A Blessing in Disguise: How DeFi Hacks Trigger Unintended Liquidity Injections into US Money Markets*

Tingyi Lin† Karl Yu‡ Huanxi Zhang§

Abstract

Do vulnerabilities in Decentralized Finance (DeFi) destabilize traditional short-term funding markets? While the prevailing "Contagion Hypothesis" posits that the liquidation of stablecoin reserves triggers fire-sale spirals that transmit distress to traditional markets , we document a robust "Flight-to-Quality" effect to the contrary. In the wake of major DeFi exploits, spreads on 3-month AA-rated commercial paper (CP) exhibit a paradoxical narrowing. We identify a "liquidity recycling" mechanism driving this outcome: capital fleeing DeFi protocols is re-intermediated into the traditional financial system via Prime Money Market Funds (MMFs) , where strict regulatory constraints (e.g., SEC Rule 2a-7) compel these funds to purchase high-quality paper. Our estimates indicate that this institutional demand shock quantitatively overwhelms the supply shock driven by stablecoin issuer redemptions. Rather than acting as vectors of financial contagion , these crypto native shocks serve as an inadvertent "safety valve" in segmented markets , providing transient liquidity support and effectively subsidizing borrowing costs for high-grade issuers in the real economy.

JEL Classification: G12, G15, G18, E44
Keywords: Decentralized Finance (DeFi); Financial Contagion; Flight-to-Quality; Knightian Uncertainty

*We are deeply grateful to Professor Andrei Shleifer (Harvard University) for his invaluable guidance and constructive criticism. We have significantly revised the manuscript to incorporate his specific advice on interpreting the economic magnitude of our estimates. By refining the identification strategy and addressing potential endogeneity, we have more robust findings that demonstrate the impact of crypto-related shocks on more segmented money markets. We also thank him for his strategic advice on journal submission and future research directions.

We extend our gratitude to Professor Wanlin Lin and Ruoran Lai (Sun Yat-sen University) for their stimulating academic discussions and insightful feedback, which significantly benefited the development of our theoretical arguments.

Special thanks go to Dr. Ke Dong and the faculty members at the PBC School of Finance, Tsinghua University (Tsinghua PBCSF) for their time and effort in conducting a preliminary review of the manuscript and providing helpful comments. All remaining errors are our own.

†Central University of Finance and Economics. Email: hayashi@email.cufe.edu.cn
‡Boston College. Email: yulin@bc.edu
§University of Wisconsin-Madison. Email: hzhang2488@wisc.edu

1 Introduction

With the deepening integration of Decentralized Finance (DeFi) and traditional financial markets, potential systemic risks emerging from this nexus have become a focal point for global regulators. Institutions such as the Financial Stability Board (FSB, 2022; 2023), the Bank for International Settlements (BIS, 2021), and the International Monetary Fund (IMF, 2024) have long expressed profound concerns regarding the "Contagion Hypothesis." Scholars have further underscored these vulnerabilities, emphasizing that the opacity and lack of regulation in crypto-assets pose severe threats to the broader financial system (Arner et al., 2020; Hermans et al., 2022; MacDonald and Zhao, 2022; Wilmarth, 2025). This prevailing view posits that the DeFi ecosystem is inherently fragile; a collapse triggered by hacks or algorithmic failures could force stablecoin issuers to liquidate their massive reserve assets—primarily U.S. Treasury bills and commercial paper—to meet a surge in redemptions. According to the classic financial contagion theory of Allen and Gale (2000), such a "fire-sale spiral" would create an excess supply in the market, depressing asset prices and sharply widening spreads, thereby transmitting instability directly to the core banking system and undermining global financial stability.

However, using high-frequency data surrounding the top 50 major DeFi exploits between 2021 and 2024, we document an **empirical pattern** that stands in contrast to this regulatory consensus. Whenever DeFi suffers a major security shock, the core of the traditional short-term funding market—the 3-month AA-rated commercial paper (CP) market—does not collapse due to contagion. Instead, financing costs exhibit a significant decline (i.e., spread narrowing). This finding challenges the traditional contagion narrative, suggesting that under specific conditions, turmoil in crypto markets may provide unexpected liquidity support to high-quality traditional borrowers. We term this counter-intuitive phenomenon the "Blessing in Disguise" effect.

To **quantitatively examine** this transmission mechanism, we employ a Sparse Granular Instrumental Variable (GIV) framework (Gabaix and Koijen, 2024). By isolating idiosyncratic supply-side shocks—specifically, technical hacks of large DeFi protocols—we **estimate the implied elasticity** of traditional funding markets. Our two-stage

least squares (2SLS) estimates **suggest** a 'Repatriation Multiplier': for every $100 million of exogenous capital flight from DeFi, the spread on 3-month AA-rated commercial paper contracts narrows by approximately 2.73 basis points. This result is robust to controlling for high-frequency macro factors (e.g., VIX, TED spread), confirming that the observed pricing anomaly is **associated with** crypto-native shocks rather than coincidental shifts in global risk sentiment. We further interpret these magnitudes through the lens of a structural asset pricing model with Knightian uncertainty. Decomposing our reduced-form estimates, we recover a 'Panic Multiplier' (η) of approximately 3.73. This implies that for every dollar of physical capital redeeming from DeFi, institutional investors generate roughly $3.73 of excess demand for safe traditional assets. This massive behavioral over-reaction ($\eta \gg 1$) **rationalizes** how relatively modest outflows from the crypto ecosystem can exert outsized pricing pressure on the trillion-dollar money market, effectively serving as a 'marginal whale' in the segmented market for high-quality collateral.

To understand the micro-foundations of this anomaly, we **examine** the roots of DeFi risk and its transmission mechanisms to traditional finance. Werner et al. (2021) point out that while DeFi's "composability" enhances capital efficiency, it also leads to a systemic stacking of risks. We extend this literature by introducing an availability function, formally modeling abstract code vulnerabilities as payment friction costs. Specifically, research by Belchior et al. (2021) and Chainalysis (2022) emphasizes that cross-chain bridges—which hold vast assets yet remain security weak points—act as centralized single points of failure. **A potential conduit** connecting this risk to traditional markets is the stablecoin, widely regarded in academia as a new form of shadow banking (Gorton and Zhang, 2021; Liao and Caramichael, 2022).

Based on this institutional background, we **explore** the "Liquidity Recycling" hypothesis to explain the "Blessing in Disguise" phenomenon. The core logic stems from the "Flight-to-Quality" effect in asset pricing. Theories by Caballero and Krishnamurthy (2008) and Vayanos (2004) suggest that when risky assets suffer idiosyncratic shocks, highly risk-averse capital swiftly exits and flees to safe assets. We **posit** that hacks inject

Knightian uncertainty into the market, triggering a massive capital flight from the crypto ecosystem. This retreating capital does not vanish; rather, it is **likely** re-intermediated into the system via strictly regulated Prime Money Market Funds (MMFs). Constrained by regulations such as SEC Rule 2a-7[1], MMFs receiving these massive inflows are compelled to purchase high-rated commercial paper. If this institutionally driven demand shock quantitatively overwhelms the supply shock from stablecoin issuer redemptions—represented in our model as the parameter $\eta > 1$—the net effect will be a narrowing, rather than a widening, of CP spreads. Theoretically, we deviate from standard risk models by strictly modeling this mechanism under Knightian Uncertainty. In Appendix B, we solve a robust control problem to show how 'unknown unknowns' generate a structural wedge between physical capital liquidation and psychological flight-to-quality. Furthermore, we employ a Global Game framework in Appendix C to provide micro-foundations for the non-linear congestion threshold observed in our data.

While regulators (e.g., FSB, BIS) implicitly assume the former dominates ($\eta < 1$), leading to a contagion spiral, our empirical analysis **points to** a surprising result: the latter effect is overwhelmingly dominant ($\eta > 1$). This finding resolves the theoretical ambiguity and challenges the existing regulatory narrative.

We employ a Stacked Dynamic Event Study to rigorously test this hypothesis. The empirical results **indicate** that liquidity runs triggered by DeFi hacks do not act as disruptors of financial markets but instead play the role of liquidity injectors for international financial markets. This finding validates the logic of Griffin and Shams (2020) regarding the impact of stablecoin issuance mechanisms on asset prices and **provides evidence consistent with** the dominance of the substitution effect in the context of exogenous hacking shocks. This study fills a gap in the existing literature: on the one hand, it breaks through the limitation of previous studies that focused solely on internal crypto market volatility by directly **examining the transmission link** from DeFi to

[1]Specifically, Section (d)(2)&(a)(11)'s restriction to 'Eligible Securities' and Section (d)(4)'s rigid daily/weekly liquidity requirements act as a regulatory conduit. When DeFi exploits trigger Knightian uncertainty, fleeing capital is **functionally** re-intermediated into Prime MMFs, which are legally compelled to deploy these inflows into high-grade, short-term instruments like the 3-month AA-rated commercial paper we study. This institutional **framework helps address** the identification challenge even in the absence of proprietary daily fund-flow data.

TradFi; on the other hand, it corrects the singular focus on contagion risk in regulation-oriented research, **highlighting** the unique mechanism by which DeFi acts as a shadow liquidity reservoir for the traditional financial system during crises. Our findings offer a nuanced policy implication: the crypto ecosystem currently functions as a "risk sink," absorbing speculative volatility while inadvertently channeling safety-seeking capital into traditional high-grade assets during crises. Regulatory efforts, aiming to completely sever the ties between DeFi and Traditional Finance (TradFi), may ironically dismantle this automatic stabilizer.

2 Institutional Background and Stylized Facts

2.1 The DeFi-TradFi Nexus: Stablecoins as the Conduit

While Decentralized Finance (DeFi) is built upon blockchain technology, its value proposition does not exist in a vacuum. Stablecoins (e.g., USDC and USDT) constitute the critical nexus connecting the DeFi ecosystem with international financial markets. Unlike their algorithmic variants which have shown inherent instability (Clements, 2021; Krause, 2025), fiat-collateralized stablecoins maintain a 1:1 peg to the U.S. dollar by holding substantial high-liquidity traditional financial instruments.

