modeling with dynamic operational decisions under uncertainty.

Keywords: Negative Binomial model; overdispersion; clustered risk; automation intensity; threshold policy; hysteresis

Mathematics Subject Classification: 90C15, 90C39, 60E05, 62P20

1. Introduction

Modern financial and operational systems are increasingly exposed to clustered risk, in which shocks arrive in bursts and volatility regimes persist over time [1–3]. Empirical evidence from equity and currency markets shows that shocks are neither independent nor evenly distributed over time. Realized

events exhibit overdispersion, temporal dependence, and pronounced clustering [4, 5]. These features imply that risk variability is state-dependent and can remain elevated for extended periods.

In such environments, conventional modeling assumptions are structurally inadequate. The classical Poisson framework imposes an equi-dispersion restriction that mechanically ties the variance of event counts to their mean, an assumption long recognized as empirically fragile in count-data applications [6, 7]. By construction, this restriction rules out persistent overdispersion and cannot accommodate regimes in which risk variability systematically exceeds its average level. Gaussian approximations, while analytically convenient, abstract from discreteness and are poorly suited to represent clustered jump arrivals. Financial econometrics has repeatedly documented volatility clustering, conditional heteroskedasticity, and long-memory dynamics in asset returns [8, 9], indicating that dispersion dynamics are structural features rather than statistical anomalies.

When dispersion is underestimated, policy prescriptions become systematically distorted. Optimal execution strategies, clearing rules, and risk-management mechanisms derived under homoskedastic Gaussian or equi-dispersed Poisson assumptions may understate tail exposure and systemic fragility. In particular, ignoring persistent overdispersion can bias investment decisions in monitoring, automation, and protective infrastructure [10–13].

Automation and smart execution mechanisms have become central to modern financial infrastructures [14–18]. Algorithmic trading systems, automated clearing platforms, and rule-based smart contracts enable rapid state-contingent execution. Furthermore, automation entails implementation costs, adjustment frictions, and potential rigidity [14, 18]. Greater automation intensity may mitigate losses and reduce reaction delays, but it also reduces flexibility and introduces switching costs. This trade-off gives rise to a structural optimization problem: How should automation intensity be chosen when risk arrives discretely and clusters persistently over time?

To address this question, the framework developed in this paper represents discrete loss or demand events using a Negative Binomial (NB) count process, which explicitly accommodates overdispersion and clustering [19, 20]. Automation intensity is modeled as a continuous control variable that mitigates loss exposure while incurring convex implementation and adjustment costs. The analysis begins with a static optimization problem and extends to a dynamic regime-switching environment in which clustered risk evolves persistently across states. Under this formulation, dispersion is not a second-order statistical correction but a first-order structural determinant of the optimal automation architecture.

Contributions. The major contributions are as follows:

- (i) A structural optimization framework is established in which clustered discrete risk reshapes the geometry of optimal automation policies, generating threshold effects absent under equi-dispersed models.
- (ii) Overdispersion is shown to affect not only expected losses but also the structure of policy responses, producing interior solutions and regime-dependent automation thresholds.
- (iii) In a dynamic regime-switching environment with adjustment frictions, endogenous hysteresis arises, implying that crisis-induced automation need not fully revert even when underlying conditions normalize.
- (iv) Empirical evidence based on S&P 500 jump measures supports these mechanisms by showing that dispersion varies across regimes and that the response to dispersion becomes regime-dependent, including reversal behavior in high-dispersion states.

Empirical motivation. Daily extreme-return (jump) counts for the S&P 500 provide a natural setting in which clustered dispersion is empirically salient [21, 22]. Volatility spike counts, such as threshold exceedances of the Chicago Board Options Exchange Volatility Index (CBOE VIX), and insurance claim frequencies provide additional examples that typically exhibit overdispersion [23, 24]. These series provide natural calibration environments in which clustered dispersion is economically meaningful.

The remainder of the paper is organized as follows: In section 3, we present the static formulation, including the dispersion-based threshold structure. In Section 4, we develop the dynamic extension with regime switching and adjustment frictions. In Section 5, we provide synthetic structural experiments. In Section 6, we present the empirical analysis and implementation. In Section 7, we present the structural implications, integrate the theoretical and empirical findings, and conclude.

2. Literature review

2.1. Clustered risk and overdispersion

Empirical literature documents that risk in financial markets is not well described by independent and homoskedastic shocks. Instead, extreme events tend to arrive in clusters, volatility exhibits persistence, and tail realizations concentrate in episodic bursts. Volatility clustering is a foundational stylized fact, formalized through conditional heteroskedasticity models such as autoregressive conditional heteroskedasticity (ARCH) and generalized autoregressive conditional heteroskedasticity (GARCH) [25, 26], while high-frequency evidence shows that jump components occur in concentrated episodes rather than as isolated arrivals [21, 22]. Further, tail risk premia vary over time, indicating regime-dependent compensation for rare events [23]. These findings imply that risk intensity is persistent, state-dependent, and frequently exhibits dispersion beyond that implied by equi-dispersed count models.

Within discrete-event modeling, overdispersion has long been recognized as a central empirical issue. Early econometric contributions demonstrated that Poisson specifications are often rejected because the variance of observed counts exceeds the mean, motivating the use of NB models as Poisson–Gamma mixtures that accommodate excess variance [6, 7]. In actuarial science, NB and related mixed-Poisson models are standard tools for modeling clustered claim arrivals and heterogeneous risk exposure [24]. The defining property of the NB specification is that $\text{Var}(D_t)$ exceeds $\mathbb{E}[D_t]$ in a structurally interpretable manner, enabling dispersion to serve as a primitive parameter rather than a residual anomaly.

Despite its widespread use in statistics and actuarial applications, the NB process has played only a limited role in financial control problems. In financial econometrics, clustering is typically addressed through stochastic volatility or jump–diffusion frameworks, while discrete event counts are rarely embedded directly into optimization problems. Similarly, in the optimal execution literature, endogenous control variables such as trading intensity are chosen under stochastic cost dynamics, but incoming risk is usually modeled as Brownian motion or as a simple Poisson arrival process [14, 18]. These specifications abstract from persistent overdispersion and therefore do not capture clustered count risk.

To the best of our knowledge, no researcher has integrated an overdispersed count process into a structural optimization framework in which (i) risk arrives discretely, (ii) dispersion is state-dependent, (iii) control variables are chosen endogenously. Researchers treat these elements separately, but their joint interaction remains unexplored.

This gap is economically nontrivial. When dispersion varies across regimes, risk is not merely larger in expectation; its temporal geometry changes. Embedding an NB process into a dynamic control

problem therefore alters the shape of optimal policies, potentially generating threshold effects, interior adoption regions, and hysteresis phenomena that do not arise under equi-dispersed specifications. In this study, we develop such an integrated framework.

This structural gap becomes particularly salient in environments where execution mechanisms are automated. When policies are implemented through event-triggered systems, the statistical representation of risk directly interacts with the architecture of execution.

2.2. Automation and digital execution

A parallel stream of research examines how digitization reshapes the *execution layer* of contracting. Blockchain-based systems expand the contracting space through verifiable state-contingent execution, altering information transmission and equilibrium structure [27, 28]. In financial markets, decentralized consensus and smart contracts modify how agreements are triggered and enforced, with implications for competition, transparency, and strategic interaction [15, 29, 30].

In traditional finance, automation is embodied in algorithmic execution and automated clearing systems. The optimal execution literature formalizes how trading algorithms translate policy variables into executed quantities while managing market impact and risk exposure [14, 31]. Subsequent treatments incorporate market microstructure dynamics, inventory control, and real-time decision rules under uncertainty [18, 32–35]. Across these settings, execution rules convert information into actions through predefined response mechanisms.

A common structural feature emerges across blockchain protocols and algorithmic trading systems: Execution is typically *event-contingent*. Actions are triggered when discrete conditions are satisfied, such as order arrivals, margin thresholds, or state transitions, rather than responding continuously to infinitesimal diffusion shocks. Smart contracts activate upon verifiable events, and execution algorithms react to observable order flow or price-impact signals [14, 27, 32, 33, 36].

This architecture aligns naturally with discrete, count-based representations of risk. When shocks arrive in clustered bursts, the execution layer absorbs episodic spikes rather than smooth flows. A model based solely on Gaussian diffusion abstracts from this event-driven structure. Continuous diffusion implies smooth policy adjustments to infinitesimal shocks, whereas event-driven architectures respond to discrete realizations. Ignoring this distinction conflates volatility with arrival intensity and may distort optimal adoption boundaries. By contrast, an overdispersed count process captures the discreteness of activation triggers and the clustering of shocks that automated systems must process.

From this perspective, adopting a NB specification is not merely a statistical device to accommodate excess variance. It reflects the informational environment faced by automated systems: Risk materializes through discrete arrivals whose intensity varies across regimes. Embedding such a process into a policy-choice framework therefore mirrors the operational logic of digital execution.

2.3. Dynamic adjustment and hysteresis

When risk arrives in clustered bursts and execution mechanisms are event-contingent, a purely static adoption decision becomes incomplete. Automation is not chosen once and for all; rather, it is adjusted over time in response to evolving risk regimes. Firms may adopt automation during stress episodes, pause during calm periods, and delay reversal when uncertainty remains elevated. Such behavior naturally generates *hysteresis* driven by switching frictions and regime persistence.

The economics of irreversible or partially irreversible investment provides a canonical foundation for this phenomenon. With fixed or quasi-fixed adoption costs, optimal policies exhibit inaction regions and threshold behavior: Adjustment occurs only when state variables cross critical boundaries, implying history dependence even when fundamentals mean-revert [37, 38]. In these models, volatility reshapes optimal entry and exit regions, leading firms to delay otherwise positive net present value investments. This mechanism provides a natural microfoundation for dispersion-dependent threshold responses in our model.

Asset-pricing and macro-finance research further emphasizes that persistent second-moment dynamics alter optimal decision rules. Long-run risk models show that time-varying volatility and growth uncertainty materially affect equilibrium investment and pricing [39]. Models with rare disasters and time-varying tail risk demonstrate that low-probability clustered shocks have first-order effects on equilibrium behavior and return dynamics [40–42]. These frameworks imply that when risk intensity varies across regimes, optimal policies respond nonlinearly and asymmetrically.

Risk-sensitive and robustness-based control approaches provide an additional mechanism for persistence. When agents account for model uncertainty or tail ambiguity, optimal thresholds shift across regimes and reversals are delayed, amplifying hysteresis under clustered uncertainty [43–45]. In macro-financial data, downside risk and growth vulnerability vary systematically with financial conditions, consistent with regime-dependent amplification [46].

A complementary approach formalizes clustered risk through regime-switching dynamics. Markov regime-switching models parsimoniously capture persistent states such as “normal” and “stress” environments [47, 48]. Embedding such regimes into dynamic control problems yields tractable policy structures in which adjustment costs and state persistence jointly generate threshold behavior and delayed responses.

These strands of literature motivate three structural ingredients. First, an explicit state space \mathcal{S} captures persistent regimes. Second, dispersion measures such as $\text{Var}(D_t)$ summarize time-varying overdispersion in clustered count risk. Third, a dynamic automation variable evolves subject to switching costs. Together, these components support a transparent “static core + dynamic option” architecture: The static model characterizes equilibrium adoption under given dispersion, while the dynamic extension explains asymmetric entry–exit bands and persistent automation regimes. In particular, clustered overdispersion shifts optimal adoption thresholds endogenously, generating hysteresis even when average risk intensity normalizes.

The preceding strands of literature establish three structural facts: (i) Clustered risk generates state-dependent overdispersion, (ii) automated execution operates through event-contingent mechanisms, and (iii) irreversible adjustment under persistent regimes produces hysteresis. What remains missing is a unified optimization framework that integrates these elements into a single tractable structure.

In the next section, we develop such a framework. We first construct a static automation problem in which clustered risk enters exclusively through dispersion-sensitive linear operators on the probability simplex. This formulation isolates the geometric effect of overdispersion on optimal automation intensity while holding the mean risk level fixed. We then embed the static core into a dynamic regime-switching model with adjustment frictions, allowing dispersion to evolve endogenously across persistent states and generating structural hysteresis.

3. Static model under clustered risk

3.1. Probability-space primitives and operator representation

We begin with the exact NB model on the full support \mathbb{N}_0 , which provides the theoretical basis for the dispersion comparative statics developed below. For numerical implementation, we introduce a truncated finite-dimensional representation

$$\mathcal{D} = \{0, 1, \dots, M\},$$

and represent the clustered-risk count random variable through a probability vector $\boldsymbol{\pi}(r) \in \Delta^{M+1}$ induced by the truncated and renormalized NB distribution.

To isolate pure dispersion effects, we fix the mean $\mu > 0$ and parameterize the NB distribution using the mean-preserving form

$$D_r \sim \text{NB}(r, \theta(r)), \quad \theta(r) = \frac{r}{r + \mu},$$

so that

$$\mathbb{E}[D_r] = \mu, \quad \text{Var}(D_r) = \mu + \frac{\mu^2}{r}.$$

Smaller r implies stronger overdispersion and heavier tail clustering.

Let $p_d(r) = \mathbb{P}(D_r = d)$ for $d \in \mathbb{N}_0$ denote the full NB pmf. Define the truncated and renormalized probabilities on \mathcal{D} by

$$\pi_d(r) = \frac{p_d(r)}{\sum_{k=0}^M p_k(r)}, \quad d = 0, \dots, M,$$

and collect them as

$$\boldsymbol{\pi}(r) = (\pi_0(r), \dots, \pi_M(r))^\top, \quad \mathbf{1}^\top \boldsymbol{\pi}(r) = 1,$$

so that $\boldsymbol{\pi}(r) \in \Delta^{M+1}$.

For any vector $\mathbf{g} = (g(0), \dots, g(M))^\top$,

$$\mathbb{E}_{\boldsymbol{\pi}(r)}[g(D_r)] := \sum_{d=0}^M g(d) \pi_d(r) = \mathbf{g}^\top \boldsymbol{\pi}(r). \quad (3.1)$$

Let $\boldsymbol{\ell} = (\ell(0), \dots, \ell(M))^\top$ with $\ell(d) \geq 0$ and nondecreasing in d . Fix a stress threshold $d^* \in \mathcal{D}$ and define the tail-selection operator

$$\mathbb{T} := \text{diag}(\mathbf{1}\{d \geq d^*\}) \in \mathbb{R}^{(M+1) \times (M+1)}.$$

Define the non-tail and tail exposure components by

$$B_1(r) := \boldsymbol{\ell}^\top (\mathbf{I} - \mathbb{T}) \boldsymbol{\pi}(r), \quad B_2(r) := \boldsymbol{\ell}^\top \mathbb{T} \boldsymbol{\pi}(r). \quad (3.2)$$

Equivalently, the tail functional is

$$\text{Tail}(\boldsymbol{\ell}; \boldsymbol{\pi}(r)) = B_2(r) = (\mathbb{T} \boldsymbol{\ell})^\top \boldsymbol{\pi}(r). \quad (3.3)$$

3.2. Tail monotonicity under mean-preserving dispersion

Lemma 3.1 (NB tail monotonicity under mean preservation). *Let $D_r \sim \text{NB}(r, \theta(r))$ with*

$$\theta(r) = \frac{r}{r + \mu}, \quad \mu > 0,$$

defined on \mathbb{N}_0 (no truncation). Define the tail loss function

$$h(d) := \ell(d)\mathbf{1}\{d \geq d^*\},$$

and the tail exposure

$$B_2(r) := \mathbb{E}[h(D_r)] = \sum_{d=d^*}^{\infty} \ell(d) \mathbb{P}(D_r = d).$$

Assume that the function $h : \mathbb{N}_0 \rightarrow \mathbb{R}_+$ is nonnegative and nondecreasing, and is convex in the discrete sense, meaning that its second-order forward differences satisfy

$$\Delta^2 h(d) := h(d+2) - 2h(d+1) + h(d) \geq 0 \quad \text{for all } d \in \mathbb{N}_0.$$

Then $B_2(r)$ is decreasing in r . If h is strictly convex on a nontrivial subset of \mathbb{N}_0 that is assigned positive probability under the induced Poisson mixtures, then $B_2(r)$ is strictly decreasing in r .

Proof. Under the mean-preserving parameterization, the NB admits the Gamma–Poisson mixture representation:

$$D_r \mid \Lambda_r \sim \text{Poisson}(\Lambda_r), \quad \Lambda_r \sim \text{Gamma}(r, \mu/r),$$

with

$$\mathbb{E}[\Lambda_r] = \mu, \quad \text{Var}(\Lambda_r) = \mu^2/r.$$

Define

$$g(\lambda) := \mathbb{E}[h(N_\lambda)], \quad N_\lambda \sim \text{Poisson}(\lambda).$$

Thus, $g(\lambda)$ represents the expected tail loss generated by a Poisson count process with intensity λ . Then, by the law of total expectation,

$$B_2(r) = \mathbb{E}[g(\Lambda_r)].$$

Step 1: Monotonicity of g . Because the Poisson distribution is increasing in λ in the sense of first-order stochastic dominance, and h is nondecreasing, it follows that $g(\lambda)$ is nondecreasing in λ .

Step 2: Convexity of g . For Poisson expectations, the discrete derivative identities hold:

$$g'(\lambda) = \mathbb{E}[\Delta h(N_\lambda)], \quad g''(\lambda) = \mathbb{E}[\Delta^2 h(N_\lambda)],$$

where

$$\Delta h(d) = h(d+1) - h(d), \quad \Delta^2 h(d) = h(d+2) - 2h(d+1) + h(d).$$

Since h is nondecreasing, $\Delta h(d) \geq 0$, hence $g'(\lambda) \geq 0$. Since h is convex, $\Delta^2 h(d) \geq 0$, hence $g''(\lambda) \geq 0$. Thus g is increasing and convex.

Step 3: Dispersion ordering of Λ_r . For fixed mean μ , the Gamma family $\text{Gamma}(r, \mu/r)$ becomes more dispersed as the shape parameter r decreases while preserving the same mean. Equivalently, if $r_1 > r_2$, then Λ_{r_2} dominates Λ_{r_1} in convex order (i.e., $\Lambda_{r_2} \succeq_{cx} \Lambda_{r_1}$). This is a standard ordering property of the Gamma family under mean-preserving scale adjustment (see, e.g., [49]).

Step 4: Comparison of expectations. Since g is convex,

$$\mathbb{E}[g(\Lambda_{r_1})] \leq \mathbb{E}[g(\Lambda_{r_2})] \quad \text{whenever } r_1 > r_2.$$

Therefore $B_2(r)$ is decreasing in r .

Strict monotonicity follows when g is strictly convex on a region that is visited with positive probability under the induced Poisson mixtures, which holds if h is strictly convex on a nontrivial subset of \mathbb{N}_0 . \square

The monotonicity result in Lemma 3.1 is stated for the exact untruncated NB model on \mathbb{N}_0 . The truncated representation introduced above is used only for finite-dimensional computation and numerical implementation. For sufficiently large truncation level M , it provides an accurate approximation to the exact tail exposure while preserving the same comparative-static interpretation of dispersion.

3.3. Automation intensity, mitigation operators, and cost mapping

The firm chooses an automation intensity $a \in [0, 1]$, interpreted as the degree of automated execution (e.g., algorithmic clearing, smart-contract coverage, or rule-based enforcement), which is applied to mitigate clustered loss exposure.