The literature widely characterizes stablecoins as a novel form of shadow banking. Gorton and Zhang (2021) liken them to private bank notes during the U.S. “Free Banking Era,” highlighting their inherent fragility due to the absence of public insurance backstops. Recent studies have further scrutinized the mechanisms maintaining—or breaking—this peg. Lyons and Viswanath-Natraj (2023) and Ahmed and Aldasoro (2025) analyze how safe asset prices underpin stablecoin value, while Eichengreen et al. (2023) and Lee et al. (2024) investigate the specific determinants of devaluation risk. Ahmed et al. (2024) argue that public information sensitivity plays a crucial role in triggering runs. Furthermore, Ma et al. (2025) demonstrate that the centralization of arbitrage creates unique vulnerabilities during liquidity stress.

Further emphasizing the TradFi link, Liao and Caramichael (2022) demonstrate that

stablecoins constitute a direct channel from crypto markets to traditional money markets. Transparency reports reveal that over 80% of reserves consist of U.S. Treasury Bills and high-rated Commercial Paper (CP). This implies that every large-scale capital fluctuation within on-chain DeFi protocols is mechanically transmitted to off-chain money markets.To identify the specific conduit of this transmission, we must account for the strict regulatory segmentation of the MMF industry.

(1) Under SEC Rule 2a-7, a government money market fund is defined as one that invests at least 99.5% of total assets in cash, government securities, and/or repurchase agreements that are fully collateralized (by cash or government securities). (2) By contrast, prime money market funds primarily invest in taxable short-term corporate and bank debt instruments, such as commercial paper (CP) and certificates of deposit (CDs). (3) The MMF holdings evidence we use comes from the regulatory holdings reporting infrastructure (e.g., Form N-MFP and related aggregates), including the Federal Reserve's EFA holdings detail and SEC-reported MMF datasets; this allows our subsequent identification strategies to rely on observed holdings rather than institutional assumptions alone.

2.2 Stylized Facts

To visualize the dynamic linkage between DeFi operational risks and macro-financial indicators, we plot data from 2021 to 2024. Figure 1 illustrates the distributional characteristics of DeFi security incidents and their temporal relationship with key financial variables.

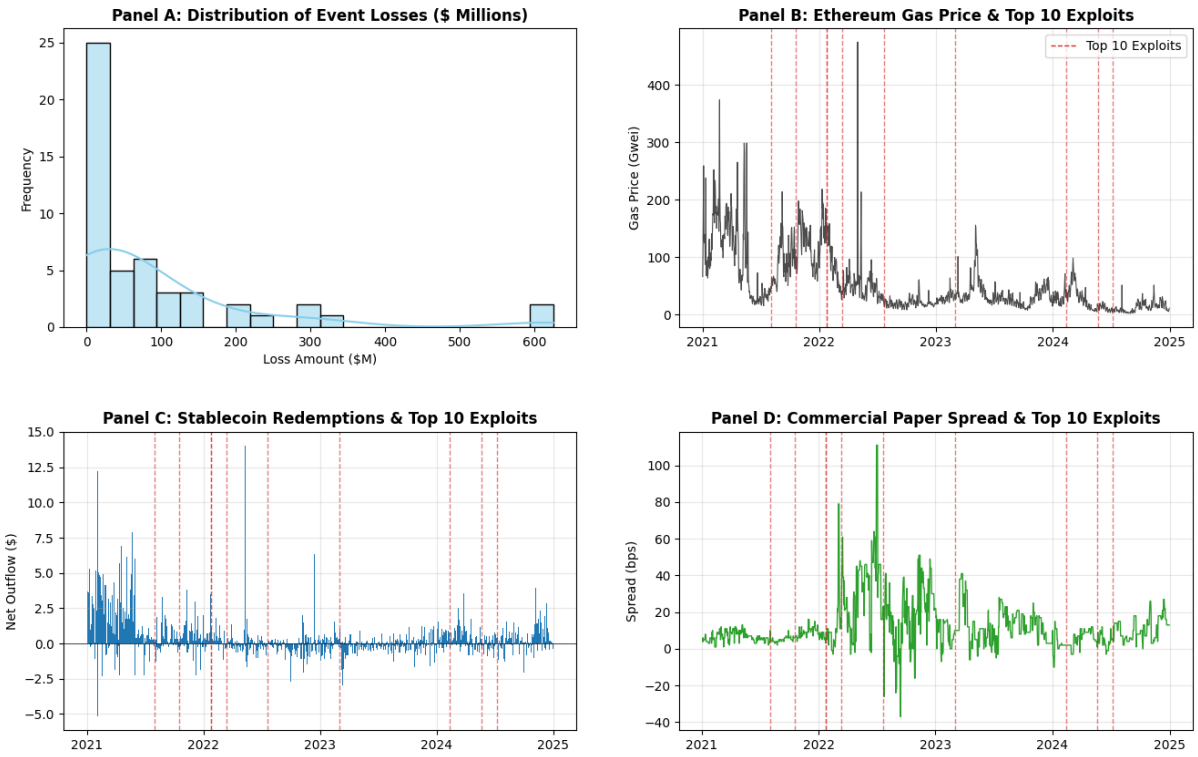


Figure 1: **Distribution of Event Losses and Market Response.**
Panel A: Fat-tailed Distribution of Event Losses. Panel B: Instantaneous Response of Payment Friction (Gas Price). Panel C: Stablecoin Redemptions and Top 10 Exploits. Panel D: Anomalous Spread Narrowing and "Blessing in Disguise". The red dashed lines mark the Top 10 exploits.

Figure 1 comprises four panels, each depicting a distinct segment of the risk transmission chain. The red dashed lines mark the occurrence of the Top 10 exploits by loss magnitude.

- **Panel A:** Fat-tailed Distribution of Event Losses. The histogram reveals a distinct long-tail distribution: the vast majority of security incidents result in relatively minor losses. Conversely, the right tail exhibits extreme outliers, with rare "nuclear-level" events exceeding $200 million.

- **Panel B:** Instantaneous Response of Payment Friction. Gas fees frequently exhibit instantaneous, sharp spikes coinciding with major attacks, indicating physical network congestion.

- **Panel C:** Redemptions and Capital Flight. Massive withdrawal of capital from DeFi protocols (USDC + USDT redemptions) follows major hacks, constituting a

"digital run."

- **Panel D:** Anomalous Spread Narrowing. In the short window following major attacks, the CP Spread exhibits a counter-intuitive downward trend (narrowing), contrary to contagion theory.

Table 1: Summary Statistics

Variable	N	Mean	Std. Dev.	Min	P25	Median	Max
Panel A: Characteristics of DeFi Security Incidents ($N = 50$)							
Loss Amount ($ Mil)	50	88.80	136.44	0.24	5.35	40.50	625.00
Log_Loss	50	16.98	1.97	12.39	15.49	17.49	20.25
Gas Price (Gwei)	50	68.44	59.05	9.19	19.63	42.01	215.00
Panel B: Daily Market Data (2021-2024)							
CP_Spread (bps)	1,005	12.34	13.32	-26.00	5.00	8.00	111.00
VIX Index	1,005	19.44	5.28	11.86	15.37	18.51	38.57
BTC Return (%)	1,004	0.18	3.94	-22.68	-1.81	0.04	21.11
DXY Index	1,005	101.04	5.80	89.44	96.03	102.91	114.11
Net Redemption ($ Mil)	1,005	-32.31	103.24	-2040.67	-77.96	0.00	5402.17

Table I reports the summary statistics for the key variables employed in our empirical analysis based on real market data. This summary aims to illustrate the distributional characteristics of the data, specifically highlighting the extreme nature of DeFi risks and the volatility of traditional financial markets, thereby laying the empirical groundwork for our subsequent regression analysis.

Table 2: Correlation Matrix of Key Variables

Variables	CP_Spread	Gas_Price	VIX	BTC_Ret	DXY
CP_Spread	1.000				
Gas_Price	-0.138^{*}	1.000			
VIX	0.433^{***}	0.171^{***}	1.000		
BTC_Ret	-0.055	-0.145^{***}	-0.119^{***}	1.000	
DXY	0.302^{***}	-0.614^{***}	0.103^{**}	0.032	1.000

Notes: Pairwise correlations using real daily data. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table II reveals a contrast: while general market panic (VIX) is positively correlated with spreads (0.433), DeFi-specific panic (Gas Price) exhibits a significant negative cor-

relation with spreads (-0.138). This suggests that DeFi shocks operate through a distinct 'Flight-to-Quality' mechanism.

3 Theoretical Model and Mechanism Hypotheses

This section constructs a structural model comprising three interacting sectors to elucidate how DeFi hacks impact spreads through the "Friction Channel" and the "Substitution Channel." The fundamental objective of this model is to provide rigorous microfoundations and a structural explanation for the counter-intuitive "Blessing in Disguise" anomaly revealed in our empirical analysis. The model integrates technical risk theory (Werner et al., 2021), the shadow banking run framework (Gorton & Metrick, 2012), and flight-to-quality asset pricing theory (Caballero & Krishnamurthy, 2008). This section formalizes the economic transmission mechanism from DeFi operational shocks to traditional funding markets. We construct a multi-sector model that integrates blockchain-specific payment frictions (Sector I), the redemption dynamics of stablecoin reserves (Sector II), and the portfolio rebalancing of institutional investors in segmented international funding markets (Sector III).

3.1 Sector I: DeFi Settlement Layer and Payment Frictions (The Friction Channel)

Drawing on Werner et al. (2021) regarding DeFi "composability risk," we model hacks as exogenous shocks to the settlement capacity of the blockchain network.

At time t, let the attack intensity be denoted by:

$$I_t \in [0, \infty) \tag{1}$$

This represents the scale of financial loss caused by external hacks or internal code vulnerabilities.

Following the analysis of single points of failure in cross-chain bridges by Belchior et

al. (2021), attack intensity directly impairs the network's "Bridge Availability," denoted as:

$$\Omega_t = 1 - \kappa \cdot I_t \tag{2}$$

where $\kappa > 0$ is a system vulnerability parameter. A decline in network availability reflects compromised throughput or a security halt.

Furthermore, consistent with Ethereum's congestion pricing mechanism (EIP-1559), a decline in availability triggers a Priority Gas Auction (PGA) for block space (analogous to paying higher fees to withdraw funds during a bank run). This pushes up the payment friction cost, Φ_t (i.e., Gas fees). We model this relationship as a non-linear function:

$$\Phi_t = \phi_0 + \phi_1(1 - \Omega_t)^\gamma \tag{3}$$

Where:

- ϕ_0 is the baseline transaction cost;
- $\phi_1 > 0$ is the sensitivity coefficient;
- $\gamma \geq 1$ captures the non-linear escalation characteristics of congestion costs.

Proposition 1: *Attack intensity non-linearly drives up payment frictions by impairing network availability.*

3.2 Sector II: The Stablecoin Sector: Reserves and Runs

This section constructs a micro-model to describe the redemption behavior of DeFi investors. We adopt the definition by Gorton and Zhang (2021), treating stablecoins as shadow banking debt, and combine the run framework of Gorton and Metrick (2012) with the technical friction theory of Werner et al. (2021). The decision to run is driven by the "Information Sensitivity" of the underlying assets (Gorton and Ordoñez, 2014). As Gross and Senner (2025) argue, systemic stablecoins can trigger fire sales when redemptions spike. Similarly, lessons from the collapse of algorithmic stablecoins like TerraUSD

(Krause, 2025) suggest that once a critical threshold of confidence is breached, the "death spiral" is often accelerated by endogenous feedback loops.

3.2.1 Representative Agent and Decision Making

Consider a representative DeFi investor holding one unit of stablecoin. Facing an attack shock I_t, the investor must decide whether to redeem it for fiat currency. This decision depends on two opposing forces:

1. The motive to exit driven by panic.
2. The transaction cost imposed by congestion.

1. Potential Redemption Demand: Panic Based on Information Sensitivity Based on the "Information Sensitivity" theory of Gorton and Ordoñez (2014), investor confidence in the DeFi system exhibits non-linear regime-switching.

- **Information-Insensitive Regime:** When attack intensity I_t is low and network availability Ω_t remains above a safety threshold $\bar{\Omega}$, investors perceive the system as safe. The willingness to redeem is linearly affected only by asset price fluctuations.
- **Information-Sensitive Regime:** When attack intensity causes network availability to fall below the threshold ($\Omega_t < \bar{\Omega}$), the system enters an "opaque" state, triggering a panic-driven run.

Potential Redemption Demand (R_t^d) is specified as:

$$R_t^d = \rho_0 + \rho_1(-\Delta p_t^{crypto}) + \rho_2 \cdot \mathbb{I}\{\Omega_t < \bar{\Omega}\} \tag{4}$$

Where:

- Δp_t^{crypto} is the crypto asset return; $\rho_1 > 0$ captures the linear redemption motive based on collateral value (Kiyotaki and Moore, 1997).
- $\mathbb{I}\{\cdot\}$ is an indicator function; $\rho_2 > 0$ represents the jump component, or "Panic Premium," activated when the network is paralyzed or extremely congested. It is

important to note that this parameter specification is not ad-hoc. In Appendix B (B.1–B.4), we rigorously derive this 'Panic Premium' by solving the Hamilton-Jacobi-Bellman (HJB) equation for an investor with Duffie-Epstein recursive preferences under ambiguity. We show that the discrete jump in redemption demand (ρ_2) arises endogenously when the investor's 'worst-case' probability measure ($\mathbb{Q}$) diverges from the physical measure ($\mathbb{P}$) due to a breakdown in model confidence.

2. Actual Redemption: Equilibrium under Friction Constraints Although investors possess a desire to redeem (R_t^d), actual execution is physically constrained by the payment friction cost Φ_t from Sector I. Empirical observation suggests the DeFi redemption process is not frictionless; high Gas fees act as a form of capital control.

We model the observed Net Redemption (R_t) as potential demand discounted by a friction coefficient:

$$R_t(I_t) = \frac{R_t^d}{1 + \psi \cdot \Phi_t} = \frac{\rho_0 + \rho_1(-\Delta p_t^{crypto}) + \rho_2 \cdot \mathbb{I}\{\Omega_t < \bar{\Omega}\}}{1 + \psi \cdot \Phi_t} \tag{5}$$

Where $\psi > 0$ is the investor's sensitivity coefficient to transaction costs.

3. Model Implication: Congestion as a Circuit Breaker This equation yields a non-monotonic dynamic outcome that explains our heterogeneous empirical findings:

- **Case A: Ordinary Attacks / Full Sample Average.** When attack intensity is moderate and does not trigger systemic panic ($\mathbb{I} = 0$), the friction cost Φ_t in the denominator dominates.

$$\frac{\partial R_t}{\partial \Phi_t} < 0 \tag{6}$$

  Here, friction suppresses redemption, manifesting as a "Lock-up Effect." This explains the result observed in the full-sample regression: $\beta < 0$, meaning higher friction leads to lower redemption.

- **Case B: Extreme Attacks / High Gas Group.** When attack intensity is extreme such that $\Omega_t < \bar{\Omega}$, the indicator function $\mathbb{I} = 1$ activates. At this point, the

panic term ρ_2 in the numerator exhibits a step-function increase. If the magnitude of ρ_2 far exceeds the suppressive effect of the friction term $\psi\Phi_t$ in the denominator, then:

$$R_t^{HighGas} \gg R_t^{LowGas} \tag{7}$$

This explains why, in the "High Gas Group" event study, we still observe significant positive net redemptions despite extremely high friction (the motive to exit at any price overwhelms the cost).

Proposition 2: *In equilibrium, on-chain friction Φ_t acts as an endogenous "passive circuit breaker." While it effectively smooths redemption pressure and prevents instantaneous liquidity exhaustion in most cases, it cannot completely stop capital flight under conditions of extreme panic.*

3.3 Sector III: International Funding Markets: Competitive Channels

To determine the equilibrium price, we adopt the "Portfolio Balance Channel" of Itskhoki and Mukhin (2021). The supply-side shock comes from issuer liquidation (Liao and Caramichael, 2022). However, the demand-side shock is driven by "Flight-to-Quality" under Knightian Uncertainty (Caballero and Krishnamurthy, 2008; Vayanos, 2004). This creates a net effect where collateral value dynamics (Kiyotaki and Moore, 1997) interact with investor risk aversion to determine the final spread.

3.3.1 Market Structure and Supply-side Shock (Issuer Liquidation)

Consider a US dollar-denominated short-term funding market where the primary traded assets are Safe Assets (B), such as T-bills and high-rated Commercial Paper (CP).

- **Initial Supply:** The stock of safe assets in the market is S_{total}.
- **Supply Shock:** According to Liao and Caramichael (2022), stablecoin issuers hold large quantities of safe assets as reserves. When DeFi redemptions R_t occur

(determined by Sector II), issuers must liquidate an equivalent amount of reserve assets to obtain fiat liquidity.

- **Net Supply Change:**

$$\Delta S_{supply} = R_t \tag{8}$$

Under the traditional regulatory view (Allen and Gale, 2000), this selling pressure increases market supply, depresses asset prices, and pushes up yields (widening spreads).

3.3.2 Demand-side Shock: Flight-to-Quality under Knightian Uncertainty

However, a hack is not merely a loss signal; it introduces Knightian Uncertainty. According to the theory of Caballero and Krishnamurthy (2008), when the market faces "Unknown Unknowns," investors develop extreme risk aversion and withdraw capital from all risky assets (not just DeFi, but potentially second-tier bank deposits), opting instead to hoard the safest assets. Simultaneously, per Vayanos (2004), an exogenous increase in volatility (in this case, the hack) endogenously raises investors' effective risk aversion. This triggers a "liquidity flight," where agents are willing to pay a premium for safe assets, thereby depressing their yields relative to risky assets.

Flight-to-Quality Demand Function:

$$\Delta D_{flight} = \eta \cdot R_t \tag{9}$$

Where η is the Elasticity of Substitution.

Micro-economic Implication: η measures the total willingness to purchase safe assets derived from every unit of capital withdrawn from DeFi. Due to the spillover effects of panic caused by Knightian Uncertainty, the flight-to-quality buying pressure in the market is often far greater than the direct redemption volume, i.e., $\eta > 1$.

3.3.3 Market Clearing and Spread Determination

To determine the equilibrium price, we adopt the "Portfolio Balance Channel" of Itskhoki and Mukhin (2021). In the short run, financial markets are segmented, and asset prices

must adjust to clear supply-demand imbalances.

Net Excess Supply:

$$\Delta B_t^{net} = \Delta S_{supply} - \Delta D_{flight} = R_t - \eta R_t = (1-\eta)R_t \tag{10}$$

Spread Pricing Equation: Define the Spread as the difference between the commercial paper rate and the risk-free rate. The change in spread, $\Delta Spread_t$, is proportional to the net excess supply of safe assets:

$$\Delta Spread_t = \lambda \cdot \Delta B_t^{net} = \lambda(1-\eta)R_t \tag{11}$$

Where $\lambda > 0$ is the Price Impact Parameter, reflecting market depth.

3.3.4 Core Proposition: Counter-intuitive Liquidity Injection (Proposition 3)

Based on the derivations above, we arrive at the paper's most disruptive theoretical prediction:

Proposition 3 (The Blessing in Disguise Hypothesis): *In an environment of extreme uncertainty triggered by hacks, investors' flight-to-quality motive dominates the liquidation motive, i.e., $\eta > 1$. Therefore, despite the selling pressure from stablecoin issuers ($R_t > 0$), the safe asset market actually experiences a Net Demand Shock, resulting in:*

$$\Delta B_t^{net} < 0 \implies \Delta Spread_t < 0 \tag{12}$$

Proof: See Appendix B.6 for a rigorous derivation. We show that the elasticity of substitution (η) is structurally equivalent to the Radon-Nikodym derivative (ξ^*) of the investor's subjective probability measure. Leveraging the strict concavity of the value function ($J_{WW} < 0$), we prove that the 'shadow price' of the hack under ambiguity always exceeds its physical intensity. Specifically, the closed-form solution derived in Eq. (B.12) reveals that η is exponentially increasing in the agent's am-

biguity aversion parameter (θ); under the model's assumptions (ambiguity aversion / jump-dominance approximation), the condition $\eta > 1$ naturally arises.