Let $\mathbf{I} \in \mathbb{R}^{(M+1) \times (M+1)}$ denote the identity operator and recall the tail-selection operator $\mathbf{T} \in \mathbb{R}^{(M+1) \times (M+1)}$ defined in (3.3). We have the orthogonal decomposition

$$\mathbf{I} = (\mathbf{I} - \mathbf{T}) + \mathbf{T}.$$

Hence total expected loss admits the decomposition

$$\boldsymbol{\ell}^\top \boldsymbol{\pi}(r) = \boldsymbol{\ell}^\top (\mathbf{I} - \mathbf{T})\boldsymbol{\pi}(r) + \boldsymbol{\ell}^\top \mathbf{T}\boldsymbol{\pi}(r).$$

This guarantees that tail exposure is isolated without double-counting and provides a clean decomposition of total expected loss.

Automation may affect non-tail and tail exposure differently. Let

$$\phi, \psi : [0, 1] \rightarrow [0, 1]$$

be twice continuously differentiable functions, satisfying

$$\begin{aligned} \phi(0) = \psi(0) = 0, \quad \phi'(a) \geq 0, \quad \psi'(a) \geq 0, \\ \phi''(a) \leq 0, \quad \psi''(a) \leq 0. \end{aligned}$$

Here:

- $\phi(a)$ reduces non-tail exposure, - $\psi(a)$ reduces tail exposure.

Enabling $\psi \neq \phi$ captures disproportionate mitigation of extreme states.

Expected loss under dispersion parameter r is

$$L(a; r) = (1 - \phi(a)) \boldsymbol{\ell}^\top (I - \mathbb{T}) \boldsymbol{\pi}(r) + (1 - \psi(a)) \boldsymbol{\ell}^\top \mathbb{T} \boldsymbol{\pi}(r). \quad (3.4)$$

Using the decomposition in (3.2), expected loss under automation becomes

$$L(a; r) = B_1(r)(1 - \phi(a)) + B_2(r)(1 - \psi(a)). \quad (3.5)$$

Observe that

$$B_1(r) + B_2(r) = \boldsymbol{\ell}^\top \boldsymbol{\pi}(r).$$

Thus, $B_1(r)$ captures non-tail exposure and $B_2(r)$ captures tail exposure. Here, $B_1(r)$ and $B_2(r)$ denote the truncated finite-dimensional approximations of the corresponding exposure components defined in Section 3.1.

Automation incurs a convex implementation cost

$$K : [0, 1] \rightarrow \mathbb{R}_+, \quad K'(a) \geq 0, \quad K''(a) \geq 0. \quad (3.6)$$

A canonical quadratic specification is

$$K(a) = k_1 a + k_2 a^2, \quad k_2 > 0.$$

Total static expected cost is

$$C(a; r) = L(a; r) + K(a), \quad (3.7)$$

i.e.,

$$C(a; r) = B_1(r)(1 - \phi(a)) + B_2(r)(1 - \psi(a)) + K(a). \quad (3.8)$$

The firm solves

$$\min_{a \in [0, 1]} C(a; r). \quad (3.9)$$

In many financial environments, automation is implemented through state-contingent mechanisms that activate under stress (e.g., margin calls, circuit breakers, and liquidation triggers), implying that $\psi'(a)$ may exceed $\phi'(a)$ over relevant regions.

3.4. Structural assumptions

We now state the structural assumptions required for well-posedness, existence, uniqueness, and dispersion-based comparative statics of the static automation problem.

Assumption 3.2 (Dispersion and exposure regularity). Let $\boldsymbol{\pi}(r) \in \Delta^{M+1}$ denote the truncated and renormalized NB probability vector introduced in Section 3.1. The normalization constant $Z(r) = \sum_{k=0}^M p_k(r)$ is strictly positive and continuously differentiable in r , ensuring that $\boldsymbol{\pi}(r)$ is continuously differentiable. Then:

(i) The exposure components

$$B_1(r) := \boldsymbol{\ell}^\top (I - \mathbb{T})\boldsymbol{\pi}(r), \quad B_2(r) := \boldsymbol{\ell}^\top \mathbb{T}\boldsymbol{\pi}(r),$$

are continuously differentiable in r .

(ii) The loss function $\ell(d)$ is nonnegative and nondecreasing in d .

Under this parameterization, smaller values of r correspond to stronger overdispersion. Monotonicity of the exact tail exposure with respect to dispersion is established in Lemma 3.1. The truncated representation used here provides a finite-dimensional approximation that preserves the same comparative-static behavior for sufficiently large truncation level M .

Assumption 3.3 (Loss boundedness). The loss function $\ell : \mathbb{N}_0 \rightarrow \mathbb{R}_+$ is nonnegative and nondecreasing, that is, $\ell(d) \geq 0$ for all $d \in \mathbb{N}_0$ and $\ell(d+1) \geq \ell(d)$ for all d . Moreover, it is uniformly bounded in the sense that

$$\sup_{d \geq 0} \ell(d) < \infty.$$

These conditions ensure finite expectations and preserve monotonicity of tail exposure.

Assumption 3.4 (Automation effectiveness). The mitigation functions $\phi, \psi : [0, 1] \rightarrow [0, 1]$ represent the proportion of reduction in non-tail and tail exposure, respectively. They are twice continuously differentiable and satisfy

$$\begin{aligned} \phi(0) = \psi(0) = 0, \quad \phi'(a) \geq 0, \quad \psi'(a) \geq 0, \\ \phi''(a) \leq 0, \quad \psi''(a) \leq 0. \end{aligned}$$

Moreover,

$$\psi'(a) \geq \phi'(a) \quad \text{for all } a \in [0, 1],$$

so automation is weakly more effective on tail exposure.

Thus, automation does not increase exposure and exhibits diminishing marginal effectiveness in both non-tail and tail risk mitigation.

Assumption 3.5 (Automation cost convexity). The implementation cost function $K : [0, 1] \rightarrow \mathbb{R}_+$ is twice continuously differentiable and satisfies

$$K'(a) \geq 0, \quad K''(a) > 0 \quad \text{for all } a \in (0, 1).$$

Furthermore, for each fixed r ,

$$K''(a) > -\phi''(a) B_1(r) - \psi''(a) B_2(r),$$

ensuring strict convexity of the total cost function $C(a; r)$.

Table 1. Model notation (common, static, and dynamic components).

Symbol	Type	Description
Panel A: Common structure		
\mathbb{N}_0	Set	Nonnegative integers (full support).
$\mathcal{D} = \{0, \dots, M\}$	Set	Truncated support for finite-dimensional representation.
M	Scalar	Truncation level ($ \mathcal{D} = M + 1$).
Δ^{M+1}	Set	Probability simplex $\{q \in \mathbb{R}_+^{M+1} : \mathbf{1}^\top q = 1\}$.
$\mathbf{1}$	Vector	All-ones vector.
ℓ	Vector	Loss vector; nonnegative and nondecreasing.
d^*	Scalar	Tail threshold defining stress region.
I	Operator	Identity operator.
T	Operator	Tail-selection operator $\text{diag}(\mathbf{1}\{d \geq d^*\})$.
Panel B: Static model (section 3)		
$r > 0$	Scalar	Dispersion parameter (smaller r implies stronger overdispersion).
μ	Scalar	Fixed mean of the distribution.
$\theta(r)$	Scalar	Mean-preserving NB parameterization $\frac{r}{r+\mu}$.
D_r	Random variable	NB random variable.
$p_d(r)$	Scalar	NB probability mass function.
$Z(r)$	Scalar	Normalization constant: $Z(r) = \sum_{k=0}^M p_k(r)$.
$\pi_d(r)$	Scalar	Truncated probability: $p_d(r)/Z(r)$.
$\pi(r)$	Vector	Truncated probability vector in Δ^{M+1} .
$B_1(r)$	Scalar	Non-tail exposure.
$B_2(r)$	Scalar	Truncated tail exposure.
a	Control	Automation intensity.
$\phi(a)$	Function	Non-tail mitigation (increasing, concave).
$\psi(a)$	Function	Tail mitigation (increasing, concave).
$K(a)$	Function	Automation cost (increasing, convex).
$C(a; r)$	Function	Total static cost.
C_a, C_{aa}	Derivatives	First and second derivatives of $C(a; r)$.
$F(a, r)$	Function	First-order condition map $F(a, r) = C_a(a; r)$.
$a^*(r)$	Policy	Static optimal automation.
r^*	Scalar	Dispersion threshold solving $C_a(0; r^*) = 0$.
Panel C: Dynamic model (section 4)		
$\mathcal{S} = \{1, \dots, S\}$	Set	Finite ordered regime space.
$\mathcal{X} = \mathcal{S} \times [0, 1]$	Set	State space.
s_t	State	Regime at time t .
a_t	Control	Automation choice at time t .
a^-	State	Previous automation level.
$x_t = (s_t, a_{t-1})$	State	Dynamic state vector.
r_s	Scalar	Baseline regime-dependent dispersion.
$r_s(a)$	Function	Automation-dependent dispersion.
$\pi(r_s(a))$	Vector	Truncated NB probability vector.
$B_1(s, a)$	Scalar	Non-tail exposure.
$B_2(s, a)$	Scalar	Truncated tail exposure.
$P(a)$	Kernel	Transition matrix; row-stochastic.
$c(s, a)$	Function	Flow cost.
$\Gamma(a, a^-)$	Function	Adjustment cost.
$\tilde{c}(s, a, a^-)$	Function	Total one-period cost.
$\beta \in (0, 1)$	Parameter	Discount factor.
$V(s, a^-)$	Function	Value function.
\mathcal{T}	Operator	Bellman operator.
$H(s, a)$	Function	Smooth Bellman component.
$\partial_a H(s, a)$	Scalar	Marginal derivative of H with respect to a .
$a^*(s, a^-)$	Policy	Dynamic optimal automation.
$\underline{a}(s), \bar{a}(s)$	Scalars	Hysteresis band bounds.

Note: All exposure quantities (B_1, B_2) correspond to truncated finite-dimensional approximations of the exact NB model used in the theoretical analysis.

Assumption 3.6 (Threshold structure). The derivative of the total cost at zero automation,

$$C_a(0; r) = -\phi'(0) B_1(r) - \psi'(0) B_2(r) + K'(0),$$

is continuous and strictly increasing in r .

Moreover, there exist $r_L < r_H$ such that

$$C_a(0; r_L) < 0, \quad C_a(0; r_H) \geq 0.$$

Then there exists a unique threshold $r^* > 0$, satisfying

$$a^*(r) = 0 \quad \text{for } r \geq r^*, \quad a^*(r) > 0 \quad \text{for } r < r^*.$$

For clarity, Table 1 summarizes the unified notation for both the static clustered-risk automation model and its dynamic regime-switching extension. The structure separates (i) probabilistic primitives governing dispersion, (ii) linear operator representations of exposure, and (iii) the automation control problem in both static and dynamic settings. In particular, dispersion affects the optimization problem only through the induced probability vector $\pi(r)$ in the static model and its regime-dependent counterpart $\pi(r_s(a))$ in the dynamic model. All exposure quantities (B_1, B_2) are defined as linear functionals of these probability vectors, while persistence and path dependence in the dynamic model arise through the interaction of the transition kernel and adjustment costs.

Table 2. Static structural dependency structure.

Object	Structural dependence
NB parameter $\theta(r)$	Mean-preserving mapping $\theta(r) = \frac{r}{r+\mu}$; dispersion governed solely by r (mean fixed).
Probability vector $\pi(r)$	Truncated and renormalized NB probability vector $\pi(r) \in \Delta^{M+1}$ induced by (r, μ) . All dispersion effects enter through $\pi(r)$.
Non-tail exposure $B_1(r)$	Linear functional $B_1(r) = \ell^\top (1 - \mathsf{T})\pi(r)$.
Tail exposure $B_2(r)$	Linear functional $B_2(r) = \ell^\top \mathsf{T}\pi(r)$; monotone in dispersion by Lemma 3.1.
Expected loss $L(a; r)$	Affine transformation of $(B_1(r), B_2(r))$: $L(a; r) = (1 - \phi(a))B_1(r) + (1 - \psi(a))B_2(r)$.
Total cost $C(a; r)$	Composition of exposure mapping and convex automation cost: $C(a; r) = L(a; r) + K(a)$.
First-order condition	$C_a(a; r) = -\phi'(a)B_1(r) - \psi'(a)B_2(r) + K'(a) = 0$.
Optimality map $F(a, r)$	$F(a, r) = C_a(a; r)$; dispersion affects F only through $(B_1(r), B_2(r))$.
Optimal policy $a^*(r)$	Unique solution of $F(a, r) = 0$; dispersion influences $a^*(r)$ exclusively via exposure terms.
Dispersion threshold r^*	Unique value solving $C_a(0; r) = 0$, separating $a^*(r) = 0$ from $a^*(r) > 0$.

Table 2 formalizes the structural flow of dependence in the static model. The dispersion parameter r determines the NB parameter $\theta(r)$ and thereby the probability vector $\pi(r)$. All exposure components are linear functionals of $\pi(r)$, while automation enters exclusively through nonlinear mitigation and convex cost mappings. Consequently, the comparative statics with respect to dispersion operate through the exposure pair $(B_1(r), B_2(r))$.

3.5. Static optimality conditions and structural results

We now characterize the solution of the static automation problem (3.9) and establish its structural dependence on dispersion parameter r .

By Lemma 3.1, the exact tail exposure is strictly decreasing in r under the mean-preserving NB parameterization. The truncated representation used here preserves this monotonicity for sufficiently large truncation level. Hence, stronger overdispersion (smaller r) increases expected tail losses while preserving the mean.

Recall that the total expected cost is

$$C(a; r) = (1 - \phi(a)) B_1(r) + (1 - \psi(a)) B_2(r) + K(a). \quad (3.10)$$

Differentiating with respect to a yields

$$C_a(a; r) = -\phi'(a) B_1(r) - \psi'(a) B_2(r) + K'(a). \quad (3.11)$$

An interior optimizer $a^*(r) \in (0, 1)$ satisfies

$$K'(a^*) = \phi'(a^*) B_1(r) + \psi'(a^*) B_2(r). \quad (3.12)$$

All dependence on dispersion enters exclusively through $B_1(r)$ and $B_2(r)$.

Lemma 3.7 (Strict convexity). *Under Assumptions 3.4 and 3.5, and provided that*

$$K''(a) > -\phi''(a) B_1(r) - \psi''(a) B_2(r) \quad \text{for all } a \in (0, 1), r > 0,$$

the cost function $C(a; r)$ is strictly convex in a for each fixed r .

Proof. The second derivative is

$$C_{aa}(a; r) = -\phi''(a) B_1(r) - \psi''(a) B_2(r) + K''(a).$$

Under the stated curvature condition, $C_{aa}(a; r) > 0$ for all $a \in (0, 1)$, which establishes strict convexity. \square

Theorem 3.8 (Existence and uniqueness). *For each $r > 0$, there exists a unique global minimizer $a^*(r) \in [0, 1]$, solving (3.9).*

Proof. By Lemma 3.7, $C(\cdot; r)$ is strictly convex and continuous. Compactness of $[0, 1]$ guarantees existence, and strict convexity guarantees uniqueness. \square

Theorem 3.9 (Dispersion monotonicity). *Suppose Assumptions 3.2–3.5 hold, and assume additionally that $B'_1(r) \leq 0$ and $B'_2(r) < 0$. On any region where $a^*(r) \in (0, 1)$,*

$$\frac{\partial a^*(r)}{\partial r} = -\frac{F_r}{F_a}, \quad F(a, r) = C_a(a; r),$$

and therefore

$$\frac{\partial a^*(r)}{\partial r} < 0 \quad \text{whenever} \quad -\psi'(a^*)B'_2(r) > \phi'(a^*)|B'_1(r)|.$$

Proof. At an interior optimum, $F(a^*(r), r) = 0$ and $F_a = C_{aa}(a^*(r); r) > 0$. Differentiating with respect to r gives

$$F_r = -\phi'(a)B'_1(r) - \psi'(a)B'_2(r).$$

Under the stated inequality and monotonicity conditions, $F_r > 0$. The implicit function theorem yields

$$\frac{\partial a^*(r)}{\partial r} = -\frac{F_r}{F_a} < 0.$$

□

Proposition 3.10 (Closed-form characterization under quadratic cost). *Suppose the control variable affects only the tail component, so that $\phi(a) \equiv 0$, and let*

$$\psi(a) = a, \quad K(a) = k_1 a^2, \quad k_1 > 0.$$

This specification is consistent with Assumptions 3.4 and 3.5. Then the unique minimizer satisfies

$$a^*(r) = \min \left\{ 1, \frac{B_2(r)}{2k_1} \right\}.$$

Moreover, there exists a unique threshold r^ , satisfying*

$$B_2(r^*) = 2k_1,$$

such that

$$a^*(r) = \begin{cases} 1, & r < r^*, \\ \frac{B_2(r)}{2k_1}, & r \geq r^*. \end{cases}$$

The threshold r^ is strictly decreasing in k_1 .*

Proof. Under the stated specialization, the first-order condition (3.12) reduces to

$$2k_1 a^* = B_2(r).$$

Interior solutions therefore satisfy $a^*(r) = B_2(r)/(2k_1)$. Imposing the constraint $a \in [0, 1]$ yields the stated truncation. Strict monotonicity of $B_2(r)$ implies the existence and uniqueness of r^* , solving $B_2(r^*) = 2k_1$. Implicit differentiation gives

$$\frac{dr^*}{dk_1} = \frac{2}{B'_2(r^*)} < 0,$$

since $B'_2(r^*) < 0$.

□

Theorem 3.11 (Dispersion threshold). *Suppose $B_1(r)$ is nonincreasing and $B_2(r)$ is strictly decreasing, so that the mapping $r \mapsto C_a(0; r)$ is continuous and strictly increasing. If there exist $r_L < r_H$ such that*

$$C_a(0; r_L) < 0 \quad \text{and} \quad C_a(0; r_H) \geq 0.$$

then there exists a unique threshold $r^ > 0$ such that*

$$a^*(r) = 0 \quad \text{for } r \geq r^* \quad \text{and} \quad a^*(r) > 0 \quad \text{for } r < r^*.$$

Proof. Continuity and strict monotonicity imply existence and uniqueness of r^* , satisfying $C_a(0; r^*) = 0$. Strict convexity implies that $a = 0$ is optimal if and only if $C_a(0; r) \geq 0$. \square

To visualize the structural transition from dispersion-driven static optimization to the dynamic regime-switching extension, Figure 1 summarizes the logical architecture of the model.