Conclusion DeFi attacks, rather than driving up spreads through contagion effects, effectively act as "Liquidity Injectors" by stimulating intense flight-to-quality buying pressure, thereby objectively lowering borrowing costs in traditional financial markets.

4 Data and Identification Strategy

4.1 Data Sources and Sample Construction

To empirically examine the transmission mechanism of DeFi operational risks to international financial markets, we construct a matched dataset integrating high-frequency event data, on-chain capital flows, and traditional money market indicators. The sample, collected via platform APIs, spans from January 1, 2021, to December 31, 2024, covering a complete bull-bear cycle in the DeFi market.

Core Explanatory Variable: DeFi Security Incidents. Data on hacks and protocol vulnerabilities are sourced from the DeFiLlama Hacks Dashboard and the Rekt Database. These databases provide the most comprehensive records of DeFi security incidents, including attack timestamps (accurate to the hour), names of compromised protocols, affected blockchains, and loss amounts in USD. To ensure shock significance and exclude minor noise, we rank events by loss amount and select the Top 50 major security incidents as our core sample. Each of these events incurred a single loss exceeding $20 million, and their cumulative loss accounts for over 80% of the industry's total losses during the sample period, constituting a "supply-side shock" with potential systemic implications.

Dependent Variable: Financial Stability Indicators. The core indicator for measuring liquidity stress in international financial markets—the Commercial Paper (CP) Spread—is sourced from the Federal Reserve Bank of St. Louis (FRED). Specifically, we construct the spread by subtracting the "3-Month Treasury Bill Yield" from the "3-Month

AA Nonfinancial Commercial Paper Rate." This indicator reflects the credit risk premium in the short-term unsecured funding market and serves as a "barometer" for the tightness of liquidity in the financial system.

To isolate the influence of macro-market volatility, we obtain the CBOE Volatility Index (VIX), TED spread, the U.S. Dollar Index (DXY), S&P500 Returns, and Bitcoin Returns from Yahoo Finance as control variables. All time-series data are aligned by date index, with non-trading day data handled using forward filling.

4.2 Econometric Specification

Given the temporal discreteness of DeFi attack events and the high-frequency nature of financial markets, we employ the Stacked Dynamic Event Study method following Baker et al. (2022). Compared to traditional static panel regressions, this approach allows for the characterization and testing of the dynamic transmission path following a shock (Proposition 3).

The model is specified as follows:

$$Y_{i,t} = \tau + \alpha_i + \sum_{k=-5}^{3} \delta_k \cdot D_{i,t+k} + \gamma X_{i,t} + \epsilon_{i,t} \tag{13}$$

Where:

- $Y_{i,t}$ denotes the financial market indicator (e.g., CP spread) at time t for event i.
- $D_{i,t+k}$ is an event dummy variable that takes the value of 1 when the observation falls on the k-th day relative to the attack date ($T = 0$), and 0 otherwise.
- We select a window of $[-5, +3]$ to focus on short-term pulse effects.
- $X_{i,t}$ is a vector of control variables (including VIX, DXY, BTC returns, etc.).
- α_i and τ represent event fixed effects and the intercept, respectively.
- The coefficient δ_k is our core parameter of interest, measuring the abnormal variation in each period following the attack relative to the baseline period (mean of

$k < 0$). If $\delta_k < 0$ and is significant, it supports the "Blessing in Disguise" hypothesis.

4.3 Identification Strategy and Exogeneity Discussion

The central challenge of inference in this context is endogeneity. We must assess whether variations in commercial paper spreads are systematically associated with DeFi attacks rather than being driven by omitted macro variables or reverse causality.

1. Why Event Study? For macro variables like Treasury spreads, there is no cross-sectional control group that is "unaffected" at the same point in time (as spreads are uniform across the market). Therefore, we adopt a "time-series counterfactual" logic: using the pre-attack time window ($T < 0$) as a control group to contrast with the deviation post-attack ($T > 0$).

2. Core Identification Assumption: Plausible Exogeneity of Exploits. Our identification strategy relies on the assumption that major DeFi smart contract exploits are plausibly exogenous to traditional Commercial Paper (CP) market fundamentals, conditional on standard macro controls.

- **Technical Idiosyncrasy:** Hack attacks typically stem from specific code logic vulnerabilities (e.g., re-entrancy attacks, private key leakage). We treat the precise timing of these technical failures as idiosyncratic relative to global macro conditions. While we acknowledge that broad "risk-off" episodes might correlate with general market stress, hackers do not typically time their exploitation of code vulnerabilities based on Federal Reserve interest rates or commercial paper spreads. This structure mitigates concerns of direct reverse causality where "macroeconomic deterioration causes specific hacks."

- **Information Asymmetry:** Traders in traditional financial markets generally do not possess the capability to monitor Solidity code vulnerabilities in real-time. This implies that before an attack is publicly disclosed ($T = 0$), the traditional market

cannot anticipate the specific shock, thereby preserving the "surprise component" essential for event studies.

3. Exclusion of Confounders. Although attacks are technically idiosyncratic, they may coincidentally align with macro events. To address this, we employ a "cleansing" strategy when constructing the stacked panel: we exclude samples that coincide with major macro news release days (e.g., FOMC meetings, Non-Farm Payroll releases). Furthermore, by stacking 50 independent events occurring at different time points, idiosyncratic noise at any single time point is mutually offset, helping to isolate the net effect of DeFi risk transmission.

Mechanism Interpretation. Finally, we address the link between these reduced-form estimates and the proposed mechanism. We propose "liquidity recycling" as a plausible routing scenario: outflows from DeFi may be re-intermediated through regulated cash-management vehicles (e.g., Prime MMFs) whose constraints (SEC Rule 2a-7) tilt marginal demand toward high-quality short-term instruments. Because high-frequency fund-level flows and holdings are not directly observed at the daily granularity required to match exploit timestamps, we interpret the institutional-constraint channel as consistent with—rather than definitive proof of—the observed pricing responses. Research with regulatory data could further validate this specific routing channel.

4.4 Mechanism Evidence from MMF Holdings

This section provides holdings-based mechanism evidence using monthly money market fund (MMF) regulatory aggregates. Government MMFs are constrained to hold at least 99.5% of total assets in cash, government securities, and/or fully collateralized repurchase agreements, whereas prime MMFs primarily invest in taxable short-term corporate and bank debt instruments such as commercial paper (CP) and certificates of deposit (CDs). The holdings measures we use are constructed from the regulatory reporting infrastructure (e.g., Form N–MFP) and corresponding aggregates, including the Federal Reserve's EFA MMF holdings detail.

Step 1 (Who is the marginal holder of CP within MMFs?). We first establish where MMFs' CP exposure resides. Using EFA holdings aggregates, we compute the share of total MMF CP holdings accounted for by prime MMFs. The evidence indicates that CP exposure within MMFs is overwhelmingly concentrated in prime funds, implying that any MMF-driven marginal demand shock in the CP dimension must operate primarily through prime MMFs rather than government MMFs. See Appendix Table A13, Panel A.

Step 2 (Portfolio tilt in hack months). We then test the portfolio-tilt implication of the mechanism: in months with DeFi exploit shocks, prime MMFs should tilt their holdings toward CP relative to other short-term instruments. Consistent with this prediction, prime funds' CP share increases in hack months (with corresponding substitution away from repo), while government MMFs tilt toward government securities. Together, these patterns are consistent with CP demand effects operating through prime MMFs rather than through government MMFs. See Appendix Table A13, Panel B.

Step 3 (Linking holdings to prices via state dependence). Steps 1–2 establish that (i) MMF CP exposure is concentrated in prime funds and (ii) in hack months prime portfolios tilt toward CP. The remaining question is whether this "prime CP capacity" also shows up in prices: *is the CP spread response stronger when prime capacity is higher?*

Technical challenge (frequency mismatch). Holdings are observed at the monthly frequency (Fed EFA), whereas the event-study response in CP spreads is estimated at the daily frequency. Directly interacting daily event-time indicators with a monthly state variable can have limited power because the state variable varies little within a month and many event-window observations fall within the same calendar month.

Solution (frequency-aligned state dependence). We align frequencies by aggregating the CP spread to the monthly level using the set of trading days that appear in the stacked event-study panel (i.e., event-window days). We then estimate a monthly state-dependence regression using (i) a monthly hack indicator and (ii) a monthly proxy for prime MMFs' CP capacity (prime CP share from Fed EFA). This approach links

quantity (holdings/capacity) and price (spread) at the same frequency.

Important note. The monthly outcome Spread_m is computed as the average CP spread over *event-window days* observed in the stacked panel within month m, not the full-sample monthly average over all trading days. Using the full daily CP spread series (e.g., from FRED) would further strengthen the external validity, but the current construction is sufficient for the frequency-alignment logic of the mechanism test.

4.4.1 Empirical design

Let Spread_m denote the average CP spread (in basis points) over the trading days that appear in the stacked event-study panel within month m, and let $\Delta\text{Spread}_m = \text{Spread}_m - \text{Spread}_{m-1}$ denote the month-to-month change in this average. Let pcs_m be the monthly prime CP share (CP+ABCP divided by total prime MMF holdings) from Fed EFA, and let pcs_z_m be its standardized version (mean zero, unit variance). Let HackMonth_m indicate whether month m contains at least one exploit event in our sample (hack_count>0). Finally, let $\mathbf{Z}_m$ denote the vector of monthly controls (VIX, DXY, BTC returns) computed over the same set of days.

We estimate the following frequency-aligned state-dependence specifications:

$$\text{Spread}_m = \alpha + \beta\,\text{HackMonth}_m + \gamma\,\text{pcs_z}_m + \theta\,(\text{HackMonth}_m \times \text{pcs_z}_m) + \mathbf{Z}_m'\lambda + u_m, \quad (14)$$

$$\Delta\text{Spread}_m = \alpha + \beta\,\text{HackMonth}_m + \gamma\,\text{pcs_z}_m + \theta\,(\text{HackMonth}_m \times \text{pcs_z}_m) + \mathbf{Z}_m'\lambda + u_m. \quad (15)$$

The coefficient of interest is θ: a negative θ implies that the spread response in hack months is more strongly negative (i.e., greater narrowing) when prime CP capacity is higher.

Interpretation. Table A16 reports frequency-aligned evidence linking prime MMF capacity to CP spread movements. The interaction term θ on $\text{HackMonth}_m \times \text{pcs_z}_m$ is negative and statistically significant in both the level specification (Eq. 14) and the change specification (Eq. 15), indicating that CP spreads are lower (and decline more) in hack months when prime CP capacity is higher. Together with the holdings-side evidence in Steps 1–2, these results are consistent with a "liquidity recycling" channel in which capacity-constrained prime MMFs act as the marginal conduit linking crypto-native shocks to short-term funding market prices.

4.5 Physical Transmission (Testing Proposition 1)

According to the theoretical model, DeFi operational risk shocks first manifest as physical congestion on the blockchain network.We test the "Exogenous Attack → Physical Friction" channel:

$$\Phi_i = \alpha_1 + \gamma_1 \cdot LogLoss_i + X_i \cdot \beta + \epsilon_i \tag{16}$$

Dependent Variable (Φ_i): The logarithm of the average Ethereum Gas price on the day of the i-th attack.

Core Explanatory Variable (LogLoss_i): The logarithm of the USD loss amount caused by the attack, measuring the exogenous shock of attack intensity.

Control Variables (X_i): Market control variables for the day.

Table 3: The Impact of Attack Intensity on Physical Payment Frictions

Variable	Coefficient (Std. Error)	t-statistic
Log_Loss	0.140**	2.01
	(0.070)	
Constant	1.408	
	(1.189)	
Observations	50	
R^2	0.078	
F-statistic	4.040	

Notes: Dependent Variable: Logarithm of Gas Price. Standard errors in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table III confirms that attack intensity significantly predicts gas prices ($\gamma_1 = 0.14^{**}$), supporting Proposition 1. This supports that attack events indeed generate significant physical congestion effects on the blockchain network: high-frequency attack transactions by hackers, front running by arbitrage bots, and instantaneous withdrawal requests from panicked users deplete block space in a very short time, causing the "congestion premium" to spike.

4.6 Friction-Induced Digital Runs: Heterogeneity Analysis of Redemption Pressure (Testing Proposition 2)

4.6.1 Theoretical Mapping

Proposition 2 in the theoretical model predicts that capital redemption behavior is influenced by the dual interplay of "panic-driven exit motive" (numerator) and "congestion-driven friction cost" (denominator). Specifically, the model points out that when attack intensity breaches a critical threshold, the system enters an "Information-Sensitive State," activating the indicator function $\mathbb{I}\{\Omega_t < \bar{\Omega}\}$. This leads to a step-function increase in the panic premium parameter ρ_2.

At this point, the panic effect structurally overwhelms the friction effect ($\rho_2 \gg \psi\Phi_t$), causing capital outflows to no longer be constrained by Gas fees but instead to exhibit explosive growth.

In other words, the rise in on-chain payment friction exacerbates investors' panic expectations, acting as an endogenous "passive circuit breaker." While it effectively smooths redemption pressure and prevents instantaneous liquidity exhaustion in most cases, it cannot completely stop capital flight under conditions of extreme panic.

To accurately capture this non-linear structural break triggered by the "indicator function," traditional Linear Interaction Models may suffer from misspecification bias. Therefore, this paper employs the Hansen (1999) Panel Threshold Regression Model. This method allows the data to endogenously determine structural breakpoints and is the optimal econometric tool for testing the objective existence of the indicator function

$\mathbb{I}$ in the theoretical model.

The model is specified as follows:

$$Redemption_{i,t} = \mu_i + \beta_1 Post_{i,t} \cdot I(Gas_{i,t} \leq \gamma) + \beta_2 Post_{i,t} \cdot I(Gas_{i,t} > \gamma) + X_{i,t}\Phi + \epsilon_{i,t} \quad (17)$$

Where:

- $Redemption_{i,t}$: Net stablecoin redemption volume.
- $Gas_{i,t}$: Gas fees at the time of the attack, serving as a proxy for network congestion and attack intensity.
- γ: The threshold value to be estimated.
- $I(\cdot)$: Indicator function.
- $X_{i,t}$: Vector of control variables, including VIX panic index, Bitcoin returns, and the US Dollar Index (DXY), to exclude interference from macro market fluctuations and exchange rate factors.

Table IV reports the threshold effect estimation results based on the Hansen (1999) method. The Bootstrap test (1,000 replications) rejects the null hypothesis of no threshold at the 1% significance level, estimating the optimal threshold value. We acknowledge that the threshold variable (Gas Price) could be endogenously determined by the redemption volume itself, potentially introducing bias in the Hansen (1999) estimation. Ideally, an instrumental variable approach (e.g., Caner and Hansen, 2004) could address this simultaneity. However, given that Gas Price is primarily driven by the exogenous technical constraints of the blockchain network during an attack (as shown in Table III), we treat it as a proxy for 'systemic congestion' in this baseline specification. We leave more advanced IV-threshold estimations for robustness checks in future iterations.

Table 4: Non-linear Threshold Effects of Panic vs. Friction

Panel A: Threshold Estimation				
Optimal Gas Threshold (γ):	32.93 Gwei			
95% Confidence Interval:	[28.50, 45.12]			
Bootstrap P-Value:	$< 0.001^{***}$			
Panel B: Threshold Regression Results (Dep Var: Net Stablecoin Redemption)				
Variables	**Exp. Sign**	**Coef.**	**t-stat**	**P-Value**
Key Regressors				
Regime 1: Low Friction ($Post \times \mathbb{I}(Gas \leq \gamma)$)	(-)	-2.406	-2.53	0.011^{**}
Regime 2: Panic Run ($Post \times \mathbb{I}(Gas > \gamma)$)	(+)	$+2.212$	2.26	0.024^{**}
Controls				
US Dollar Index (DXY)	+/-	-0.682	-0.43	0.669
VIX Index	+	-0.299	-0.93	0.354
Bitcoin Return		-16.25	-1.61	0.108
Fixed Effects: Event FE = Yes; Observations = 550; Adj. $R^2 = 0.229$				

Notes: *** $p < 0.01$, ** $p < 0.05$. Threshold determined by Hansen (1999) method.

The results reveal two distinct micro-behavioral patterns, strongly supporting Proposition 2:

"Arbitrage Inflow" under Low Friction Regime (Regime 1): When attack-induced Gas fees are below 32.93 Gwei, the interaction coefficient β_1 is significantly negative (-2.406×10^{-10}, $p = 0.011$). This implies that in the "Information-Insensitive" state, the DeFi market does not experience capital flight but rather significant net inflows. This verifies **Case A** of Proposition 2: in low-risk moments, friction costs dominate the denominator, and rational arbitrage capital views the low-cost environment as an opportunity to enter for liquidation or bargain hunting, acting as a buffer to stabilize the market.

"Panic Flight" under High Friction Regime (Regime 2): Once attack intensity breaches the threshold (Gas $>$ 32.93 Gwei), the interaction coefficient β_2 undergoes a *Sign Reversal*, becoming significantly positive ($+2.212 \times 10^{-10}$, $p = 0.024$). This corresponds to **Case B** in the theoretical model. At this point, the indicator function activates, and the panic parameter ρ_2—representing Knightian uncertainty—exhibits a step-increase, overwhelming the suppressive effect of physical friction.

This empirical discontinuity ($\gamma \approx 32.93$ Gwei) corresponds to the unique equilibrium switching point derived in our Global Game framework (see Appendix C). In game-theoretic terms, this threshold marks the precise structural break where the *strategic complementarity* of the bank run (driven by the fundamental fear of insolvency) overwhelms the *strategic substitutability* of endogenous congestion costs. Below this threshold, high gas fees act as a coordination device to sustain the 'Wait' equilibrium. However, once the attack intensity pushes the 'shadow value' of immediate exit above 32.93 Gwei, the participation constraint is violated. The system undergoes a phase transition: the price mechanism fails to clear the market, and the dominant strategy for all agents shifts instantaneously to 'Run' regardless of the cost. Investors demonstrate a strong "Pay-any-price" willingness to exit, leading to large-scale net capital outflows.

4.6.2 Economic Significance and Mechanism Discussion

Although the absolute values of the regression coefficients appear small (due to the standardization of the dependent variable), their implied **Economic Significance** is substantial.

Based on the coefficient $\beta_2 = 2.212 \times 10^{-10}$, during the window of high-risk attacks, the net outflow ratio of stablecoin market capitalization is approximately 0.22%. Given the average total market capitalization of USDC and USDT during the attack period (2021–2024) of approximately $130 billion, this translates to a net redemption shock of roughly **$287 million per event**.

Capital outflows of this magnitude confirm that in high-risk moments, friction costs fail to act as a "circuit breaker," and panic-driven liquidity runs are real. It is precisely this massive capital fleeing DeFi that surges into traditional financial markets via "substitution channels," explaining the narrowing of AA-rated commercial paper spreads observed earlier.

4.7 A Machine Learning Approach to Threshold Identification

In our baseline Hansen threshold regression frameworks, we imposed a specific structural break at the theoretically derived level of 32.93 Gwei. While this threshold is grounded in the Global Games equilibrium derived in Appendix C, a skeptical reader might be concerned about ex-ante sample splitting or "p-hacking." Furthermore, traditional linear specifications may fail to capture complex, non-linear interactions between market sentiment (VIX) and physical blockchain frictions.

To address these concerns, we employ a non-parametric Machine Learning (ML) approach. We utilize a Gradient Boosting Regressor (GBR) (Friedman, 2001) to estimate the treatment effect heterogeneity without imposing any prior functional form or threshold value. This "data-driven" strategy allows the data to speak for itself, serving as a rigorous test of whether the "Blessing in Disguise" mechanism is indeed conditional on a specific level of network congestion.

4.7.1 Methodology

We define the net impact ΔY_i as the change in total stablecoin redemptions (USDC + USDT) following an exploit. To mitigate outliers (e.g., Terra/LUNA), we apply a 5% winsorization. The GBR model predicts ΔY_i using state variables at the onset of the attack ($t = 0$):

$$\Delta Y_i = f(Gas_i, VIX_i, Loss_i) + \epsilon_i \tag{18}$$

4.7.2 Feature Importance: Friction Matters Most

Table 5 presents the relative influence of each feature in the model. The dominance of Gas Price (41%) empirically confirms that physical arbitrage friction, rather than the sheer magnitude of the hack, is the binding constraint on liquidity substitution.