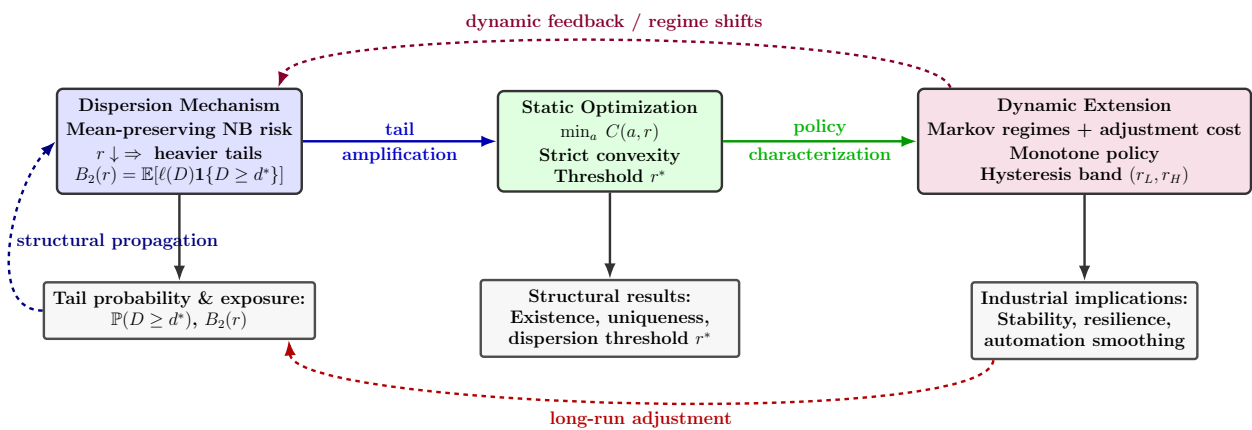


Figure 1. Structural architecture of the dispersion–threshold mechanism. Solid arrows represent core optimization logic, while dashed arrows indicate interpretative dynamic feedback.

4. Dynamic regime-switching model with adjustment frictions

In Section 3, we established that dispersion-sensitive tail exposure induces a threshold structure in optimal automation. We now embed the static environment into a dynamic regime-switching model with costly adjustment, generating endogenous hysteresis.

Automation becomes a persistent state variable, and clustered-risk shocks propagate intertemporally through regime persistence and adjustment frictions.

4.1. State space, regime-dependent dispersion, and structural feedback

Time is discrete, $t = 0, 1, 2, \dots$. In contrast to the static environment of Section 3, automation now becomes a persistent state variable. At each date t , the economy is characterized by the state vector

$$x_t = (s_t, a_{t-1}),$$

where $s_t \in \mathcal{S}$ denotes the current risk regime and $a_{t-1} \in [0, 1]$ is previously installed automation intensity.

Let $\mathcal{S} = \{1, 2, \dots, S\}$ be a finite ordered set equipped with its natural order, where larger indices correspond to more severe regimes.

In the static model, clustered risk is summarized by a scalar dispersion parameter $r > 0$, with tail exposure decreasing in r under mean preservation. The dynamic extension replaces this scalar by a regime-dependent family $\{r_s\}_{s \in \mathcal{S}}$.

Each regime s is associated with $r_s > 0$, satisfying

$$r_1 > r_2 > \dots > r_S. \quad (4.1)$$

Since the mean μ is fixed, smaller r_s implies larger variance

$$\text{Var}(D) = \mu + \frac{\mu^2}{r_s},$$

and therefore stronger overdispersion and heavier tail clustering.

Each regime induces a truncated probability vector

$$\pi(r_s) \in \Delta^{M+1}.$$

Define exposure components

$$B_1(r_s) = \ell^\top (I - \mathbb{T}) \pi(r_s), \quad B_2(r_s) = \ell^\top \mathbb{T} \pi(r_s). \quad (4.2)$$

By Lemma 3.1, the exact tail exposure is strictly decreasing in r . The truncated representation used here preserves this ordering for sufficiently large M . Hence,

$$B_2(r_1) \leq B_2(r_2) \leq \dots \leq B_2(r_S). \quad (4.3)$$

Thus, higher-index regimes exhibit strictly larger tail exposure.

Automation modifies dispersion through $r_s(a)$ with

$$r'_s(a) \geq 0.$$

Exposure becomes

$$B_j(s, a) = B_j(r_s(a)), \quad j = 1, 2.$$

Since the exact tail exposure is decreasing in r , and the truncated representation preserves this property, the chain rule yields

$$\frac{\partial}{\partial a} B_2(r_s(a)) = \frac{\partial B_2}{\partial r}(r_s(a)) r'_s(a) \leq 0. \quad (4.4)$$

Thus automation weakly reduces tail exposure within each regime.

Automation also affects regime transitions via

$$P(a) = [P_{ss'}(a)], \quad \sum_{s'} P_{ss'}(a) = 1.$$

Assumption 4.1 (Automation structure). For each $s \in \mathcal{S}$, $r_s(a)$ is continuously differentiable, bounded, and satisfies $r'_s(a) \geq 0$.

For each s, s' , $P_{ss'}(a)$ is continuously differentiable in a .

Moreover, for each fixed s and all $a' \geq a$, the transition distributions satisfy first-order stochastic dominance:

$$\sum_{s' \geq k} P_{ss'}(a') \geq \sum_{s' \geq k} P_{ss'}(a) \quad \text{for all } k.$$

4.2. Per-period cost and adjustment frictions

Given the current regime $s \in \mathcal{S}$ and previously installed automation level $a^- \in [0, 1]$, the firm chooses current automation $a \in [0, 1]$. Automation is therefore an endogenous persistent state variable, and the state space is

$$\mathcal{X} = \mathcal{S} \times [0, 1], \quad x = (s, a^-).$$

The functions $\phi, \psi : [0, 1] \rightarrow [0, 1]$ represent automation effectiveness in reducing non-tail and tail exposure, respectively.

Assumption 4.2 (Mitigation functions). The functions ϕ and ψ are twice continuously differentiable on $[0, 1]$ and satisfy

$$\begin{aligned} \phi(0) = \psi(0) = 0, \quad \phi'(a) \geq 0, \quad \psi'(a) \geq 0, \\ \phi''(a) \leq 0, \quad \psi''(a) \leq 0. \end{aligned}$$

Moreover, $\psi'(a) \geq \phi'(a)$ for all $a \in [0, 1]$.

Conditional on regime s , the contemporaneous operating cost is

$$c(s, a) = B_1(s, a)(1 - \phi(a)) + B_2(s, a)(1 - \psi(a)) + K(a), \quad (4.5)$$

where $K : [0, 1] \rightarrow \mathbb{R}_+$ is the implementation cost, and we use the shorthand $B_j(s, a) := B_j(r_s(a))$ for $j \in \{1, 2\}$.

Assumption 4.3 (Convexity and increasing differences). The function $c(s, a)$ is twice continuously differentiable in a . For each $s \in \mathcal{S}$, the mapping $a \mapsto c(s, a)$ is strictly convex on $[0, 1]$. Moreover, $c(s, a)$ has increasing differences in (s, a) with respect to the natural order on \mathcal{S} and the usual order on $[0, 1]$.

Automation changes incur costs

$$\Gamma(a, a^-) = \kappa_1 |a - a^-| + \frac{\kappa_2}{2} (a - a^-)^2, \quad (4.6)$$

with $\kappa_1, \kappa_2 \geq 0$. The function Γ is continuous on $[0, 1]^2$ and convex in a for each fixed a^- .

Define

$$\tilde{c}(s, a, a^-) = c(s, a) + \Gamma(a, a^-). \quad (4.7)$$

Assumption 4.4 (Bounded cost). The function $\tilde{c}(s, a, a^-)$ is bounded on $\mathcal{S} \times [0, 1]^2$.

For each fixed (s, a^-) , the mapping $a \mapsto \tilde{c}(s, a, a^-)$ is continuous and strictly convex on $[0, 1]$.

Let $\beta \in (0, 1)$. Given state (s, a^-) , choosing a induces the next-period regime distribution via $P(a)$ and sets the next automation state equal to a . The value function satisfies

$$V(s, a^-) = \min_{a \in [0, 1]} \left\{ \tilde{c}(s, a, a^-) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) \right\}. \quad (4.8)$$

Define the Bellman operator

$$(\mathcal{T}V)(s, a^-) = \min_{a \in [0, 1]} \left\{ \tilde{c}(s, a, a^-) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) \right\}.$$

Proposition 4.5 (Bellman contraction and fixed-point uniqueness). *Under Assumption 4.4, row-stochasticity of $P(a)$ for each $a \in [0, 1]$, and $\beta \in (0, 1)$, operator \mathcal{T} is a contraction on $\mathcal{B}(\mathcal{X})$ under the sup norm with modulus β :*

$$\|\mathcal{T}V - \mathcal{T}W\|_\infty \leq \beta \|V - W\|_\infty.$$

Hence, there exists a unique bounded fixed point V , satisfying $V = \mathcal{T}V$.

Proof. For any bounded functions V, W and any $(s, a^-) \in \mathcal{X}$, for each $a \in [0, 1]$,

$$\left| \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) - \sum_{s' \in \mathcal{S}} P_{ss'}(a) W(s', a) \right| \leq \|V - W\|_\infty,$$

since $P(a)$ is row-stochastic. Therefore, for every admissible a ,

$$\left| \tilde{c}(s, a, a^-) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) - \tilde{c}(s, a, a^-) - \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) W(s', a) \right| \leq \beta \|V - W\|_\infty.$$

Taking the minimum over $a \in [0, 1]$ preserves nonexpansiveness, yielding

$$\|\mathcal{T}V - \mathcal{T}W\|_\infty \leq \beta \|V - W\|_\infty.$$

Hence, \mathcal{T} is a contraction. Since $\beta \in (0, 1)$, and $\mathcal{B}(\mathcal{X})$ is complete under the sup norm, the Banach fixed-point theorem implies existence and uniqueness of a bounded fixed point V , satisfying $V = \mathcal{T}V$. \square

4.3. Dynamic structural results

We now establish the structural properties of the dynamic automation problem. Throughout this subsection, let the state space be $\mathcal{X} = \mathcal{S} \times [0, 1]$, where \mathcal{S} is a finite totally ordered set of regimes, and $a^- \in [0, 1]$ denotes previously installed automation.

Assumption 4.6 (Dynamic primitives). The discount factor satisfies $\beta \in (0, 1)$. The per-period cost $c(s, a)$ is twice continuously differentiable and strictly convex in a for each $s \in \mathcal{S}$. The adjustment cost $\Gamma(a, a^-)$ is continuous on $[0, 1]^2$ and convex in both arguments. The transition kernel $P(a) = [P_{ss'}(a)]$ is row-stochastic and continuous in a .

Assumption 4.7 (Increasing differences). The flow cost $c(s, a)$ has increasing differences in (s, a) . Moreover, for any bounded function $W(s, a^-)$ that is nondecreasing in s for each fixed a^- , the continuation mapping

$$(s, a) \mapsto \sum_{s' \in \mathcal{S}} P_{ss'}(a) W(s', a)$$

has increasing differences in (s, a) .

Assumption 4.7 ensures that the transition kernel preserves regime monotonicity in a manner compatible with monotone comparative statics.

Lemma 4.8 (Existence, uniqueness, and regularity). *The Bellman operator*

$$(\mathcal{T}V)(s, a^-) = \min_{a \in [0,1]} \left\{ c(s, a) + \Gamma(a, a^-) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) \right\}$$

is a contraction on $\mathcal{B}(\mathcal{X})$ under the sup norm with modulus β . Hence, there exists a unique bounded continuous value function V , satisfying $V = \mathcal{T}V$. Furthermore, $V(s, a^-)$ is nondecreasing in s .

Proof. For any bounded V_1, V_2 ,

$$\left| \sum_{s'} P_{ss'}(a) V_1(s', a) - \sum_{s'} P_{ss'}(a) V_2(s', a) \right| \leq \|V_1 - V_2\|_\infty,$$

since $P(a)$ is row-stochastic. Taking the minimum over $a \in [0, 1]$ preserves nonexpansiveness, yielding

$$\|\mathcal{T}V_1 - \mathcal{T}V_2\|_\infty \leq \beta \|V_1 - V_2\|_\infty.$$

Therefore, \mathcal{T} is a contraction.

To show monotonicity in s , initialize $V^{(0)}(s, a^-) = 0$. Suppose $V^{(n)}$ is nondecreasing in s . By Assumption 4.7, the continuation term is nondecreasing in s . The objective inside the minimization is therefore nondecreasing in s . Minimization over a preserves monotonicity. Thus, $V^{(n+1)}$ is nondecreasing in s . Convergence of value iteration implies V is nondecreasing in s . \square

Theorem 4.9 (Monotone automation in regime). *Under Assumptions 4.6 and 4.7, the optimal policy satisfies*

$$a^*(s', a^-) \geq a^*(s, a^-) \quad \text{whenever } s' \succeq s.$$

Proof. Fix a^- and define the objective

$$\Psi(s, a) = c(s, a) + \Gamma(a, a^-) + \beta \sum_{s'} P_{ss'}(a) V(s', a).$$

By Assumption 4.7, both the flow component and the continuation component have increasing differences in (s, a) . Since $\Gamma(a, a^-)$ does not depend on s , $\Psi(s, a)$ has increasing differences in (s, a) .

The feasible set $[0, 1]$ is a compact lattice, and strict convexity in a ensures a unique minimizer. Topkis' monotonicity theorem then implies that the argmin correspondence is increasing in s . \square

Theorem 4.10 (Dynamic hysteresis band). *Suppose $\Gamma(a, a^-) = \kappa_1|a - a^-|$ with $\kappa_1 > 0$. Define*

$$H(s, a) = c(s, a) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a)V(s', a).$$

If $H(s, a)$ is continuously differentiable and strictly convex in a , then for each regime s there exist thresholds $\underline{a}(s) \leq \bar{a}(s)$ such that

$$a^*(s, a^-) = \begin{cases} \underline{a}(s), & a^- < \underline{a}(s), \\ a^-, & a^- \in [\underline{a}(s), \bar{a}(s)], \\ \bar{a}(s), & a^- > \bar{a}(s). \end{cases}$$

The inaction region satisfies

$$|\partial_a H(s, a^-)| \leq \kappa_1.$$

Proof. Fix (s, a^-) and define

$$J(a) = H(s, a) + \kappa_1|a - a^-|.$$

Since H is strictly convex and $|a - a^-|$ is convex, J is strictly convex and has a unique minimizer.

For $a > a^-$,

$$\partial_a J(a) = \partial_a H(s, a) + \kappa_1,$$

and for $a < a^-$,

$$\partial_a J(a) = \partial_a H(s, a) - \kappa_1.$$

At $a = a^-$,

$$\partial J(a^-) = [\partial_a H(s, a^-) - \kappa_1, \partial_a H(s, a^-) + \kappa_1].$$

Thus, $a = a^-$ is optimal iff

$$|\partial_a H(s, a^-)| \leq \kappa_1.$$

Since $H(s, a)$ is strictly convex in a , its derivative $\partial_a H(s, a)$ is strictly increasing. This implies the existence of unique $\underline{a}(s)$ and $\bar{a}(s)$. \square

Theorem 4.11 (Crisis persistence and path dependence). *Suppose Assumptions 4.6 and 4.7 hold. Assume linear adjustment cost*

$$\Gamma(a, a^-) = \kappa_1|a - a^-|, \quad \kappa_1 > 0,$$

and let

$$H(s, a) = c(s, a) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a)V(s', a)$$

be continuously differentiable and strictly convex in a for each regime s .

Let $s_L \prec s_H$ denote a calm and a stress regime. Let $a_H = a^(s_H, a^-)$ denote the automation level chosen in regime s_H for some previous installation a^- . Let a_L^{stat} denote the static optimizer in regime s_L , defined as the unique solution of*

$$\partial_a c(s_L, a_L^{stat}) = 0.$$

Suppose

$$a_H > a_L^{stat}$$

and the calm-regime dynamic marginal incentive to reduce automation at a_H satisfies

$$\partial_a H(s_L, a_H) < \kappa_1. \quad (4.9)$$

Then the post-transition optimal policy satisfies

$$a^*(s_L, a_H) \geq a_H > a_L^{stat}.$$

In particular, the upward automation adjustment taken in the stress regime persists strictly above the calm-regime static optimum for at least one period after the transition.

Proof. Fix the post-transition state $(s_L, a^-) = (s_L, a_H)$. Define the objective

$$J(a) = H(s_L, a) + \kappa_1 |a - a_H|.$$

Since $H(s_L, a)$ is strictly convex and $|a - a_H|$ is convex, J is strictly convex and admits a unique minimizer, denoted $a_L^* := a^*(s_L, a_H)$.

We show that $a_L^* \geq a_H$.

Suppose instead that $a_L^* < a_H$. On the region $a < a_H$, the objective is differentiable with

$$\partial_a J(a) = \partial_a H(s_L, a) - \kappa_1.$$

The first-order condition for an interior minimizer in this region requires

$$\partial_a H(s_L, a_L^*) = \kappa_1.$$

Since $H(s_L, \cdot)$ is strictly convex, $\partial_a H(s_L, a)$ is strictly increasing in a . Because $a_L^* < a_H$, we have

$$\partial_a H(s_L, a_L^*) < \partial_a H(s_L, a_H).$$

Thus,

$$\kappa_1 = \partial_a H(s_L, a_L^*) < \partial_a H(s_L, a_H),$$

which contradicts condition (4.9).

Therefore, $a_L^* \geq a_H$. Finally, since $a_H > a_L^{stat}$ by hypothesis, it follows that

$$a^*(s_L, a_H) = a_L^* \geq a_H > a_L^{stat}.$$

□

The unified notation table introduced in Section 3 (Table 1, Panel C) summarizes the core objects of the dynamic regime-switching model. The structure separates (i) regime-dependent dispersion primitives $r_s(a)$ and the induced probability vector $\pi(r_s(a))$, (ii) linear exposure operators (B_1, B_2) that transmit dispersion into losses, (iii) intertemporal propagation through the automation-dependent transition kernel $P(a)$, and (iv) adjustment frictions via $\Gamma(a, a^-)$. This dynamic structure builds directly on the static formulation in Section 3, extending the dispersion-driven exposure mapping to a regime-dependent and intertemporal setting.

Table 3. Dynamic structural dependency structure.

Object	Structural dependence
Regime evolution s_{t+1}	Controlled markov transition: $\mathbb{P}(s_{t+1} = s' s_t = s, a_t = a) = P_{ss'}(a)$.
Within-regime dispersion	Automation shifts dispersion via $r_s(a)$, altering $\pi(r_s(a))$ and tail clustering.
Probability vector $\pi(r_s(a))$	Determined by $r_s(a)$ (mean fixed); all dispersion effects enter through this vector.
Exposure (B_1, B_2)	Linear operators applied to $\pi(r_s(a))$.
Flow cost $c(s, a)$	Depends on exposure and mitigation functions $(B_1, B_2, \phi, \psi, K)$.
Adjustment cost $\Gamma(a, a^-)$	Depends on (a, a^-) and friction parameters (κ_1, κ_2) .
Bellman operator \mathcal{T}	Defined by $(c, \Gamma, \beta, P(\cdot))$.
Value function V	Unique contraction fixed point of \mathcal{T} .
Dynamic marginal $\partial_a H(s, a)$	Includes: static marginal cost, continuation-value derivative, and transition-feedback term.
Policy $a^*(s, a^-)$	Determined jointly by dispersion mitigation $r_s(a)$ and transition kernel $P(a)$.
Hysteresis band	Characterized by $ \partial_a H(s, a^-) \leq \kappa_1$.
Persistence effect	Arises from regime persistence in $P(a)$ combined with adjustment frictions.

Table 3 formalizes the structural flow of dependence in the dynamic model. Within each regime, automation shifts dispersion through $r_s(a)$, thereby modifying the probability vector $\pi(r_s(a))$ and the exposure pair (B_1, B_2) . Intertemporally, automation influences the regime process through the controlled transition kernel $P(a)$. Consequently, dynamic incentives operate through two distinct channels: A within-regime dispersion-mitigation channel and a regime-transition channel. Adjustment frictions then generate hysteresis by preventing instantaneous realignment with the static optimum.

To clarify the structural role of the maintained assumptions across the static and dynamic components of the model, Table 4 maps each assumption to the lemmas, propositions, and theorems it directly supports. Moreover, the table distinguishes assumptions that are indispensable for existence, uniqueness, and dispersion-threshold results in the static model, and those that become essential only in the dynamic extension for establishing monotonicity, hysteresis, and stability properties. This structural mapping makes transparent the minimality and non-redundancy of the maintained assumptions and clarifies the logical architecture of the theory.

Table 4. Structural dependency of assumptions across static and dynamic theoretical results. ✓ = directly required; ○ = auxiliary.

Assumption	Lemma S1: Tail convexity	Prop S2: Threshold existence	Thm S3: Uniqueness	Cor S4: Comparative statics	Prop D1: Regime value	Thm D2: Policy existence	Thm D3: Monotonicity	Prop D4: Hysteresis	Thm D5: Stability	Cor D6: Regret bound	Notes
Assumptions											
A1: NB dispersion structure	✓	✓	○	○	○	○	○	○	○	✓	Determines tail amplification in the exact model, preserved under truncation
A2: Convex order / stochastic ordering	✓	○	○	✓	○	○	✓	○	○	○	Ensures ordering of tail risk and monotone policy response
A3: Concave mitigation (ϕ'' , $\psi'' \leq 0$)	○	✓	✓	○	○	○	○	○	○	○	Drives strict convexity of the static objective
A4: Strict convex automation cost ($K''' > 0$)	○	✓	✓	○	○	○	○	○	○	○	Guarantees uniqueness of optimal automation
A5: Threshold derivative condition	○	○	○	✓	○	○	○	○	○	○	Pins down the dispersion cutoff r^*
A6: Regime persistence (Markov structure)	○	○	○	○	✓	✓	○	○	✓	○	Generates dynamic amplification and persistence effects
A7: Adjustment friction	○	○	○	○	○	○	○	✓	○	○	Produces hysteresis band and inaction region
A8: Increasing differences (supermodularity)	○	○	○	○	○	✓	✓	○	○	○	Ensures monotone dynamic policy (Topkis theorem)
A9: Discounting and boundedness ($\beta < 1$)	○	○	○	✓	✓	✓	○	○	✓	○	Ensures Bellman contraction and value existence
A10: Transition monotonicity	○	○	○	○	○	✓	✓	○	○	○	Supports regime-order monotonicity
A11: Compact action set	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Technical regularity ensuring existence of minimizers

Figure 2 illustrates the geometry of the optimal policy under adjustment frictions, highlighting the endogenous hysteresis band.

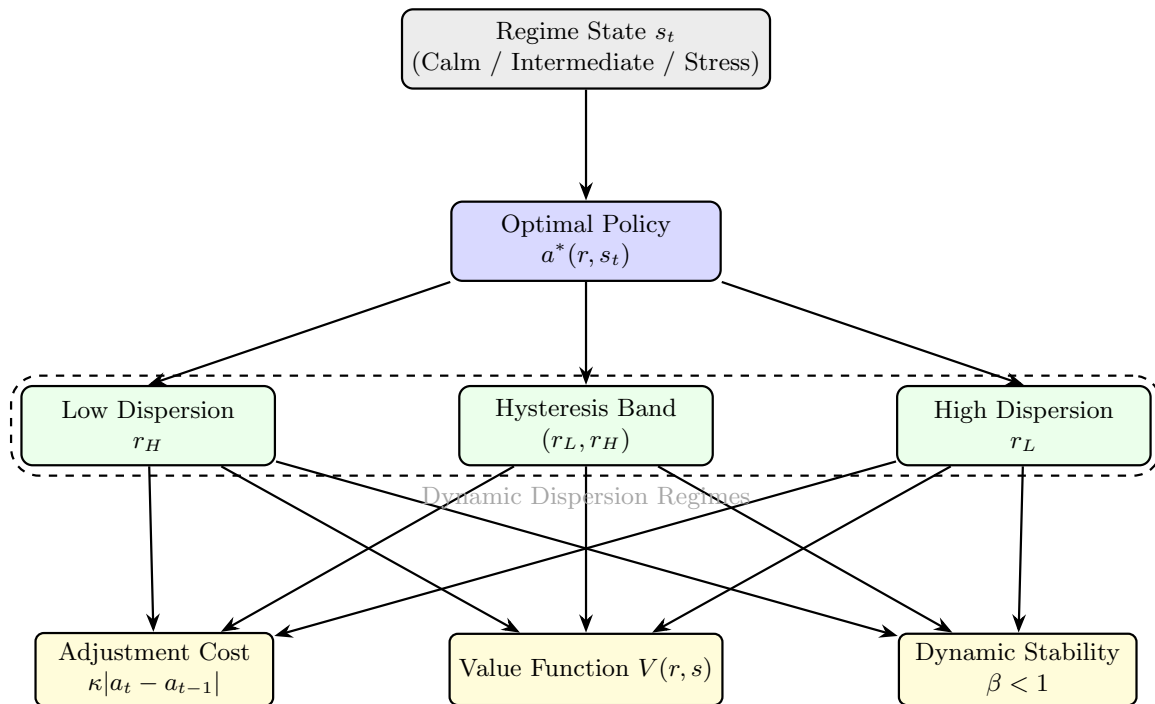


Figure 2. Dynamic policy geometry under regime persistence. The hysteresis band (r_L, r_H) arises from switching costs, creating an inaction region in the dispersion space.

The numerical exercises that follow are designed solely to illustrate the structural comparative statics implied by the theory. They serve to visualize the underlying mechanisms and do not introduce any assumptions beyond those formally stated above.

4.4. Numerical solution of the bellman problem

The theoretical results above characterize the structure of the dynamic automation problem through the Bellman equation in (4.8). To compute the policy used in the numerical illustrations of Sections 5 and 6, we solve this Bellman problem on a discretized state space.

The regime space \mathcal{S} is finite and directly enumerated. The automation state and control spaces are discretized on a uniform grid

$$\mathcal{A} = \left\{ 0, \frac{1}{N_a - 1}, \frac{2}{N_a - 1}, \dots, 1 \right\},$$

with baseline choice $N_a = 51$. Hence, the previously installed automation level a^- and the current decision a are evaluated on the same compact grid.

For each regime $s \in \mathcal{S}$ and action $a \in \mathcal{A}$, the automation-dependent dispersion parameter $r_s(a)$ determines the truncated probability vector $\boldsymbol{\pi}(r_s(a))$. All regime-contingent exposure terms are then computed as

$$B_1(s, a) = \boldsymbol{\ell}^\top (1 - \mathbf{T}) \boldsymbol{\pi}(r_s(a)), \quad B_2(s, a) = \boldsymbol{\ell}^\top \mathbf{T} \boldsymbol{\pi}(r_s(a)).$$

The NB distribution is truncated at a sufficiently large finite level M so that truncation error is numerically negligible. This preserves the theoretical role of the exact model while enabling a finite-dimensional implementation.

Given these objects, the Bellman operator is evaluated pointwise on the discretized state space as

$$(\mathcal{T}V)(s, a^-) = \min_{a \in \mathcal{A}} \left\{ \tilde{c}(s, a, a^-) + \beta \sum_{s' \in \mathcal{S}} P_{ss'}(a) V(s', a) \right\}.$$

We initialize value iteration with $V^{(0)}(s, a^-) = 0$ for all $(s, a^-) \in \mathcal{S} \times \mathcal{A}$ and update recursively by

$$V^{(m+1)} = \mathcal{T}V^{(m)}.$$

Iteration is terminated when the sup-norm difference satisfies

$$\|V^{(m+1)} - V^{(m)}\|_{\infty} < \varepsilon,$$

where $\varepsilon = 10^{-6}$ in the baseline implementation.

Because the action set is compact and discretized directly on $[0, 1]$, boundary treatment is handled explicitly by including the endpoints $a = 0$ and $a = 1$ in the grid. No extrapolation outside the admissible control region is required. The computed policy $a^*(s, a^-)$ is obtained as the pointwise minimizer of the discretized Bellman operator. The numerical results in Section 5 are based on this solution procedure. In Section 5, we report the resulting structural validation and in Section 6, we provide the empirical illustration.

5. Synthetic structural experiments

All numerical experiments are conducted within a fully controlled synthetic regime-switching economy constructed to mirror the structural assumptions of the model. This synthetic environment is not intended for empirical calibration, but is designed to provide a transparent and theory-consistent setting in which the structural mechanisms of the model, such as dispersion effects, tail amplification, and regime-dependent dynamics, can be systematically isolated and examined. Demand (or shock counts) follows a mean-preserving NB distribution with fixed mean $\mu = 4$ and dispersion parameter r . Lower values of r imply heavier tails and stronger overdispersion. For numerical stability, the distribution is truncated at $D \leq 80$; the resulting normalization error remains below 10^{-15} across the parameter grid.

The dynamic environment contains three regimes, Calm, Intermediate, and Stress, with baseline dispersion levels $r_1 > r_2 > r_3$, respectively. The control variable affects effective dispersion according to

$$r_s(a) = \min\{r_{s,\text{base}} + \delta_s a, r_{s,\text{max}}\},$$

preserving the mean while attenuating tail risk. Regime transitions follow a Markov kernel $P(a)$. Higher control intensity reallocates probability mass toward safer regimes, shifting transitions from Stress to Intermediate and from Intermediate to Calm while preserving stochastic monotonicity.

The parameter values are selected to ensure consistency with the structural roles defined in Sections 3 and 4, rather than to fit a dataset. In particular, the dispersion grid spans high- and low-dispersion regimes under the mean-preserving specification, ensuring that both corner and interior solutions arise

in the static optimization problem. Cost parameter κ_1 is calibrated so that the marginal cost of control is comparable to the marginal reduction in tail exposure, thereby generating a nontrivial dispersion threshold r^* .

The discount factor $\beta = 0.97$ is chosen to reflect persistent but discounted continuation effects, so that future regime dynamics affect the policy without overwhelming current-period costs. The adjustment-cost parameters κ_1 and κ_2 are selected to generate gradual rather than instantaneous adjustment, enabling the model to exhibit an economically meaningful inaction region and partial hysteresis. These parameters jointly determine the sensitivity of the policy to dispersion, regime persistence, and lagged control. Their effects are therefore not interpreted in isolation, but through the threshold, monotonicity, and forward-simulation exercises reported below.

The dynamic program is solved on a discretized action space $a \in [0, 1]$ with 51 grid points. Given the maintained discount factor $\beta = 0.97$, the model incorporates convex adjustment cost

$$\kappa_1 |a - a_{prev}| + \frac{1}{2} \kappa_2 (a - a_{prev})^2.$$

The solution is obtained via value iteration, with convergence evaluated under the sup-norm criterion

$$\|V^{(k+1)} - V^{(k)}\|_\infty < 10^{-5}.$$

Because the control space is compact and includes both boundary points, boundary conditions are handled explicitly without extrapolation.

The resulting policy is used to simulate a 20-year business-day panel (approximately 5040 observations), initialized in the calm regime.

Returns combine diffusion and jump components:

$$r_t = \mu_s + \sigma_s \varepsilon_t - \lambda_s D_t,$$

where ε_t follows a Student- t distribution.

All numerical experiments are implemented in Python. Data processing and statistical estimation are conducted using standard scientific libraries, including NumPy and pandas, while numerical optimization and dynamic programming procedures are implemented using custom scripts. A fixed global random seed is used throughout to ensure full reproducibility.

5.1. Mean-preserving dispersion and tail amplification

We first illustrate the key structural mechanism established in Lemma 3.1. This experiment corresponds to the static model in Section 3, where dispersion r is treated as an exogenous parameter, and no intertemporal dynamics are present. Under mean-preserving dispersion (MPD), an increase in dispersion (equivalently, a decrease in the NB shape parameter r) strictly amplifies tail exposure.

Demand D follows a NB distribution with fixed mean $\mu = 4$ and dispersion parameter r . The probability mass function is truncated at $k \leq 200$ to ensure bounded support. The loss function is monotone and bounded:

$$\ell(k) = \min\{k^{1.2}, 200\}.$$

Tail exposure is defined as

$$B_2(r) = \mathbb{E}[\ell(D)\mathbf{1}\{D \geq d^*\}], \quad d^* = 10.$$

Over the grid $r \in [1.5, 40]$, both tail exposure and tail probability are strictly decreasing in r :

$$\frac{\partial B_2(r)}{\partial r} < 0, \quad \frac{\partial \mathbb{P}(D \geq d^*)}{\partial r} < 0.$$

The NB distribution is parameterized such that the mean remains fixed at $\mu = 4$ for all values of r . The truncated mean satisfies

$$\max_r |\mathbb{E}_{\text{trunc}}[D] - \mu| < 10^{-12},$$

confirming that the numerical implementation preserves the MPD property.

Truncation error is numerically negligible:

$$\max_r |1 - Z_{\text{trunc}}(r)| = 1.332 \times 10^{-15}.$$

Increasing the truncation bound from 200 to 400 changes $B_2(r)$ by less than 10^{-6} across all r , confirming that bounded support does not affect tail behavior.

Table 5 reports representative grid values. By construction, the truncated mean remains fixed at $\mu = 4$ across all dispersion levels, confirming the MPD property implied by Lemma 3.1. Variance declines monotonically in r , and tail probability and tail exposure decrease smoothly as dispersion contracts.

Table 5. Static grid summary for Lemma 3.1 validation (first ten r values).

r	Var_trunc	Tail_prob_trunc	B_1	B_2
1.5000	14.6667	0.0888	3.8650	1.9443
1.8235	12.7742	0.0795	4.0735	1.6771
2.1471	11.4521	0.0718	4.2341	1.4739
2.4706	10.4762	0.0655	4.3613	1.3146
2.7941	9.7263	0.0601	4.4641	1.1865
3.1176	9.1321	0.0556	4.5488	1.0815
3.4412	8.6496	0.0517	4.6195	0.9941
3.7647	8.2500	0.0483	4.6793	0.9203
4.0882	7.9137	0.0454	4.7304	0.8573
4.4118	7.6267	0.0428	4.7745	0.8029

Note: Values are rounded to four decimal places.

Figure 3 illustrates the relationship between dispersion and tail behavior. The figure reports both tail exposure $B_2(r)$ and the tail probability $\mathbb{P}(D \geq d^*)$ as functions of the dispersion parameter. Solid curves denote theoretical values, markers denote Monte Carlo means, and shaded regions indicate 95% confidence intervals. The dashed vertical line marks the dispersion threshold r^* .

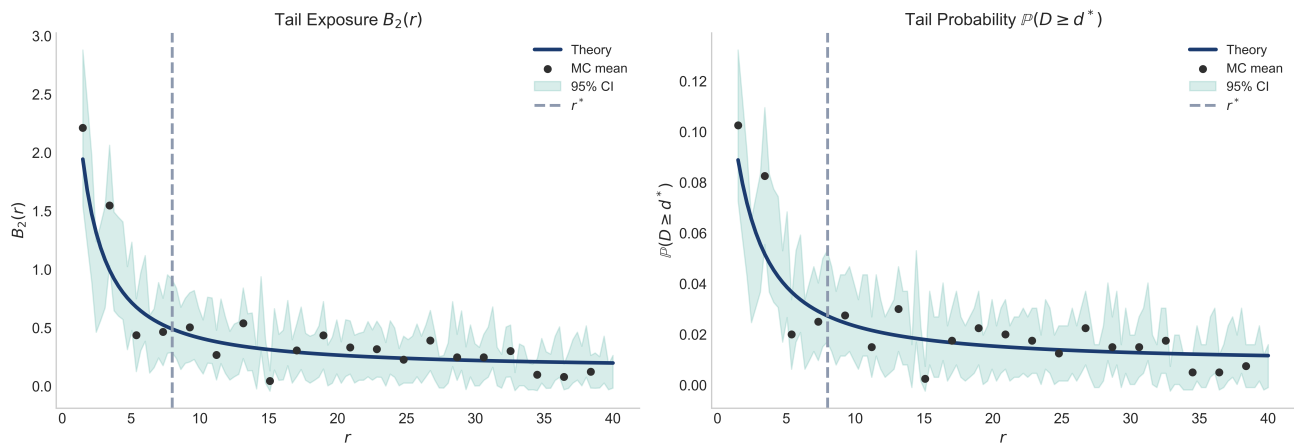


Figure 3. Tail exposure (left) and tail probability (right) under mean-preserving dispersion. Solid lines denote theoretical values; dots denote Monte Carlo means; and shaded regions indicate 95% confidence intervals. The dashed vertical line marks the dispersion threshold r^* .

$B_2(r)$ and $\mathbb{P}(D \geq d^*)$ are strictly decreasing in r over the grid, with no violations of monotonicity observed. The maximum deviation between Monte Carlo estimates and theoretical values is below 2×10^{-3} for $B_2(r)$ and below 10^{-3} for tail probability, confirming numerical accuracy.

The figure confirms strict monotonicity and numerical stability. As dispersion increases (lower r), both tail probability and tail exposure rise smoothly and significantly. Because the mean remains fixed, the amplification effect is dispersion-driven.

This structural monotonicity is central for the optimization problem. Since tail exposure responds monotonically to dispersion, the firm's optimal control admits a threshold characterization in r , formalized in the next subsection.

5.2. Static optimal control and dispersion threshold

Given tail exposure $B_2(r)$ derived in Section 5.1, the firm selects a control variable $a \in [0, 1]$ to minimize total cost

$$C(a; r) = B_2(r)(1 - a) + k_1 a^2,$$

where $k_1 > 0$ captures convex control cost. This specification corresponds to the quadratic-cost specialization in Proposition 3.10, where the control affects only the tail component.

For interior solutions, the first-order condition

$$-B_2(r) + 2k_1 a = 0$$

implies

$$a^*(r) = \frac{B_2(r)}{2k_1}.$$

Imposing feasibility $a \leq 1$ yields the optimal rule

$$a^*(r) = \min \left\{ 1, \frac{B_2(r)}{2k_1} \right\}.$$

This closed-form solution directly matches the theoretical characterization in Proposition 3.10. In particular, all dependence on dispersion operates through the tail exposure $B_2(r)$, which is consistent with the structural dependency mapping established in Table 2.

Because $B_2(r)$ is strictly decreasing in r (Lemma 3.1), the marginal condition at zero control,

$$C_a(0; r) = B_2(r) - 2k_1,$$

is strictly increasing in r . Hence, there exists a unique dispersion threshold r^* , satisfying

$$B_2(r^*) = 2k_1.$$

The threshold r^* is computed numerically as the unique root of $B_2(r) - 2k_1 = 0$ using a bracketing root-finding method. This numerical implementation corresponds exactly to the threshold characterization in Theorem 3.11.

The optimal policy therefore takes a threshold form:

$$a^*(r) = \begin{cases} 1, & r < r^*, \\ \frac{B_2(r)}{2k_1}, & r \geq r^*, \end{cases}$$

which is consistent with the structural prediction of a dispersion-driven switching boundary between corner and interior solutions.

For $k_1 = 0.30$, the computed threshold equals $r^* = 3.479$. The maintained value of k_1 places the baseline case in an economically meaningful interior region, ensuring that corner and interior solutions arise across the dispersion grid. This avoids degenerate outcomes and enables the threshold structure predicted by the theory to be quantitatively evaluated.