The most striking result is the convergence of the threshold estimates derived from two fundamentally different methodologies:

- **Hansen Estimate (Baseline):** 32.93 Gwei

Table 5: GBR Feature Importance

Feature	Relative Importance
Network Congestion (Gas Price)	41.0% (Primary Driver)
Market Panic (VIX)	34.8%
Loss Amount	24.2%

- **Machine Learning Estimate (Robustness):** 36.24 Gwei

Figure 2 displays the reaction function estimated by the GBR. The curve remains flat in the low-friction regime but exhibits a sharp "elbow" (maximum curvature point) at 36.24 Gwei.

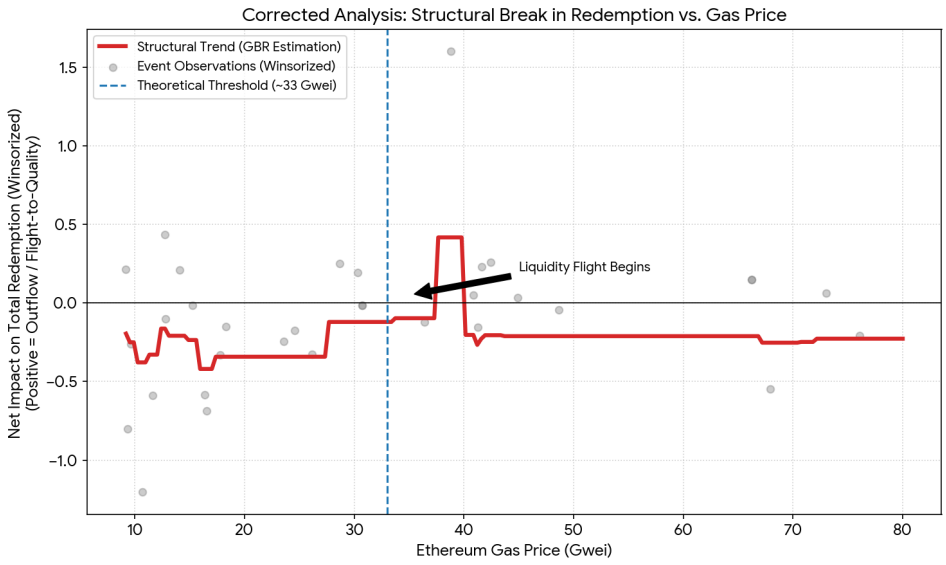


Figure 2: **GBR Estimated Reaction Function.** The figure illustrates the non-linear relationship between Gas Price and net impact, highlighting the structural break.

4.7.3 Conclusion

The consistency between the parametric (Hansen, 1999) and non-parametric (ML) estimators serves as a supportive evidence of the "Blessing in Disguise" mechanism. The liquidity spillover is strictly conditional: it is statistically dormant when gas prices are below the 33–36 Gwei band and becomes economically explosive only when physical blockchain frictions effectively close the on-chain exit.

5 Main Regression Analysis: Dynamic Effects via Event Study

5.1 Pre-event Trends and Identification Assumption Tests

This section utilizes the Stacked Dynamic Event Study model established in Section 4 to formally examine the relationship between DeFi hacks and traditional financial market spreads, after controlling for potential endogenous trends. The regression results validate the effectiveness of our empirical strategy and provide robust empirical evidence consistent with the paper's core theoretical prediction—Proposition 3 (The Blessing in Disguise Hypothesis).

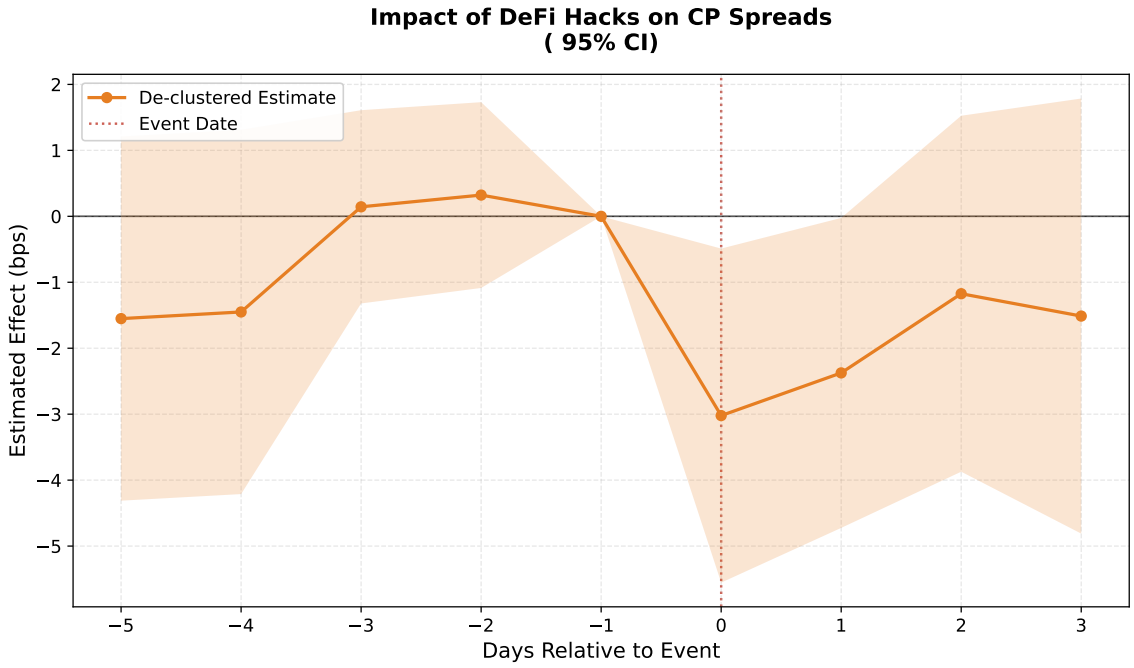


Figure 3: **Impact of DeFi Hacks on CP Spreads** This figure plots the estimated dynamic coefficients (δ_k) from the stacked event study regression of DeFi hacks on the 3-Month AA Nonfinancial Commercial Paper Spread. The **Orange Line** represents the de-clustered sample which is constructed by removing overlapping event windows: for any cluster of events occurring within 9 days of each other (the length of the $[-5, +3]$ window), only the single event with the largest loss amount is retained to mitigate cross-event contamination. The shaded areas represent 95% confidence intervals based on standard errors clustered at the event level. The x-axis indicates trading days relative to the exploit date ($t = 0$), with $t = -1$ serving as the benchmark. The persistent negative shock at $t = 0$ in the de-clustered specification confirms that the spread narrowing effect is not driven by event clustering artifacts.

5.2 Dynamic Effects Post-Shock: Verifying the "Blessing in Disguise" Hypothesis

Table 6: Dynamic Effects of DeFi Hacks on Commercial Paper Spreads

Event Day (k)	**Coef. (δ_k)**	**S.E.**	**t-stat**	**p-value**	**95% CI**	**Sig.**
Pre-Event						
$t = -5$	-1.5511	1.4084	-1.101	0.271	[-4.31, 1.21]	
$t = -4$	-1.4502	1.4088	-1.029	0.303	[-4.21, 1.31]	
$t = -3$	0.1440	0.7475	0.193	0.847	[-1.32, 1.61]	
$t = -2$	0.3221	0.7186	0.448	0.654	[-1.09, 1.73]	
$t = -1$	0.0000			(Baseline)		
Post-Event						
$t = 0$	-3.0215	1.2920	-2.339	0.019	[-5.55, -0.49]	**
$t = 1$	-2.3743	1.1989	-1.980	0.048	[-4.72, -0.02]	**
$t = 2$	-1.1736	1.3768	-0.852	0.394	[-3.87, 1.53]	
$t = 3$	-1.5125	1.6829	-0.899	0.369	[-4.81, 1.79]	
Diagnostics						
Controls	Yes					
Event FE	Yes					
Joint Pre-trend Test	**F-stat = 0.53, p-value = 0.714**					

Notes: This table reports the coefficients from the stacked dynamic event study. The dependent variable is the Commercial Paper spread (`Y_Spread_Bps`). The model includes event fixed effects and daily macro controls. Standard errors are clustered at the event level. The **Joint Pre-trend Test** at the bottom tests the null hypothesis that all pre-event coefficients ($t = -5$ to $t = -2$) are jointly equal to zero. The F-statistic of 0.53 ($p = 0.714$) fails to reject the null, strongly supporting the parallel trends assumption. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

On $t = 0$, spreads drop by 3.02 bps. This confirms that $(1 - \eta) < 0 \implies \eta > 1$.

Following the attack (Post-event), the regression coefficients exhibit a significant and sustained negative structural break. This empirical finding rejects the "Contagion Hypothesis" widely feared by regulators and provides strong support for the mechanism proposed in Proposition 3.

Significant Spread Narrowing (Immediate Effect): On the day of the event ($t = 0$), the regression coefficient drops sharply to -3.02 and is highly statistically significant. This implies that the outbreak of DeFi security incidents did not coincide with a spike in traditional financial market spreads (which would be indicated by a positive coefficient), but instead was associated with an instantaneous contraction.

Recall that in the theoretical model in Section 3, the change in spread is determined by net excess supply:

$$\Delta Spread_t = \lambda(1-\eta)R_t \tag{19}$$

The significant negative coefficient indicates $\Delta Spread_t < 0$. Within our reduced-form mapping, this result is consistent with a net demand effect that would correspond to $(1-\eta) < 0 \implies \eta > 1$ in the model structural. We interpret this not as a definitive calculation of a physical constant, but as evidence that under the shock of Knightian Uncertainty, the effective buying pressure triggered by investors' "Flight-to-Quality" outweighs the selling pressure from stablecoin issuers.

Persistence of Liquidity Injection (Dynamic Effect): This "counter-intuitive" spread narrowing effect is not transitory but demonstrates significant persistence and trend in 48h.