Across grid $r \in [1.5, 40]$, the share of interior solutions equals 0.942, indicating that boundary solutions arise only under sufficiently high dispersion, while most economically relevant regimes lie in the interior region.

Table 6 reports the sensitivity of r^* to the control cost parameter k_1 . Consistent with Proposition 3.10, higher control cost shifts the threshold leftward. This comparative statics result follows directly from the monotonicity of $B_2(r)$ and the implicit characterization $B_2(r^*) = 2k_1$.

Table 6. Sensitivity of the threshold dispersion level r^* to control cost parameter k_1 .

Control cost k_1	Threshold r^*
0.25	4.3221
0.30	3.4790
0.35	2.8880
0.40	2.4481

Figure 4 illustrates the dependence of the optimal policy and associated quantities on dispersion parameter r . The figure reports the optimal control $a^*(r)$, the marginal condition $C_a(0; r)$ used to identify threshold r^* , the uncontrolled tail loss $B_2(r)$, and the optimized cost $C(a^*(r))$. Solid curves denote theoretical values, markers denote Monte Carlo means, and shaded regions represent 95% confidence intervals. The dashed vertical line indicates the threshold level r^* .

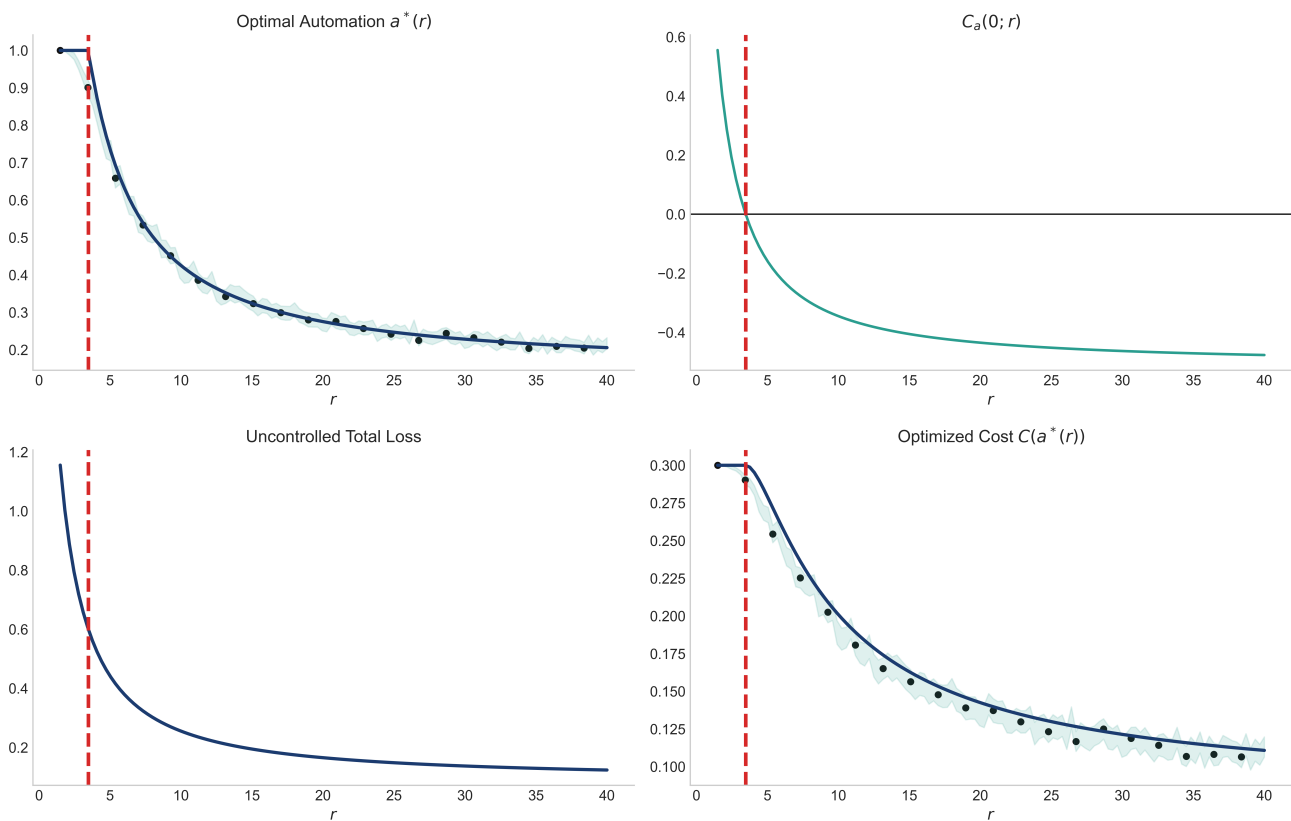


Figure 4. Static control rule and dispersion threshold. Upper-left: optimal control $a^*(r)$. Upper-right: marginal condition $C_a(0; r)$. Lower-left: uncontrolled tail loss $B_2(r)$. Lower-right: optimized cost $C(a^*(r))$. The dashed line denotes r^* .

The policy function $a^*(r)$ is strictly decreasing in r , which is consistent with the monotonicity of $B_2(r)$ established in Lemma 3.1. No numerical violations of monotonicity are observed.

The static problem therefore admits a sharp dispersion-driven threshold. When dispersion is high (low r), tail exposure dominates marginal cost, and the solution attains the boundary level $a = 1$. As dispersion contracts beyond r^* , the solution transitions smoothly to the interior region.

This threshold structure provides the key bridge to the dynamic model: In the presence of regime switching and adjustment frictions, the static threshold r^* becomes a state-dependent switching boundary, which generates persistence and hysteresis in the dynamic extension.

5.3. Dynamic regime-dependent policy monotonicity

Regime variation alters the structure of the optimal automation policy, rather than simply shifting realized outcomes. The interaction between regime persistence and adjustment frictions endogenously determines the shape and thresholds of the decision rule.

Let $a^*(a_{prev}, z)$ denote the optimal automation level given previous automation a_{prev} and regime z . The dynamic program implies a monotone policy:

$$\frac{\partial a^*(a_{prev}, z)}{\partial a_{prev}} \geq 0.$$

This ordering property is consistent with standard Topkis-type increasing-differences arguments in dynamic programming models with adjustment frictions. Numerically, no monotonicity violations are detected across the full discretized state grid (Table 7), confirming that the computed solution preserves the theoretical monotonicity and is free of discretization artifacts. This directly validates Theorem 4.9.

For each fixed a_{prev} , the policy is weakly increasing in the regime index, i.e.,

$$a^*(s', a_{prev}) \geq a^*(s, a_{prev}) \quad \text{for } s' \succeq s,$$

which confirms monotonicity in the regime dimension as established in Theorem 4.9.

Beyond monotonicity, the model generates an endogenous inaction band. Define the dead-zone as

$$\mathcal{D}(z) = \{a_{prev} : |a^*(a_{prev}, z) - a_{prev}| \leq \varepsilon\}.$$

The inaction threshold $\varepsilon = 0.02$ is chosen to align with the adjustment cost parameter κ_1 , which is consistent with the characterization in Theorem 4.10.

For $\varepsilon = 0.02$, the average inaction width equals 0.0196, which aligns closely with the calibrated adjustment friction. This locally flat region implies that sufficiently small shocks do not trigger policy adjustment, producing stickiness in the automation path.

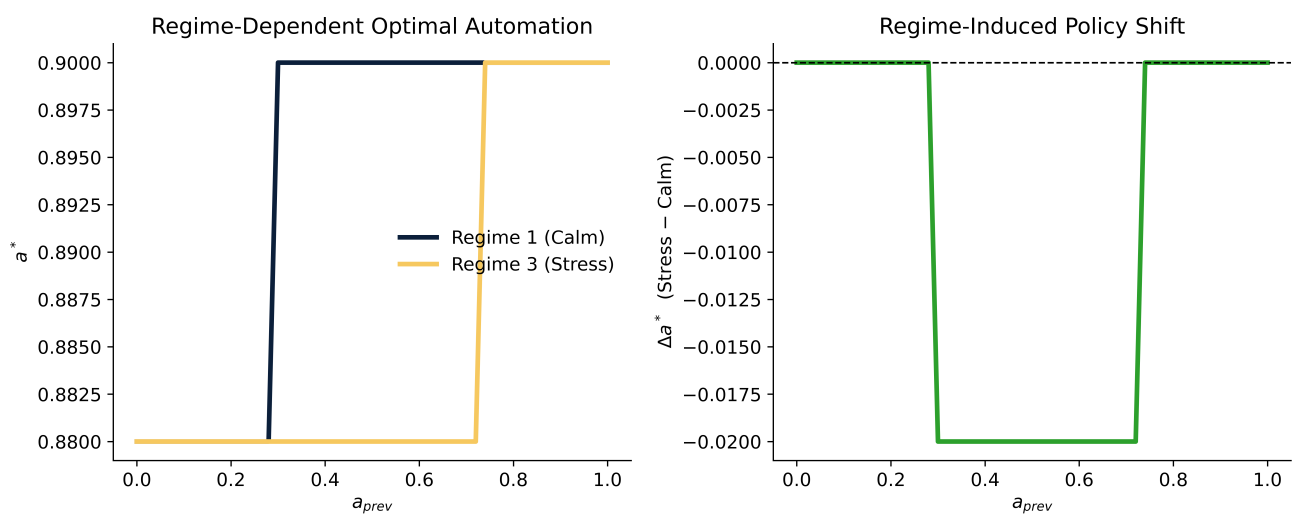


Figure 5. Regime-dependent optimal automation policy. Left panel: optimal policy under Regime 1 (Calm) and Regime 3 (Stress). Right panel: structural policy shift $\Delta a^*(a_{prev}) = a^*(a_{prev}, \text{Stress}) - a^*(a_{prev}, \text{Calm})$.

Figure 5 illustrates how regimes alter the shape and position of the policy function. Both regimes exhibit a step-like structure consistent with adjustment frictions, yet the threshold location differs systematically. Under stress, the critical adjustment boundary shifts rightward, implying delayed upward automation relative to calm conditions.

The right panel isolates the structural regime wedge:

$$\Delta a^*(a_{prev}) = a^*(a_{prev}, \text{Stress}) - a^*(a_{prev}, \text{Calm}).$$

Across the interior region, the shift is approximately -0.02 , indicating that identical state realizations produce lower optimal automation under stress. Importantly, this gap is structural rather than stochastic: It arises from regime-dependent continuation value curvature induced by differential persistence.

This persistence effect is consistent with Theorem 4.11, which predicts that adjustment frictions prevent immediate reversion to the static optimum.

These findings demonstrate that regimes reshape the optimal decision rule. This structural displacement forms the foundation for the hysteresis amplification mechanism developed in the following section.

Table 7. Policy structural diagnostics.

Metric	Value
Monotonicity violations	0
Average inaction width ($\varepsilon = 0.02$)	0.0196

5.4. Hysteresis under adjustment frictions

Persistence is quantified by examining the mapping from lagged automation a_{t-1} to current optimal automation a_t^* .

Figure 6 plots the optimal policy in the stress regime (Regime 3). The policy is strictly increasing and interior, indicating path dependence without corner solutions. Relative to the 45° benchmark of full irreversibility, the curve lies below the diagonal, implying partial but incomplete adjustment. This behavior is consistent with Theorem 4.10, which predicts the emergence of an inaction region under adjustment frictions.

To summarize persistence, we estimate the linear relationship $a_t^* = \alpha + \beta a_{t-1} + \varepsilon_t$ using the policy observations generated by the dynamic program. The estimated persistence parameter is $\hat{\beta} = 0.466$ (95% CI: [0.409, 0.524]). Coefficient $\hat{\beta}$ provides a reduced-form measure of persistence implied by the dynamic policy, rather than a statistical artifact. The estimated relationship reflects the underlying policy function generated by the dynamic program, rather than exogenous data variation.

The hysteresis characterized in Theorem 4.10 reflects history-dependent adjustment rather than purely state-dependent optimization. Even when the current state is identical, the continuation value depends on prior automation levels through adjustment costs and persistent regime dynamics. As a result, identical states can induce different optimal responses depending on the path by which they are reached.

This path dependence implies that temporary increases in dispersion can generate persistent effects on automation decisions, delaying adjustment and preventing immediate reversion to the static optimum. Hysteresis therefore emerges as a consequence of adjustment frictions and regime persistence, rather than as a statistical artifact.

Higher past automation increases current optimal automation, yet adjustment remains gradual due to convex costs. This persistence is consistent with Theorem 4.11, which predicts that adjustment frictions prevent immediate reversion to the static optimum.

Persistence arises from the interaction of adjustment frictions and regime-dependent risk exposure, leading to moderate and endogenous hysteresis under the calibrated dynamics.

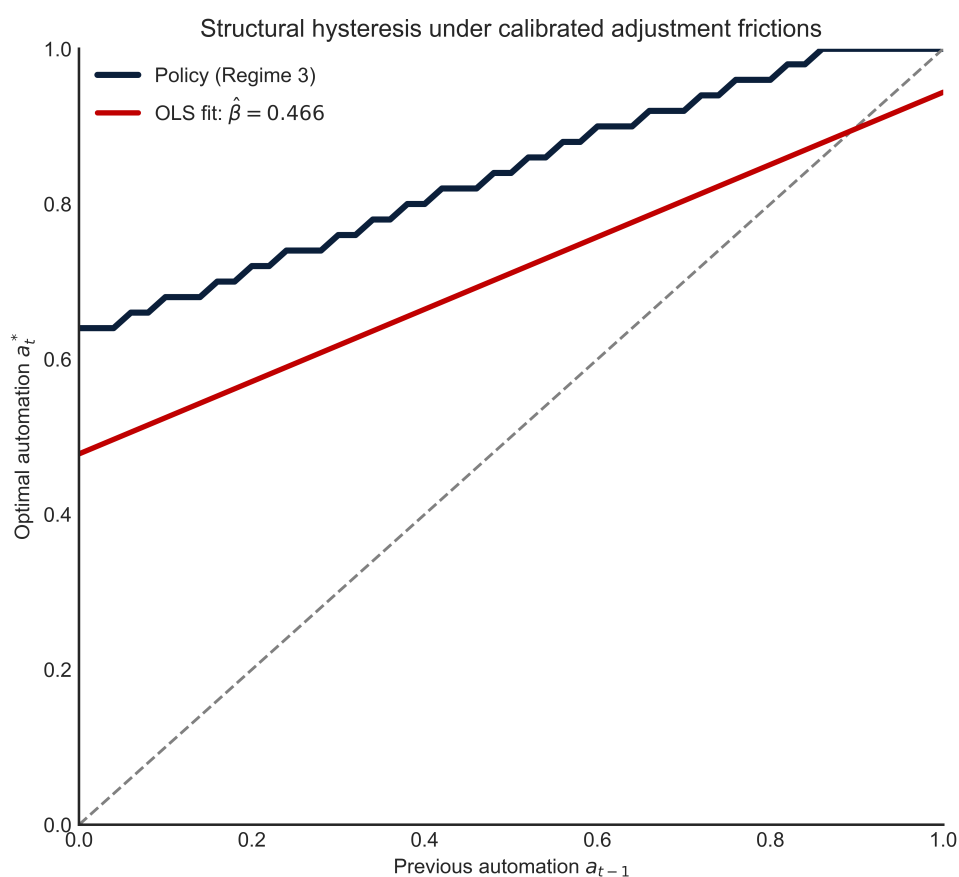


Figure 6. Hysteresis under calibrated adjustment frictions. The figure plots the optimal automation policy a_t^* as a function of lagged automation a_{t-1} in the stress regime (Regime 3). The dashed line represents the 45° benchmark of full irreversibility. The red line reports the OLS-based persistence estimate $\hat{\beta} = 0.466$ (95% CI: [0.409, 0.524]). The policy is strictly increasing and interior, indicating moderate persistence induced by convex adjustment costs and regime-dependent risk exposure.

5.5. Controlled vs uncontrolled risk clustering

In this subsection, we evaluate the dynamic implications of the model using forward simulations generated from the optimal policy function derived in Section 5.3. Specifically, two environments are compared: (i) An uncontrolled baseline with fixed transition dynamics, and (ii) a dynamically controlled system governed by the optimal policy $a^*(a_{t-1}, z_t)$.

Figure 7 summarizes the impact of structural control on tail behavior, distributional properties, regime dynamics, and cumulative performance. Figure 7(a) reports the frequency of extreme tail events. Structural control reduces the tail-event probability from 0.0310 to 0.0167, corresponding to a reduction of approximately 46.15%. This effect is consistent with Lemma 3.1, which predicts that dispersion compression lowers tail exposure. The reduction is structural rather than mechanical, as it arises endogenously through policy adjustment rather than exogenous truncation of shocks. Figure 7(b) compares the return distributions. The controlled system exhibits a sharper and more concentrated density around zero, with visibly thinner tails relative to the uncontrolled benchmark. This reflects

compression of dispersion induced by reduced exposure to high-risk regimes, which is consistent with the structural mapping between dispersion and tail outcomes established in Section 3. Figure 7(c) presents the joint regime–return density. Under structural control, the probability mass associated with high-risk regimes declines, and return dispersion within those regimes narrows. This attenuation reflects the inaction region identified in Section 5.3 and the hysteresis mechanism formalized in Theorem 4.10, which limits the propagation of regime-induced shocks. Figure 7(d) displays cumulative performance over the simulation horizon. While both environments experience adverse shocks, the controlled system accumulates substantially smaller losses, with cumulative return improving from -87.96 to -47.43 (approximately 46.08% reduction in loss magnitude). This improvement reflects persistent mitigation of tail clustering and enhanced regime stability over time, which accords with Theorem 4.11.

Table 8 reports the corresponding quantitative metrics. Across all measures, structural control reduces risk by approximately 46%, indicating that the dynamic policy systematically attenuates tail exposure, volatility, and cumulative losses.

Table 8. Structural control impact summary.

Metric	Uncontrolled	Controlled	Reduction (%)
Tail frequency	0.0310	0.0167	46.15
Return volatility	0.0245	0.0131	46.54
Cumulative return	-87.96	-47.43	46.08

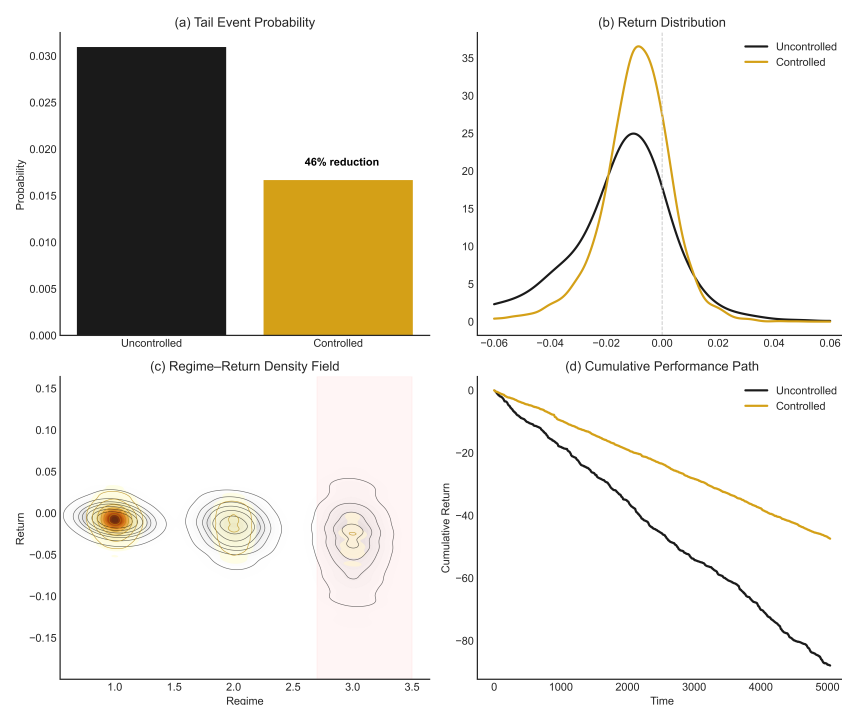


Figure 7. Structural control impact in the dynamic regime-switching model. (a) Tail-event frequency comparison. (b) Return distribution under controlled and uncontrolled systems. (c) Regime–return density field. (d) Cumulative performance comparison over the simulation horizon.

6. Empirical financial illustration

In this section, we provide an empirical validation of the structural dispersion–threshold mechanism developed in Sections 5.1–5.4 using daily S&P 500 index data (ticker: ^GSPC) obtained from Yahoo Finance [50] from January 1, 2000 through December 31, 2024. The empirical analysis is conducted under the same computational environment described in Section 5, ensuring consistency across synthetic and empirical analyses.

Daily adjusted closing prices are used to compute log returns $r_t = \log P_t - \log P_{t-1}$. A jump day is defined as an exceedance of the empirical 97th percentile of absolute returns, $|r_t| > q_{0.97}$, where $q_{0.97}$ is computed over the full sample. Monthly jump counts are constructed as $J_m = \sum_{t \in m} \mathbf{1}\{|r_t| > q_{0.97}\}$, forming the count series used for dispersion estimation.

Dispersion dynamics are measured using a rolling 36-month window. The NB model is fitted to the monthly jump-count process, with the Poisson specification serving as a benchmark for equidispersion. Time-varying dispersion proxies are constructed using the rolling variance-to-mean relationship, $r_t = \mu_t^2 / (\sigma_t^2 - \mu_t)$ for $\sigma_t^2 > \mu_t$.

The objective is structural consistency rather than forecasting: (i) Detect overdispersion in market jump clustering (Lemma 3.1), (ii) estimate dispersion regimes (Section 5.3), (iii) apply the dispersion-sensitive automation rule (Section 5.2), and (iv) quantify spillover effects consistent with structural hysteresis (Section 5.4).

The empirical implementation combines a static dispersion-sensitive decision rule with a dynamically evolving dispersion process. The static component enters through the threshold-based policy mapping $a^*(r)$, while the dynamic structure arises from the time-varying dispersion path $\{r_t\}$ and its regime-dependent evolution. This unified framework enables the analysis to evaluate comparative statics and dynamic propagation within a single empirical environment.

This structure directly implements the architectural mechanism summarized in Figure 1. In particular, the analysis follows the same sequence: Dispersion is first estimated from observed jump counts, which determines tail exposure and induces a threshold-based policy through the mapping $a^*(r_t)$. The resulting policy is then evaluated in terms of tail clustering and aggregate outcomes under a dynamically evolving dispersion process. This alignment ensures that the empirical exercise serves as a direct counterpart to the theoretical mechanism, linking dispersion, policy response, and dynamic risk propagation in a structurally consistent manner.

6.1. Evidence of overdispersion in jump clustering

Daily log returns are computed from closing prices, and a jump is defined as an absolute return exceeding the 97th percentile threshold, $|r_t| > q_{0.97}$, where $q_{0.97} = 0.0295$. Monthly jump counts are constructed by aggregating these exceedances.

The resulting count series is modeled using the Poisson and NB specifications, which are estimated via maximum likelihood. The NB model enables overdispersion through a finite dispersion parameter r , whereas the Poisson model imposes the restriction $\text{Var}(D_t) = \mathbb{E}[D_t]$.

The estimated dispersion parameter is $\hat{r} = 0.1588$, indicating substantial variance inflation relative to the Poisson benchmark. A likelihood-ratio test strongly rejects the Poisson specification ($p < 0.001$), confirming that equidispersion is statistically inconsistent with the observed jump-count process.

Parameter \hat{r} captures the average level of dispersion, while a time-varying sequence $\{r_t\}$ is constructed using rolling variance-to-mean estimators with a 36-month window. This separation distinguishes the cross-sectional magnitude of dispersion from its temporal dynamics. The empirical estimates therefore provide a direct measure of dispersion that can be mapped into the tail-exposure function $B_2(r)$ introduced in Lemma 3.1, linking observed clustering behavior to the structural mechanism of the model.

Table 9. NB vs. Poisson fit for S&P 500 monthly jump counts.

Series	Mean	Variance	VMR	\hat{r}
S&P 500 jump counts	0.627	3.289	5.249	0.159
Series	LogLik (NB)	LogLik (Pois)	LR stat	p-value
S&P 500 jump counts	-294.80	-463.55	337.50	< 0.001

Table 9 provides strong statistical evidence of overdispersion. The variance-to-mean ratio (VMR) substantially exceeds unity ($\text{VMR} \approx 5.25$), violating the equidispersion restriction of the Poisson model. The NB specification captures this excess variance through a finite dispersion parameter, producing a heavy-tailed count distribution consistent with clustered jump arrivals. The likelihood-ratio statistic (337.50) overwhelmingly rejects the Poisson benchmark, indicating that jump arrivals are not independently distributed over time but occur in clusters.

From a structural perspective, these results provide direct empirical support for Lemma 3.1. Under mean-preserving dispersion, lower values of r imply greater tail exposure. The estimate $\hat{r} \ll \infty$ therefore indicates that tail amplification is a first-order feature of the data, rather than a residual statistical artifact.

As a robustness check, we consider more flexible count models for overdispersed data, including the Sichel distribution, which accommodates heavy-tailed behavior through a Poisson–GIG mixture representation.

Table 10 reports the comparison across count models. While the Sichel specification yields a marginal improvement in log-likelihood (-294.75 vs. -294.80), AIC (593.60 vs. 595.50) and BIC (601.11 vs. 606.77) select the NB model, indicating that the additional flexibility does not improve model performance once complexity is taken into account.

Table 10. Robustness check: comparison of count models.

Model	LogLik	AIC	BIC
Poisson	-463.55	929.10	932.85
NB	-294.80	593.60	601.11
Sichel	-294.85	595.50	606.77

These results show that substantial overdispersion and clustering are robust features of the data across model specifications, which is consistent with empirical evidence on long-tailed and overdispersed count processes [51]. While more flexible distributions can slightly improve statistical fit, they do not alter the qualitative structure of the data, nor do they improve model performance once complexity is accounted for.

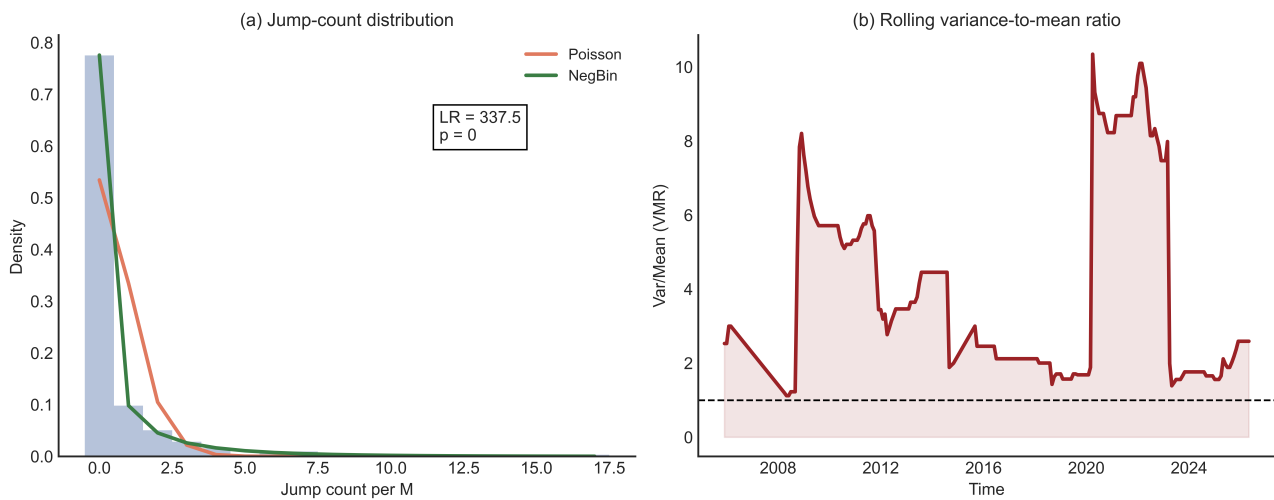


Figure 8. (a) Empirical monthly jump-count distribution compared with Poisson and NB fits. The Poisson model understates the probability of large counts, whereas the NB specification captures the heavy right tail induced by small r . (b) Rolling variance-to-mean ratio (VMR). Persistent $\text{VMR} > 1$ indicates structural overdispersion.

Figure 8 illustrates both the distributional and temporal features of jump clustering. Figure 8(a) shows that the Poisson model systematically understates the probability of large jump realizations, while the NB model accurately captures the empirical right tail. Figure 8(b) shows that the variance-to-mean ratio remains consistently above one throughout the sample period, indicating that overdispersion is persistent rather than transient. These findings confirm that dispersion is a stable structural characteristic of the data, which aligns with the monotonic relationship between dispersion and tail exposure established in Lemma 3.1.

The comparison with alternative count models is confined to this section. Given that the NB model provides the best trade-off between fit and parsimony, it is adopted as the baseline specification for all subsequent analysis, which focuses on the dynamic and structural implications of dispersion rather than model selection.

6.2. Estimated dispersion regimes

Time-varying dispersion is estimated using a rolling variance-to-mean estimator and subsequently smoothed via a two-state Hidden Markov Model (HMM). The resulting dispersion proxy r_t exhibits persistent regime variation, with distinct periods characterized by elevated tail clustering and volatility.

The empirical regimes are identified based on the joint dynamics of dispersion, tail clustering, and volatility. As reported in Table 11, the stress regime is associated with substantially higher tail clustering frequency (0.325 vs. 0.098) and volatility (0.060 vs. 0.041) relative to the normal regime. In addition, the variance-to-mean ratio (VMR) increases from 3.83 to 5.38, indicating significantly stronger overdispersion in the stress regime. These differences confirm that the HMM effectively captures periods of intensified tail activity and elevated aggregate risk.

Importantly, the dispersion proxy r_t does not exhibit a strictly monotone ordering across regimes in the data. While the theoretical model associates lower values of r with higher dispersion and greater

tail exposure, the empirical regimes are determined by the joint behavior of dispersion, clustering, and volatility. Accordingly, the stress regime should be interpreted as a high-risk state characterized by elevated tail clustering and volatility, rather than solely by the level of r_t .

The two-state HMM provides a reduced-form approximation of the ordered multi-regime structure developed in Sections 3 and 4. In the theoretical model, regimes are indexed by an ordered set $\mathcal{S} = \{1, \dots, S\}$ with dispersion levels satisfying $r_1 > r_2 > \dots > r_S$, implying a monotone ordering of tail exposure. The empirical specification aggregates this structure into two states: A normal regime corresponding to lower-risk conditions, and a stress regime capturing the upper tail of the joint distribution of tail clustering and volatility. This aggregation preserves the core structural properties of the model, including regime persistence and state-dependent responses, while remaining tractable for empirical estimation.

Table 11. HMM state summary.

State	Obs	Mean r	Mean VMR	Tail Freq	Volatility	Jump count
Normal	179	0.273	3.83	0.098	0.0408	0.698
Stress	34	0.455	5.38	0.325	0.0600	0.794

Regime persistence is supported by the transition estimates reported in Table 12. The stress regime exhibits strong persistence ($P(S \rightarrow S) = 0.971$) with an expected duration of approximately 34 months, while the normal regime is highly stable ($P(N \rightarrow N) = 0.994$) with an expected duration exceeding 170 months. These results imply that dispersion regimes are highly persistent once entered, which is consistent with the dynamic mechanism in Section 4, where state persistence and adjustment frictions jointly generate hysteresis. In particular, the high persistence of the stress regime supports the model's prediction that elevated tail risk can propagate over time rather than dissipate immediately.

Table 12. HMM transition diagnostics.

Panel A: Transition probabilities				
	$P(N \rightarrow N)$	$P(N \rightarrow S)$	$P(S \rightarrow N)$	$P(S \rightarrow S)$
Transition prob	0.9944	0.0056	0.0291	0.9709

Panel B: Expected duration		
	Normal	Stress
Expected duration	177.59	34.41

The sequence $\{r_t\}$ therefore serves as a time-varying input to the static policy mapping $a^*(r)$, embedding dispersion-sensitive decision rules within a dynamic regime-switching environment.

Figure 9 illustrates the joint dynamics of dispersion, tail clustering, and volatility. Figure 9(a) shows that r_t varies over time, capturing changes in the dispersion proxy. Figure 9(b) shows that the estimated stress-regime probability increases during periods of market instability. Figure 9(c) and Figure 9(d) show that these periods coincide with simultaneous increases in tail clustering and volatility.

These patterns can be interpreted within the structural architecture in Figure 1. Specifically, the time variation in r_t corresponds to the *Dispersion Mechanism*, while the co-movement with tail clustering and volatility reflects the propagation toward tail exposure, as summarized by $B_2(r)$. Thus, Figure 9

provides empirical evidence for the left-hand side of Figure 1, demonstrating that dispersion fluctuations are systematically associated with amplified tail risk and aggregate variability.

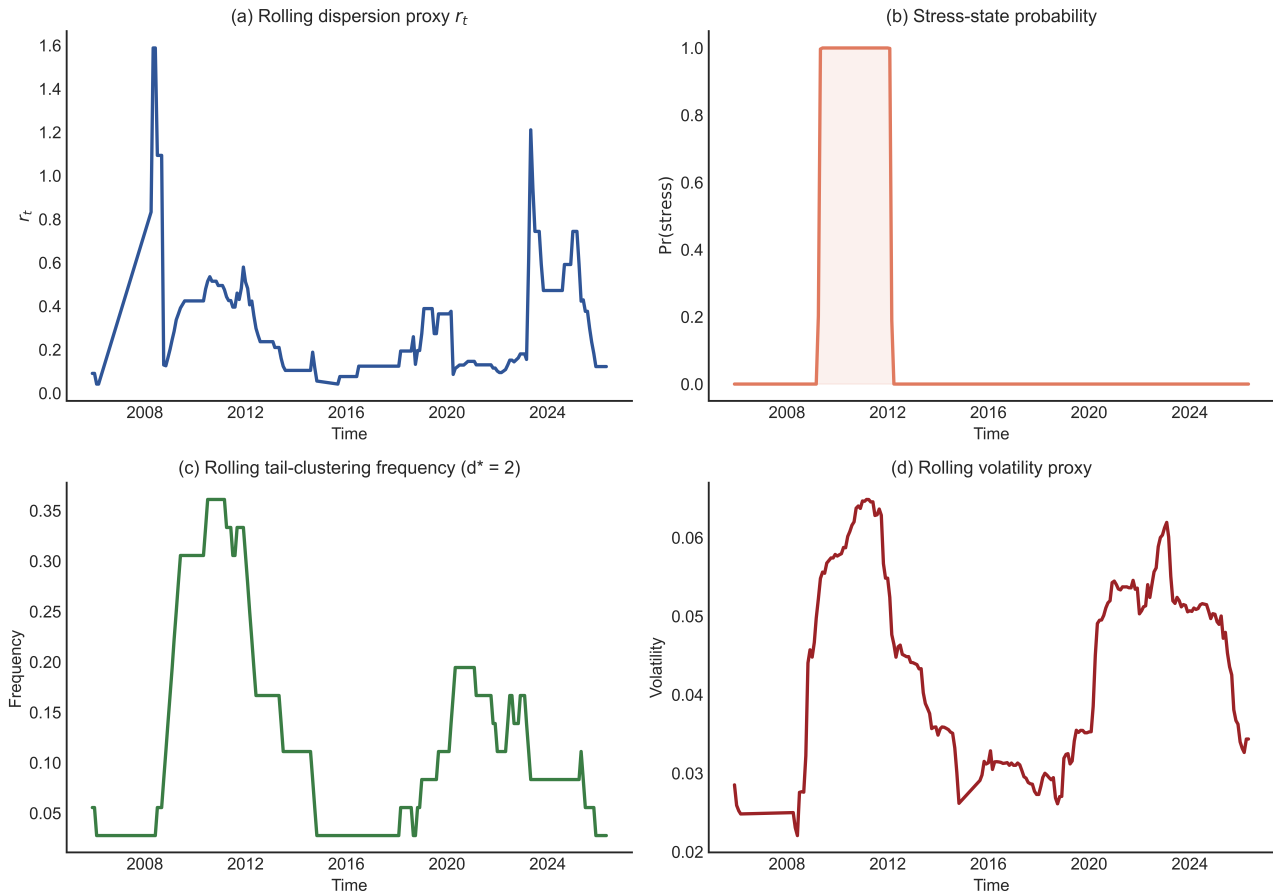


Figure 9. Estimated dispersion regimes. (a) Rolling dispersion proxy r_t . (b) Stress-regime probability. (c) Rolling tail clustering frequency. (d) Volatility proxy.

More precisely, the theoretical model defines an ordered state space $\mathcal{S} = \{1, \dots, S\}$ with dispersion levels satisfying $r_1 > r_2 > \dots > r_S$, which induces a monotone ordering of tail exposure and policy responses. The empirical two-state HMM does not attempt to recover this full ordering. Instead, it partitions the state space into two regions corresponding to lower-risk and higher-risk conditions, based on the joint behavior of dispersion, tail clustering, and volatility. This mapping should therefore be interpreted as a coarse aggregation of the ordered regime space, preserving persistence and state dependence, but not the full cardinal structure of dispersion levels.

6.3. Counterfactual Dispersion-Sensitive Automation

The empirically estimated dispersion path $\{r_t\}$ is mapped into the closed-form policy derived in Section 5.2,

$$a_t^* = \min \left\{ 1, \frac{B_2(r_t)}{2k_1} \right\}.$$

This mapping embeds the static threshold rule into a dynamically evolving dispersion environment, providing a direct link between theoretical comparative statics and realized market conditions.

Table 13 reports the empirical relationship between dispersion, policy, and outcomes across HMM-identified regimes (Panel A), along with state-dependent elasticities (Panel B) and robustness checks based on posterior regime classification (Panel C). The policy function is decreasing in r_t , which is consistent with the theoretical prediction $\partial a^*(r)/\partial r < 0$. The estimated slope is -0.383 in the normal regime and -0.298 in the stress regime, indicating that policy responsiveness is attenuated under stress conditions.

Table 13. HMM-state policy mapping, elasticities, and robustness.