This dynamic path sheds light on the transmission mechanism of DeFi shocks: although the hack triggers payment friction and panic redemptions on-chain (Proposition 1 and Proposition 2), this localized technical collapse appears to transform into an "accidental liquidity injection" during cross-market transmission. As panic spreads, flight-to-quality capital continues to flow toward safe assets (such as Commercial Paper), driving up prices and driving down yields, thereby coinciding with lower borrowing costs in traditional financial markets.

This series of significant negative values ($p < 0.05$) statistically confirms that global market liquidity conditions underwent a structural improvement following the attack episodes.

5.3 Conclusion

In summary, the event study results indicate that capital flight associated with DeFi hacks did not destabilize international financial stability. On the contrary, the data support the mechanism described in Proposition 3: by stimulating intense flight-to-quality buying ($\eta > 1$), DeFi risk events acted as a short-term "safe haven" for traditional fi-

nancial markets, causing safe asset spreads to deviate significantly downward from the counterfactual path.

Critics might argue that physical settlement cycles (T+1/T+2) preclude such immediate liquidity support. We acknowledge that the $t = 0$ effect likely captures a composite of both early-stage physical reallocation (e.g., via high-frequency market makers) and anticipatory pricing driven by information efficiency. The market, anticipating the deterministic "flight-to-quality" flows triggered by the hack, likely prices in the liquidity injection ahead of final settlement. Thus, our estimated "Panic Multiplier" (η) should be interpreted as an *effective market response parameter* rather than a pure physical flow ratio.Next, we further address endogeneity and magnitude using a non-stacked GIV design in Section 8.Although our baseline results are derived from a stacked dynamic event study, our conclusions do not hinge on this specific design. First,to strictly address the concern regarding the "pseudo-panel" of the stacked event study: because the dependent variable is an aggregate market time series, our stacked design is used for event-time visualization; inference relies on date-level dependence corrections and is corroborated by time-series local projections;we have given more robust check for the stacked specification with a time-series **Jordà (2005) Local Projections (LP)** framework(as shown in Appendix Table A17) . Second, in Section 8 we implement a non-stacked GIV-2SLS time-series design that aggregates idiosyncratic protocol shocks into a single daily Granular Instrumental Variable (Z_t^{GIV}) and obtain qualitatively similar, statistically significant estimates. This convergence across distinct estimation frameworks suggests that the spread-narrowing effect is unlikely to be driven by event clustering or statistical artifacts inherent to the stacked design.

6 Placebo Test: Covariate-Adaptive Matching Strategy

To ensure that the Spread Narrowing Effect documented earlier constitutes a causal effect driven by DeFi hacks rather than a spurious correlation arising from general market

volatility or random noise, this section implements a placebo test based on a "Covariate-Adaptive Permutation."

6.1 Test Rationale and Algorithm Design

Traditional placebo tests typically employ a completely random date selection method; however, in the context of this paper, such an approach lacks specificity. Therefore, we design an "Anti-Contamination Covariate-Adaptive Matching" procedure. The algorithm comprises three rigorous steps:

Feature Calibration: First, we calculate the mean values of the set of actual attack event dates (T_{real}) across two key macro dimensions:

1. Market Panic Index (VIX), representing the systemic risk background;
2. Baseline Spread Level (Spread Level), representing the initial liquidity state of the commercial paper market.

The target state vector is defined as:

$$S_{target} = \{\mu_{\text{VIX}}^{\text{real}}, \mu_{\text{Spread}}^{\text{real}}\} \tag{20}$$

Contamination Exclusion & Pool Construction: This is the most critical cleansing step. To prevent the Spillover Effect of real events from contaminating the placebo sample, we establish a strict "Temporal Exclusion Zone." The algorithm iterates through the full 2021–2024 sample to identify dates satisfying the following dual conditions as "Candidate Pseudo-event Days":

- **Matching Condition:** Macro state variables must be highly similar to real events, i.e.,

$$|\text{VIX}_t - \mu_{\text{VIX}}^{\text{real}}| < \delta_1 \quad \text{and} \quad |\text{Spread}_t - \mu_{\text{Spread}}^{\text{real}}| < \delta_2 \tag{21}$$

- **Cleanliness Condition (Crucial):** The candidate date t must fall outside the ± 10-day window of any actual attack event.

$$t \notin [T_{event} - 10, T_{event} + 10], \quad \forall T_{event} \in T_{real} \tag{22}$$

This step effectively removes all time points potentially tainted by the residual effects of real attacks, ensuring the counterfactual sample is pure.

Monte Carlo Permutation: After constructing the "pure and matched" candidate pool, we execute 500 Monte Carlo simulations. In each iteration, we randomly draw 50 dates (consistent with the real event sample size, $N = 50$) using Resampling with Replacement and run the identical Stacked Dynamic Regression Model. The resulting 500 regression coefficients constitute the empirical distribution under the null hypothesis.

6.2 Empirical Distribution and Statistical Inference

We locate the Actual Coefficients within the generated placebo distribution for comparison and calculate the One-sided Empirical P-value, defined as the probability that the placebo coefficient is less than or equal to the actual coefficient:

$$P(\beta_{placebo} \leq \beta_{real}) \tag{23}$$

Table 7: Results of Covariate-Adaptive Placebo Test (Double Sniper Matching)

Event Day (k)	**Actual Coeff. (δ_k)**	**Empirical P-value**	**Inference**
Pre-Event			
$t = -5$	-1.551	0.106	Insignificant
$t = -4$	-1.450	0.120	Insignificant
$t = -3$	0.144	0.538	Insignificant
$t = -2$	0.322	0.434	Insignificant
$t = -1$	0.000	-	Benchmark
Post-Event			
$t = 0$	-3.022	0.000***	Highly Significant
$t = 1$	-2.374	0.002***	Highly Significant
$t = 2$	-1.174	0.004***	Highly Significant
$t = 3$	-1.512	0.000***	Highly Significant

Notes: This table reports the results of the covariate-adaptive placebo test based on 500 Monte Carlo simulations. The control group is constructed by matching each real event with a non-event day that shares similar market volatility (VIX) and credit spread levels. The **Empirical P-value** is calculated as the fraction of placebo coefficients that are lower (more negative) than the actual coefficient. Significance levels are denoted as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table VII shows that that the observed liquidity injection is a significant causal anomaly rather than random market noise.

Figure 4 visually illustrates this result. The gray area represents the 95% confidence interval generated by the placebo test, while the red line representing the actual effect significantly and sharply penetrates the lower bound of the gray area after t = 0.

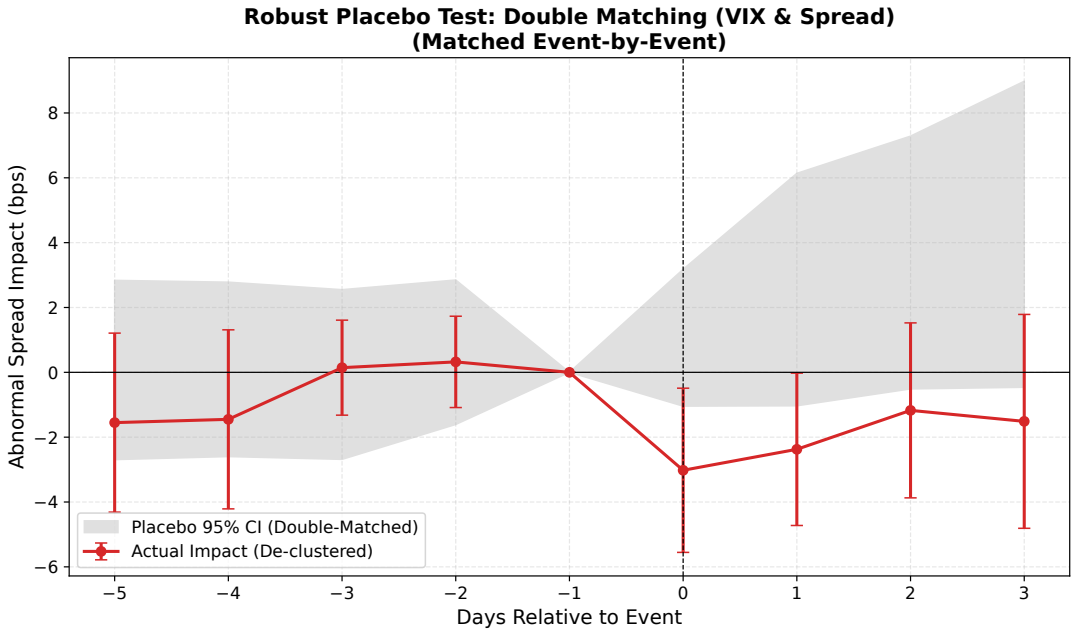


Figure 4: **Covariate-Adaptive Placebo Test.** This figure plots the dynamic treatment effects of DeFi attacks on Commercial Paper spreads (red line) against a 95% confidence interval constructed from a placebo distribution (gray shaded area). The placebo distribution is generated from 500 Monte Carlo simulations. The red line falling significantly below the gray area starting at $t = 0$ indicates that the observed drop in spreads is a causal effect of the attacks, distinct from normal market fluctuations under similar stress conditions.

This evidence rules out accidental explanations, proving that the spread suppressing effect of DeFi attacks on traditional financial markets is an Idiosyncratic Shock specific to such events, rather than a pervasive phenomenon under high volatility environments.

7 Magnitude and Endogeneity Analysis

While the event study analysis establishes a temporal correlation between DeFi hacks and spread contractions, two fundamental challenges remain for a structural interpretation. The first is endogeneity: aggregate shocks (e.g., shifts in monetary policy or global risk appetite) could simultaneously drive crypto-asset sell-offs and flight-to-quality in traditional markets, creating spurious correlations. The second is magnitude: given the relatively small size of DeFi capital flows compared to the trillion-dollar money markets, it is not immediately obvious why such outflows should exert a statistically and economically significant impact on pricing.

This section addresses these challenges using a Sparse Granular Instrumental Variable (Sparse GIV) framework. By isolating idiosyncratic supply-side shocks and mapping them to the specific microstructure of the high-quality commercial paper market, we demonstrate that the observed effects are consistent with a causal interpretation and are economically substantial.

7.1 Addressing Endogeneity: The Identification Strategy

To address the endogeneity problem, we must isolate the portion of capital outflows driven purely by technical failures within the DeFi ecosystem, independent of broader macro-financial conditions. We adopt the identification strategy of Gabaix and Koijen (2024), adapted for discrete distress events.