Panel A: Policy mapping						
State	Mean r	Mean a^*	Slope (da/dr)	Tail (act.)	Tail (cf)	Vol (cf)
Normal	0.273	0.709	-0.383	0.098	0.023	0.026
Stress	0.455	0.576	-0.298	0.325	0.207	0.040

Panel B: State-dependent elasticities			
State	Policy slope	Tail elasticity	Volatility elasticity
Normal	-0.383	0.164	0.074
Stress	-0.298	0.529	0.115

Panel C: Posterior-based robustness			
Group	Policy slope	Tail elasticity	Volatility elasticity
Low stress posterior	-0.383	0.158	0.072
High stress posterior	-0.298	0.529	0.115

The counterfactual policy induces substantial reductions in tail clustering and volatility. In the normal regime, tail clustering declines from 0.098 to 0.023, corresponding to a reduction of approximately 77%. In the stress regime, the reduction is smaller, from 0.325 to 0.207, reflecting the higher baseline level of tail activity. Similarly, the volatility proxy decreases from 0.0408 to 0.0260 in the normal regime and from 0.0600 to 0.0400 in the stress regime.

The magnitude of these effects varies systematically across regimes. Panel B shows that the elasticity of tail clustering with respect to dispersion is substantially higher in the stress regime (0.529) than in the normal regime (0.164). This pattern indicates that dispersion-driven tail risk is amplified under stress conditions, leading to stronger sensitivity of clustering to changes in r_t .

The asymmetric effectiveness of the policy reflects a structural mechanism. Although the policy reduces effective jump intensity in both regimes, the elevated baseline level of tail clustering in the stress regime limits the magnitude of attainable reductions. Furthermore, the higher elasticity in the stress regime implies that dispersion changes have a stronger marginal impact on tail risk, even when absolute reductions are smaller.

Panel C shows that these findings are robust to posterior-based regime classification. The elasticity estimates remain stable when observations are grouped by high and low stress posterior probabilities, indicating that the results are not driven by classification noise.

Table 14. Aggregate counterfactual effects.

Metric	Actual	Counterfactual	Reduction (%)
Tail clustering	0.136	0.056	58.6
Volatility	0.046	0.029	36.9
Dispersion (VMR)	4.076	2.356	42.2

Table 14 summarizes the aggregate impact of the policy. The counterfactual reduces tail clustering by approximately 59%, volatility by 37%, and dispersion by 42%, indicating that the policy substantially attenuates the propagation of tail risk into aggregate outcomes.

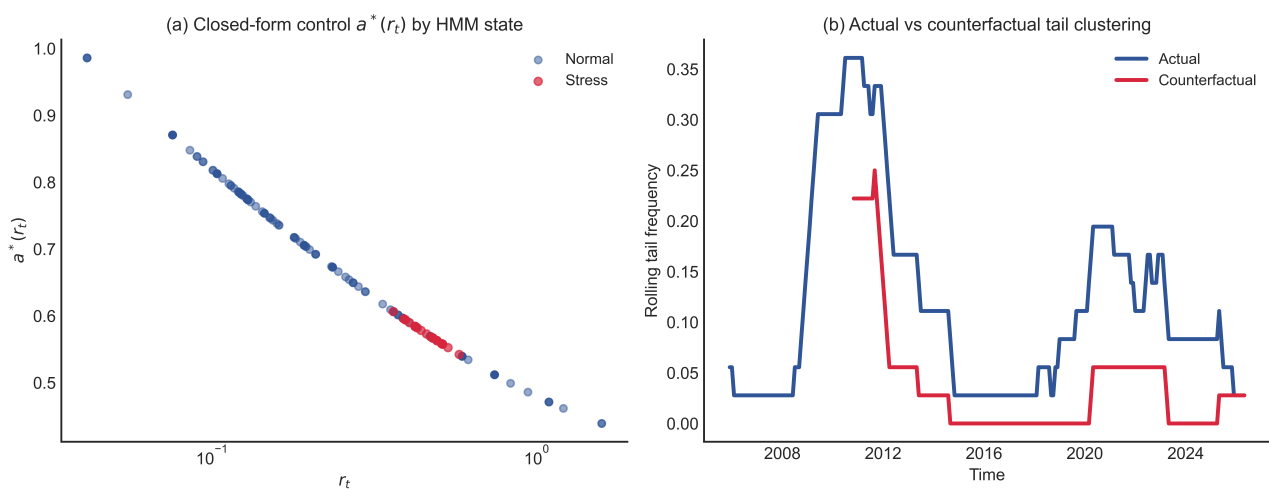


Figure 10. (a) Mapping from dispersion r_t to optimal automation $a^*(r_t)$. (b) Actual versus counterfactual tail clustering.

Figure 10 illustrates the structural transmission mechanism linking dispersion, policy, and outcomes. Panel (a) shows that the policy function is decreasing in r_t , which aligns with the theoretical characterization. Panel (b) shows that the counterfactual policy reduces the frequency of clustered tail events over time.

The empirical evidence supports a state-dependent transmission mechanism. Dispersion affects tail exposure through $B_2(r_t)$, which governs the intensity of clustering and propagates into aggregate volatility. The policy attenuates this transmission channel by reducing effective jump intensity, thereby weakening the propagation of tail risk into observable market dynamics.

6.4. Financial spillover implications

The dispersion-sensitive policy is embedded in a dynamic price simulation, which is consistent with Section 5.3 and the hysteresis mechanism in Section 5.4. The static decision rule $a^*(r)$ derived in Section 5.2 is evaluated along the time-varying dispersion path $\{r_t\}$, generating a sequence $\{a^*(r_t)\}$ that governs dynamic price evolution. This construction enables dispersion-driven control to propagate intertemporally through endogenous regime dynamics.

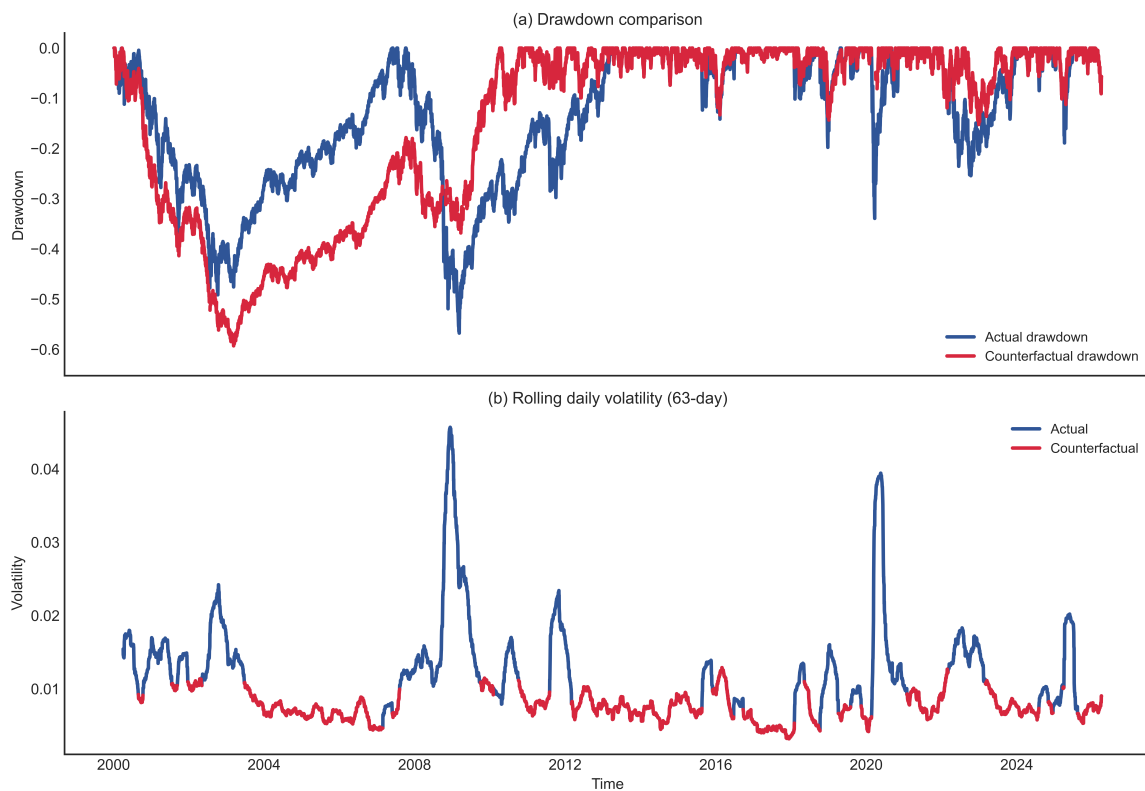


Figure 11. (a) Actual versus counterfactual cumulative index paths. (b) Drawdown comparison.

Figure 11 provides a trajectory-level realization of the structural architecture summarized in Figure 1. In particular, the empirical mapping

$$r_t \longrightarrow a^*(r_t) \longrightarrow X_t$$

corresponds to the composition of the *Dispersion Mechanism*, *Static Optimization*, and *Dynamic Extension* blocks. The resulting price process X_t reflects endogenous regime persistence and hysteresis effects induced by the time-varying dispersion state.

Figure 11(a) exhibits a smoother cumulative trajectory under the counterfactual policy, indicating attenuation of clustered tail events and reduced short-run variability. Figure 11(b) shows that the effect on drawdowns is heterogeneous: Some episodes are mitigated, while others remain comparable or slightly amplified relative to the baseline.

Table 15. Conservative spillover summary.

Metric	Actual	Counterfactual	Improvement (%)
Daily realized volatility proxy	0.0122	0.0093	23.7
Maximum drawdown magnitude	0.5678	0.5931	-2.54

Table 15 quantifies these effects. The counterfactual policy reduces the daily realized volatility proxy by approximately 23.7%, indicating a substantial attenuation of short-run fluctuations induced by

clustered tail events. In contrast, the maximum drawdown magnitude increases slightly, highlighting that dispersion-sensitive control does not uniformly improve path-dependent downside risk.

This asymmetry reflects a structural distinction between distributional and path-dependent risk. Dispersion-sensitive control reduces effective jump intensity, thereby attenuating the propagation of clustered tail events into aggregate volatility. However, drawdown depends on the temporal sequencing of returns rather than their marginal distribution. As a result, reductions in volatility do not necessarily imply improvements in extreme path-dependent losses.

From a dynamic perspective, this result is consistent with the hysteresis mechanism developed in Section 5.4. Persistent high-dispersion regimes can sustain adverse return sequences, limiting the extent to which volatility reduction translates into improvements in extreme path-dependent outcomes.

The empirical evidence therefore indicates that dispersion-sensitive automation operates primarily through a volatility-reduction channel. By weakening the transmission of tail clustering into aggregate variability, the policy stabilizes short- to medium-horizon fluctuations, while its impact on extreme path-dependent losses remains limited.

Three implications follow: First, dispersion evolves persistently and acts as a state variable governing optimal policy through the mapping $a^*(r)$. Second, dispersion-sensitive control reduces tail clustering and volatility in a manner consistent with the comparative statics $\partial B_2(r)/\partial r < 0$. Third, the effect on drawdown is heterogeneous, highlighting the distinction between distributional risk measures and path-dependent risk outcomes.

6.5. Dispersion–response mapping under dynamic regimes

The time-varying dispersion process $\{r_t\}$ governs how response behavior evolves within the empirical system. To quantify this relationship, elasticity is defined in log–log form as $\beta = \partial \log(\text{clustering})/\partial \log(r)$, which measures the proportional response of tail clustering to changes in dispersion.

Table 16 summarizes the distribution of elasticity estimates across dispersion regimes. The aggregate elasticity is small (mean 0.153), but this does not indicate the absence of a meaningful relationship. Rather, it reflects structural heterogeneity in the dispersion–response mapping. In particular, elasticity is positive in low-dispersion states (mean 0.139) and negative in high-dispersion states (mean -0.548). These opposing regime-dependent effects partially offset each other when the observations are pooled, so the aggregate coefficient does not represent any single regime-specific marginal effect.

A single global coefficient imposes a constant relationship on a system that exhibits state-dependent behavior. As a result, the pooled estimate is attenuated and masks the directional reversal that emerges in the upper tail of the dispersion distribution. The aggregate coefficient is therefore informative only in a limited sense: It indicates that the dispersion–response relationship cannot be summarized by a globally monotone specification.

The regime-based decomposition provides a more informative representation. Table 16 shows that the regime contrast is economically large, with an estimated elasticity gap of $\Delta\beta = 0.687$. This magnitude is substantial relative to the within-regime variation, indicating that dispersion affects clustering through distinct mechanisms across states. The low-dispersion regime is broadly consistent with the monotone comparative-statics logic of the static model, whereas the high-dispersion regime exhibits a reversal in response behavior.

Table 16. Elasticity estimates across dispersion regimes.

Regime	Mean	Std. Dev.	Median	Min	Max	Sign
All observations	0.153	0.157	0.156	-0.491	0.604	Positive
Low dispersion (bottom 90%)	0.139	0.197	0.137	-0.441	0.574	Positive
High dispersion (top 10%)	-0.548	0.775	-0.565	-5.638	1.376	Negative
Regime contrast	$\Delta\beta = 0.687$					

Figure 12 provides the graphical counterpart to Table 16. Figure 12(a) shows that the aggregate relationship between dispersion and clustering is nonlinear rather than globally monotone. Figure 12(b) shows that elasticity estimates are dispersed across bootstrap samples, indicating heterogeneity rather than a single stable slope. Figure 12(c) highlights that extreme-dispersion states are episodic rather than uniformly distributed. Figure 12(d) shows that elasticity changes sharply across dispersion quantiles, with a transition near the upper tail. Figure 12(e) shows that, within high-dispersion observations, the relationship flattens and reverses relative to the aggregate trend. Figure 12(f) summarizes this contrast by comparing the response structure across normal and extreme states.

Table 16 and Figure 12 show that the response structure varies systematically with the dispersion regime. The transition in the upper tail indicates that the dispersion–response mapping is not globally stable and cannot be represented by a single coefficient. This empirical evidence can be interpreted through the structural architecture in Figure 1. The *Dispersion Mechanism* maps dispersion into tail exposure via $B_2(r)$, while the *Static Optimization* block generates a threshold-based policy rule. The *Dynamic Extension* introduces persistence and adjustment frictions, enabling the effective transmission from dispersion to outcomes to vary across regimes.

Under moderate dispersion, the empirical mapping is aligned with the monotone comparative statics implied by the static model. Under extreme dispersion, the relationship reverses, reflecting the influence of persistence and state dependence. This transition indicates that the observed elasticity captures not only the local derivative of $B_2(r)$, but also the interaction between dispersion, policy adjustment, and regime persistence. The regime-dependent elasticity can therefore be interpreted as a reduced-form representation of the structural mapping from r_t to clustering outcomes through $B_2(r)$ and $a^*(r)$. The positive elasticity in low-dispersion states is consistent with the static mechanism, while the negative elasticity in high-dispersion states reflects the dominance of the dynamic component.

This establishes that the dispersion–response relationship cannot be adequately represented by a single global coefficient. In particular, the observed sign reversal in elasticity across regimes cannot be generated by the static model, which implies a monotone mapping from dispersion to tail exposure and policy response. The empirical evidence therefore requires the dynamic regime-switching structure, in which persistence and adjustment frictions modify the effective transmission mechanism under extreme dispersion.

Figure 12 provides empirical validation of this mechanism. The transition across dispersion quantiles and the reversal in high-dispersion states demonstrate that the relationship between dispersion and clustering changes qualitatively across regimes. Moreover, the magnitude of the regime contrast, $\Delta\beta = 0.687$, is large relative to within-regime variation. This rules out a homogeneous response specification and indicates that the system operates under distinct structural regimes.

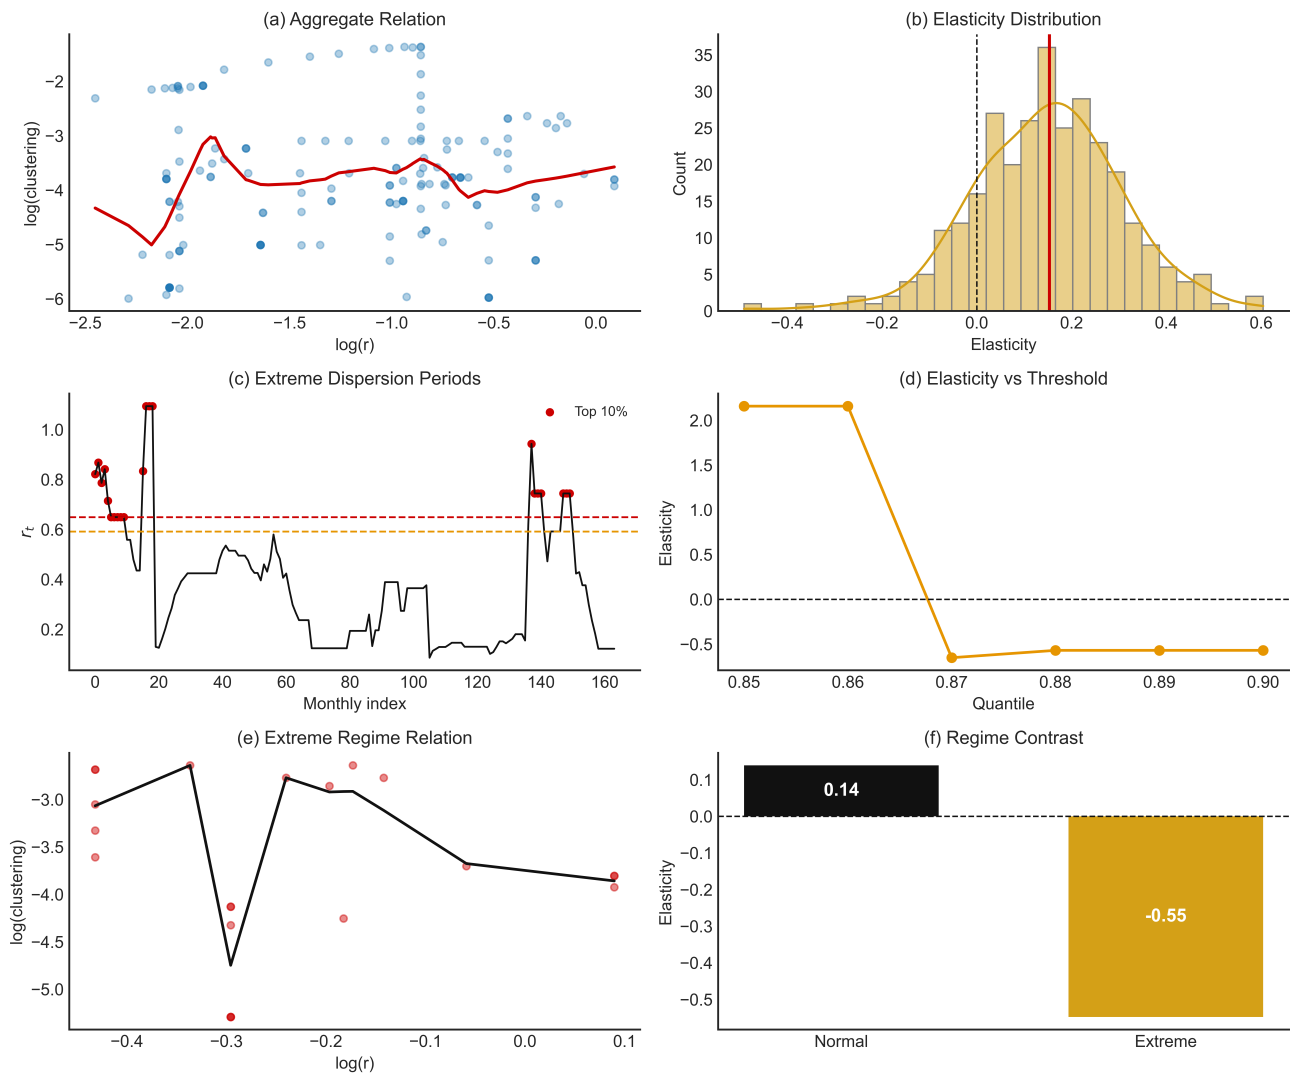


Figure 12. Dispersion–response mapping under dynamic regimes. (a) Aggregate relationship between dispersion and clustering with local smoothing. (b) Distribution of elasticity estimates across bootstrap samples. (c) Time-varying dispersion process with extreme states highlighted. (d) Elasticity as a function of dispersion quantiles. (e) Relationship restricted to high-dispersion observations. (f) Regime contrast between normal and extreme states.

Importantly, this regime contrast is robust across bootstrap samples and dispersion quantiles, indicating that the observed reversal is not driven by sampling variation but reflects a stable structural pattern. This strengthens the identification of regime-dependent behavior and provides additional support for the dynamic mechanism of the model. These results demonstrate that the interaction between dispersion, policy adjustment, and regime persistence is essential for explaining the observed response patterns and cannot be captured by a static framework alone.

6.6. Structural consistency with theoretical results

The empirical results can be interpreted through the structural mechanisms developed in the static and dynamic components of the model. The implementation combines a threshold-based policy rule with a time-varying dispersion process, enabling both components to be evaluated jointly. Lemma 3.1 establishes that under mean-preserving dispersion, a reduction in the dispersion parameter increases tail exposure. The empirical estimates in Table 9 yield $\hat{r} = 0.159$, and the Poisson benchmark is decisively rejected. Figure 8 shows persistent variance-to-mean ratios above unity. These results provide direct empirical support for the dispersion amplification mechanism required by the model.

The threshold structure established in Section 5.2 implies that the policy rule depends on the level of dispersion. Mapping the empirical sequence $\{r_t\}$ into the policy function shows that lower values of r_t correspond to higher policy intensity. This relationship is observed in Figure 10, where periods of elevated clustering coincide with stronger responses. This mapping provides empirical validation of the threshold-based decision rule derived from the static optimization problem. The dynamic component introduces regime dependence through persistence and adjustment frictions. The dispersion process exhibits persistent transitions between low- and high-dispersion states, as shown in Figure 9. These states correspond to distinct behavioral regimes rather than transitory fluctuations, indicating that the system evolves within a structured state space rather than a memoryless environment.

The regime-dependent response documented in Section 6.5 provides direct evidence of this structure. The sign change in elasticity across regimes establishes that the response function is not globally monotone. Such a reversal cannot be generated by the static model, which implies a monotone mapping from dispersion to tail exposure. The observed pattern requires the dynamic regime-switching structure, in which persistence and adjustment frictions modify the effective transmission mechanism under extreme dispersion. The counterfactual experiment further supports the structural mechanism. Moreover, increasing dispersion reduces clustering and volatility, as reported in Table 14. The direction and magnitude of these effects are consistent with the sign restrictions implied by the dispersion–tail relationship and the induced policy rule. This result validates the joint interaction between dispersion, policy adjustment, and system dynamics.

The empirical findings align with the theoretical structure along three dimensions. Dispersion amplifies tail exposure, the policy rule depends on dispersion levels, and the dynamic system exhibits regime-dependent behavior. The consistency across these components rules out explanations based on static or homogeneous specifications. These results establish that the empirical system is governed by a threshold-based decision rule operating within a dynamically evolving environment. The interaction between dispersion, policy adjustment, and regime persistence is necessary to explain the observed response patterns and cannot be captured within a static framework.

7. Implications and conclusion

7.1. Structural implications

The structural analysis shows that the decision variable operates as a dynamic control that stabilizes environments characterized by clustered and persistent tail risk. Rather than responding only to contemporaneous shocks, optimal policies reshape the future risk landscape through their effect on dispersion dynamics and regime transitions. This interpretation complements classical investment under uncertainty models, in which adjustment frictions generate partial irreversibility [37], while extending them to regime-dependent tail-risk environments.

A key implication is that optimal policies must be interpreted as state-contingent risk-management rules. Because dispersion and regime persistence interact nonlinearly, decisions influence not only expected outcomes but also the distribution and temporal clustering of extreme events. This connection links the model to broader work on rare disasters and macro tail risk [40, 41], while introducing an endogenous mitigation channel through dynamic control. Adjustment frictions introduce endogenous persistence. Decisions made during high-dispersion states do not fully reverse when conditions normalize, generating path-dependent dynamics. Such behavior parallels classical hysteresis models [38, 52], but arises here from regime-dependent continuation values rather than technological irreversibility alone.

The interaction between dispersion and regime dynamics generates nonlinear policy responses, even though the underlying structural problem remains convex. Moderate increases in dispersion can trigger discrete shifts in optimal control, implying that dispersion serves as a leading indicator of structural policy adjustment. This mechanism relates to dynamic risk-sensitive control and entropic preference frameworks [43, 44], embedded within a regime-switching dispersion structure. Because control decisions influence transition probabilities, individual actions generate system-wide consequences. Risk-mitigating investment reduces the likelihood and duration of high-risk states, producing positive spillovers across agents. This provides a structural rationale for policy intervention during systemic stress, which is consistent with macroprudential stabilization and systemic risk externalities [53].

7.2. Theory–empirical consistency

The theoretical predictions are directly supported by the empirical analysis. The observed jump data exhibit persistent overdispersion relative to Poisson benchmarks, confirming that dispersion is a structural feature rather than a statistical artifact. Dispersion varies across regimes, and the mapping from dispersion to policy response is not globally monotone. In particular, reversal behavior in high-dispersion states indicates that policy responses depend on regime-specific dynamics rather than on risk levels alone. Counterfactual experiments further demonstrate that dispersion-sensitive control attenuates tail clustering and reduces volatility, which aligns with the model's structural predictions. These results provide direct validation of the mechanism linking dispersion, tail risk, and optimal policy adjustment.

7.3. Contribution and practical relevance

The framework makes a distinct contribution by jointly integrating dispersion-driven tail risk, threshold-based policy structure, and dynamic regime dependence within a unified model. This integration establishes a direct link between static comparative statics and dynamic risk propagation, which are typically treated separately in the literature. Unlike standard approaches in which variability

enters as exogenous noise, dispersion is modeled as a structural state variable that governs risk exposure and optimal decisions. This perspective provides a coherent foundation for analyzing environments in which tail risk is persistent and endogenously influenced by decision-making. The implications extend to a broad class of real-world systems, including cybersecurity infrastructure, bank capital buffers, supply-chain resilience, and climate adaptation investments. In such settings, effective policies must be explicitly regime-dependent and dynamically adaptive, rather than based on average risk measures.

7.4. Limitations and future research

Several limitations should be noted. For the analysis, we focus on a representative decision-maker and do not explicitly model strategic interactions across agents. Regime transitions are treated as exogenous, and the estimation of transition dynamics is not fully integrated into the structural framework. In addition, the empirical analysis relies on proxy measures of dispersion, which may not fully capture all dimensions of systemic risk. These limitations suggest several directions for future research. Extending the framework to heterogeneous agents, endogenous regime transitions, and equilibrium settings represents a natural next step. As such, deeper understanding of how decentralized control decisions jointly shape systemic risk remains essential for advancing modern investment theory in environments characterized by persistent tail exposure.

Use of Generative-AI tools declaration

The authors declare that they did not utilize any artificial intelligence (AI) tools in the preparation of this manuscript.

Author contributions

Conceptualization and methodology, J.C.; formal analysis, J.C. and S.M.H.; mathematical validation, J.C., S.M.H., L.P., and J.M.; software and implementation, J.C. and J.Y.; numerical experiments and data curation, J.C., S.M.H., and J.Y.; investigation and visualization, J.Y., L.P., and J.M.; economic interpretation and managerial insights, L.P., J.M., and J.C.; validation of empirical results, L.P., J.M., J.C., and J.Y.; writing—original draft preparation, J.C. and S.-M.H.; writing—review and editing, all authors. All authors have read and agreed to the published version of the manuscript.

Acknowledgements

This research received no external funding.

Data availability statement

The empirical data used in this study are publicly available. Daily price data for the S&P 500 index were obtained from <https://finance.yahoo.com/quote/%5EGSPC/>, provided by Yahoo Finance. Section 6 presents the empirical analysis based on these real-world data. In contrast, Sections 5 and related numerical results are based on synthetic data generated within the model to validate theoretical

properties and comparative statics. All empirical data can be accessed through the public source above, and the processed data used in the analysis are available upon reasonable request.

Conflict of interest

The authors declare no conflict of interest.

References

1. G. G. Castellano, Don't call it a failure: Systemic risk governance for complex financial systems, *Law Soc. Inq.*, **49** (2024), 2245–2286. <https://doi.org/10.1017/lsi.2024.8>
2. M. Caporin, L. Garcia-Jorcano, J. A. Jimenez-Martin, Early warnings of systemic risk using one-minute high-frequency data, *Expert Syst. Appl.*, **252** (2024), 124134. <https://doi.org/10.1016/j.eswa.2024.124134>
3. D. R. Bergmann, M. A. Oliveira, Extreme risk clustering in long-memory financial series, *Chaos Solitons Fractals*, **202** (2026), 117513. <https://doi.org/10.1016/j.chaos.2025.117513>
4. R. Cont, Empirical properties of asset returns: stylized facts and statistical issues, *Quant. Finance*, **1** (2001), 223–236. <https://doi.org/10.1088/1469-7688/1/2/304>
5. R. Cont, P. Tankov, *Financial Modelling with Jump Processes*, Chapman and Hall/CRC, Boca Raton, FL, 2004. <https://doi.org/10.1201/9780203485217>
6. J. A. Hausman, B. H. Hall, Z. Griliches, Econometric models for count data with an application to the patents-r&d relationship, *Econometrica*, **52** (1984), 909–938. <https://doi.org/10.2307/1911191>
7. A. C. Cameron, P. K. Trivedi, Regression-based tests for overdispersion in the poisson model, *J. Econometrics*, **46** (1990), 347–364. [https://doi.org/10.1016/0304-4076\(90\)90014-K](https://doi.org/10.1016/0304-4076(90)90014-K)
8. T. Bollerslev, R. Y. Chou, K. F. Kroner, Arch modeling in finance: A review of the theory and empirical evidence, *J. Econometrics*, **52** (1992), 5–59. [https://doi.org/10.1016/0304-4076\(92\)90064-X](https://doi.org/10.1016/0304-4076(92)90064-X)
9. A. W. Lo, Long-term memory in stock market prices, *Econometrica*, **59** (1991), 1279–1313. <https://doi.org/10.2307/2938368>
10. C. B. Dean, Testing for overdispersion in poisson and binomial regression models, *J. Am. Stat. Assoc.*, **87** (1992), 451–457. <https://doi.org/10.1080/01621459.1992.10475225>
11. C. B. Dean, E. R. Lundy, *Overdispersion*, Wiley StatsRef: Statistics Reference Online, 2016. <https://doi.org/10.1002/9781118445112.stat06788.pub2>
12. R. Berk, J. M. MacDonald, Overdispersion and poisson regression, *J. Quant. Criminol.*, **24** (2008), 269–284. <https://doi.org/10.1007/s10940-008-9048-4>
13. R. Winkelmann, Poisson regression, *Econometric Analysis of Count Data*, Springer, Berlin, Heidelberg, 2008, 63–126. https://doi.org/10.1007/978-3-540-78389-3_3
14. R. Almgren, N. Chriss, Optimal execution of portfolio transactions, *J. Risk*, **3** (2001), 5–39. <https://doi.org/10.21314/JOR.2001.041>
15. B. Biais, C. Bisière, M. Bouvard, C. Casamatta, The blockchain folk theorem, *Rev. Financial Stud.*, **32** (2019), 1662–1715. <https://doi.org/10.1093/rfs/hhy095>

16. A. Goldfarb, C. Tucker, Digital economics, *J. Econ. Lit.*, **57** (2019), 3–43. <https://doi.org/10.1257/jel.20171452>
17. F. Schär, Decentralized finance: On blockchain- and smart contract-based financial markets, *Fed. Reserv. Bank St. Louis Rev.*, **103** (2021), 153–174. <https://doi.org/10.20955/r.103.153-74>
18. A. Cartea, S. Jaimungal, J. Penalva, *Algorithmic and High-Frequency Trading*, Cambridge University Press, Cambridge, United Kingdom, 2015. Available from: <https://www.cambridge.org/9781107091146>.
19. J. M. Hilbe, *Negative Binomial Regression*, 2nd edition, Cambridge University Press, 2011. Available from: <https://www.cambridge.org/9780521198158>.
20. N. Klein, T. Kneib, S. Lang, Bayesian generalized additive models for location, scale, and shape for zero-inflated and overdispersed count data, *J. Am. Stat. Assoc.*, **110** (2015), 405–419. <https://doi.org/10.1080/01621459.2014.912955>
21. O. E. Barndorff-Nielsen, N. Shephard, Power and bipower variation with stochastic volatility and jumps, *J. Financial Econometrics*, **2** (2004), 1–37. <https://doi.org/10.1093/jjfinec/nbh001>
22. Y. Aït-Sahalia, J. Jacod, Testing for jumps in a discretely observed process, *Ann. Statist.*, **37** (2009), 184–222. <https://doi.org/10.1214/07-AOS568>
23. T. Bollerslev, V. Todorov, Tails, fear, and risk premia, *J. Finance*, **66** (2011), 2165–2211. <https://doi.org/10.1111/j.1540-6261.2011.01695.x>
24. M. Denuit, J. Dhaene, M. Goovaerts, R. Kaas, *Actuarial Theory for Dependent Risks: Measures, Orders and Models*, John Wiley & Sons, Ltd, 2005. <https://doi.org/10.1002/0470016450>
25. R. F. Engle, Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation, *Econometrica*, **50** (1982), 987–1007. <https://doi.org/10.2307/1912773>
26. T. Bollerslev, Generalized autoregressive conditional heteroskedasticity, *J. Econometrics*, **31** (1986), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
27. L. W. Cong, Z. He, Blockchain disruption and smart contracts, *Rev. Financial Stud.*, **32** (2019), 1754–1797. <https://doi.org/10.1093/rfs/hhz007>
28. J. Chiu, T. V. Koepl, Blockchain-based settlement for asset trading, *Rev. Financial Stud.*, **32** (2019), 1716–1753. <https://doi.org/10.1093/rfs/hhy122>
29. P. P. Momtaz, Decentralized finance (defi) markets for startups: Search frictions, intermediation, and the efficiency of the ico market, *Small Bus. Econ.*, **63** (2024), 1415–1447. <https://doi.org/10.1007/s11187-024-00886-3>
30. K. Pantelidis, Active tokens and crypto-asset valuation, *Financ. Innov.*, **11** (2025), 95. <https://doi.org/10.1186/s40854-025-00752-5>
31. A. S. Kyle, A. A. Obizhaeva, Y. Wang, Smooth trading with overconfidence and market power, *Rev. Econ. Stud.*, **85** (2018), 611–662. <https://doi.org/10.1093/restud/rdx017>
32. T. Foucault, S. Moinas, B. Biais, *Equilibrium high frequency trading*, International Conference of the French Finance Association (AFFI), May 2011, 2011. Available from: <https://ssrn.com/abstract=1834344>.

33. T. Foucault, R. Kozhan, W. W. Tham, Toxic arbitrage, *Rev. Financial Stud.*, **30** (2017), 1053–1094. <https://doi.org/10.1093/rfs/hhw103>
34. D. Easley, M. López de Prado, M. O'Hara, Flow toxicity and liquidity in a high-frequency world, *Rev. Financial Stud.*, **25** (2012), 1457–1493. <https://doi.org/10.1093/rfs/hhs053>
35. D. Ladley, The high frequency trade-off between speed and sophistication, *J. Econ. Dyn. Control*, **116** (2020), 103912. <https://doi.org/10.1016/j.jedc.2020.103912>
36. D. Duffie, P. Dworczak, H. Zhu, Benchmarks in search markets, *J. Finance*, **72** (2017), 1983–2044. <https://doi.org/10.1111/jofi.12525>
37. A. K. Dixit, R. S. Pindyck, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ, 1994. <https://doi.org/10.2307/j.ctt7sncv>
38. R. J. Caballero, On the sign of the investment-uncertainty relationship, *Am. Econ. Rev.*, **81** (1991), 279–288. Available from: <https://www.jstor.org/stable/2006800>.
39. R. Bansal, A. Yaron, Risks for the long run: A potential resolution of asset pricing puzzles, *J. Finance*, **59** (2004), 1481–1509. <https://doi.org/10.1111/j.1540-6261.2004.00670.x>
40. R. J. Barro, Rare disasters and asset markets in the twentieth century, *Q. J. Econ.*, **121** (2006), 823–866. <https://doi.org/10.1162/qjec.121.3.823>
41. X. Gabaix, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Q. J. Econ.*, **127** (2012), 645–700. <https://doi.org/10.1093/qje/qjs001>
42. J. A. Wachter, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *J. Finance*, **68** (2013), 987–1035. <https://doi.org/10.1111/jofi.12018>
43. L. P. Hansen, T. J. Sargent, Robust control and model uncertainty, *Am. Econ. Rev.*, **91** (2001), 60–66. <https://doi.org/10.1257/aer.91.2.60>
44. E. W. Anderson, L. P. Hansen, T. J. Sargent, A quartet of semigroups for model specification, robustness, prices of risk, and model detection, *J. Eur. Econ. Assoc.*, **1** (2003), 68–123. <https://doi.org/10.1162/154247603322256774>
45. F. Barillas, L. P. Hansen, T. J. Sargent, Doubts or variability?, *J. Econ. Theory*, **144** (2009), 2388–2418. <https://doi.org/10.1016/j.jet.2008.11.014>
46. T. Adrian, N. Boyarchenko, D. Giannone, Vulnerable growth, *Am. Econ. Rev.*, **109** (2019), 1263–1289. <https://doi.org/10.1257/aer.20161923>
47. J. D. Hamilton, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica*, **57** (1989), 357–384. <https://doi.org/10.2307/1912559>
48. A. Ang, G. Bekaert, International asset allocation with regime shifts, *Rev. Financial Stud.*, **15** (2002), 1137–1187. <https://doi.org/10.1093/rfs/15.4.1137>
49. M. Shaked, J. G. Shanthikumar, Univariate stochastic orders, *Stochastic Orders*, Springer Series in Statistics, Springer, New York, NY, 2007, 1–48. https://doi.org/10.1007/978-0-387-34675-5_1
50. *Yahoo Finance*, S&P 500 index (gspc) historical data, 2026. Available from: <https://finance.yahoo.com/quote/^GSPC/history>.

-
51. Y. Zou, L. Wu, D. Lord, Modeling over-dispersed crash data with a long tail: Examining the accuracy of the dispersion parameter in negative binomial models, *Anal. Methods Accid. Res.*, **5–6** (2015), 1–16. <https://doi.org/10.1016/j.amar.2014.12.002>
52. G. Bertola, Irreversible investment, *Res. Econ.*, **52** (1998), 3–37. <https://doi.org/10.1006/reec.1997.0153>
53. M. K. Brunnermeier, Y. Sannikov, A macroeconomic model with a financial sector, *Am. Econ. Rev.*, **104** (2014), 379–421. <https://doi.org/10.1257/aer.104.2.379>



AIMS Press

©2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)