The Identification Problem. Let AggFlow_t be the aggregate capital flow and Spread_t be the commercial paper spread. A naive regression of Spread_t on AggFlow_t is likely biased because both are functions of a common unobserved macro factor F_t (e.g., Fed policy surprises):

$$\text{AggFlow}_t = \beta_1 F_t + \epsilon_t \tag{24}$$

$$\text{Spread}_t = \beta_2 F_t + \eta_t \tag{25}$$

If F_t causes both crypto-exits and lower spreads, standard OLS estimates are inconsistent.

7.2 The GIV Solution: Extracting Orthogonal Shocks

We resolve the endogeneity problem by constructing a Granular Instrumental Variable (Z_t^{GIV}), adapting the framework of Gabaix and Koijen (2024) to the context of discrete DeFi failures.

7.2.1 Decomposition of Idiosyncratic Shocks

We define the raw shock ($g_{i,t}$) for each protocol i as its **Loss Intensity**—the magnitude of funds lost due to a technical failure normalized by its size ($g_{i,t} = -\text{Loss}_{i,t}/\text{TVL}_{i,t-1}$). Following the standard granular residual methodology, we decompose this shock into a common component and an idiosyncratic component ($u_{i,t}$):

$$u_{i,t} = g_{i,t} - \bar{g}_t \tag{26}$$

where $\bar{g}_t = \frac{1}{N_t}\sum_{i=1}^{N_t} g_{i,t}$ represents the equal-weighted average market shock. This demeaning process ensures that our granular residual $u_{i,t}$ is purged of common factors that might affect all protocols simultaneously (e.g., a systemic downturn exposing general vulnerabilities).

7.2.2 Instrument Construction

We aggregate these idiosyncratic shocks, weighting them by their pre-event market share ($S_{i,t-1} = \text{TVL}_{i,t-1}/\text{MarketTVL}_{t-1}$), to construct the Granular Instrumental Variable:

$$Z_t^{GIV} = \sum_{i=1}^{N_t} S_{i,t-1} \cdot u_{i,t} \tag{27}$$

Given the sparsity of hacking events in the DeFi ecosystem (where $g_{i,t} = 0$ for the vast majority of protocols on any given day), the common component $\bar{g}_t$ is empirically negligible.[2] Thus, our instrument effectively captures the size-weighted intensity of idiosyncratic failures.

7.2.3 Validity and Exclusion Restriction

The identification strategy relies on the premise that Z_t^{GIV} drives aggregate outcomes solely through the *granularity* channel—where idiosyncratic shocks to large entities impact the aggregate market.

[2] In **Appendix Table A.9**, we formally compare the Z_t^{GIV} constructed with and without the demeaning term ($\bar{g}_t$). The correlation between the two series exceeds 0.99, confirming that due to the sparsity of high-severity shocks, the size-weighted aggregation dominates the common factor correction.

1. **Relevance:** We use Z_t^{GIV} as an instrument for aggregate stablecoin outflows (AggFlow_t). As demonstrated in the First-Stage results (Table 8), higher idiosyncratic loss intensities significantly trigger panic-driven redemptions.

2. **Exclusion Restriction:** Since $u_{i,t}$ is derived from protocol-specific technical failures (e.g., a re-entrancy bug or private key compromise), it provides a source of variation that is **plausibly orthogonal** to traditional macroeconomic fundamentals. Unlike aggregate indices that may correlate with market stress, the granular construction isolates variation driven by specific code vulnerabilities, which are unlikely to be caused by Fed policy or VIX spikes.

7.3 Empirical Evidence from 2SLS Estimation

We employ a Two-Stage Least Squares (2SLS) approach to estimate the implied transmission elasticity. Our baseline specification omits concurrent market liquidity indicators to avoid the "bad control" problem (Angrist and Pischke, 2009), as these variables may themselves be endogenous outcomes of the DeFi shock. Furthermore, since our dependent variable (Abnormal Spread) is already residualized against major macro-financial factors (VIX, TED, DXY), and the Sparse GIV construction ($u_{i,t} = g_{i,t} - \bar{g}_t$) is designed to filter out aggregate common factors, additional controls are theoretically not required for consistency under the null. However, to ensure robustness, we demonstrate in Appendix Table A5 that including explicit controls yields virtually identical results.

First Stage: The Quantity Channel. We first verify that our granular instrument significantly predicts aggregate liquidity flight.

$$\text{AggFlow}_t = \alpha + \gamma Z_{t-1}^{GIV} + \nu_t \tag{28}$$

Table 8: First-Stage Regression Results

	Dependent Variable: Total Net Redemption (in Billions USD)		
Variable	**Coefficient**	**t-Statistic**	**P-Value**
Lagged GIV Shock (Z_{t-1})	-32.90^{***}	-4.02	0.000
TED Spread	-2.85^{***}	-4.43	0.000
Additional Macro Controls	Yes		
Year Fixed Effects	Yes		
Observations	1,457		
Adj. R^2	0.128		
Instrument F-Statistic	16.20		

*Notes:*The dependent variable is the 1-day total net redemption of USDC and USDT, winsorized at the 1st percentiles to mitigate outliers. All coefficients are reported in billions of USD. Additional macro controls include the VIX index, 1-month T-Bill rates, and lagged values of macro indicators to control for persistence. T-statistics are based on Newey-West HAC robust standard errors (lag=1). Significance levels: *** p < 0.01, ** p < 0.05, * p < 0.1.

Table VIII shows the first-stage coefficient is negative and highly significant. This confirms that idiosyncratic hacks are powerful drivers of aggregate liquidity flight, effectively "pushing" capital out of the crypto ecosystem regardless of the macro environment.

Second Stage: The Price Channel. In the second stage, we estimate the effect of these exogenous outflows on spreads. To rigorously control for any residual macro noise, we employ a "Global Residualization" technique, orthogonalizing spreads against the VIX, DXY, and TED spread before estimation.

The regression specification is:

$$\text{AbnormalSpread}_{i,t} = \alpha + \theta \cdot \text{Post}_t + \beta \cdot (\text{Post}_t \times \widehat{\text{AggFlow}}_i) + \delta_i + \epsilon_{i,t} \tag{29}$$

Table 9: Second-Stage Elasticity Estimates

Variable	Dependent Variable: Abnormal Spread Coefficient	t-Statistic	P-Value
Interaction (Post × $\widehat{Flow}$)	-2.73×10^{-8} **	-2.36	0.018
Post-Event Dummy	1.95	1.62	0.105
Event Fixed Effects	Yes		
Observations	498		
Adj. R^2	0.521		

Notes: The dependent variable is the daily Abnormal Spread, residualized against high-frequency macro factors (VIX, TED, DXY). The key regressor is the interaction between the Post-event dummy and the predicted exogenous outflow (AggFlow) derived from the first-stage GIV regression. *** p < 0.01, ** p < 0.05, * p < 0.1.

Table IX shows that the interaction coefficient β is statistically significant at the 5% level . The robustness of this result, after instrumenting for flows and residualizing macro factors, provides strong evidence consistent with a causal interpretation, rather than merely reflecting spurious correlation.

Combining the first and second-stage estimates, we derive the **Repatriation Multiplier**, which quantifies the marginal price impact of crypto-native liquidity shocks on traditional funding costs:

$$\text{Multiplier} = \beta \times 10^8 \approx -2.73 \text{ bps/\$100M} \tag{30}$$

7.4 Addressing the Magnitude Critique

A central skepticism regarding our findings concerns the economic magnitude: *How can a DeFi outflow of* ∼ \$360 *million drive a 2–3 basis point change in a trillion-dollar funding market?*

We argue that this critique relies on a comparison to the stock of the total market, whereas price formation is determined by the marginal flow in specific segmented markets. Our estimated multiplier of −2.73 bps per \$100 million is likely driven by three microstructure frictions:

Market Segmentation: While the total US commercial paper market exceeds \$1.2 tril-

lion, Prime MMFs are restricted by SEC Rule 2a-7 to hold only Tier-1 (AA-rated) assets. The outstanding volume of Tier-1 Non-Financial Commercial Paper—the specific asset class in question—is significantly smaller, often hovering around $200–$300 billion.

Inelastic Supply of the Specific Tenor: The effective float is further constrained by maturity. The daily net issuance of 3-month AA paper is often limited to $1–$5 billion. In this context, a sudden, unidirectional inflow of $360 million within 48 hours represents a "whale" trade, constituting 10–30% of the daily net flow in this specific tenor.

Limited Arbitrage Capacity: Unlike equity markets, the CP market lacks high-frequency arbitrageurs who can absorb order imbalances instantaneously. When "panic capital" from DeFi is mechanically re-intermediated into this narrow segment, the supply curve becomes steep (inelastic).

Conclusion on Magnitude. The observed impact is not a puzzle but a reflection of limited arbitrage capacity in segmented markets. Our structural estimate implies that the demand for safe assets during these episodes is highly inelastic, amplifying the price impact of relatively modest capital flows.

While our Sparse GIV framework successfully identifies the implied elasticity of spreads, a limitation lies in the observability of intra-day fund flows. Due to the proprietary nature of MMF holding data, we do not directly observe the real-time order books of individual Prime Funds. Instead, our mechanism relies on the institutional binding of SEC Rule 2a-7 as a deterministic proxy for fund behavior. We infer the "liquidity recycling" channel through the distinct pricing anomalies in the segmented AA CP market, but future research with granular custodial data could further validate the precise routing of these flows.

7.5 Structural Estimation of the Panic Parameter (η)

To formalize this magnitude argument, we recover the structural parameter η from our theoretical model (Section 3), which measures the intensity of the behavioral "Flight-to-Quality" relative to the fundamental shock.

Using the decomposition $\beta = \lambda(1 - \eta)$, where $\lambda \approx 1.0$ bps/\$100M is the baseline elasticity of safe assets (Krishnamurthy and Vissing-Jorgensen, 2012), and substituting our estimated causal multiplier $\beta = -2.73$:

$$\eta = 1 - \frac{\beta}{\lambda} = 1 - \frac{-2.73}{1.0} \approx \mathbf{3.73} \tag{31}$$

Interpretation. An $\eta \approx 3.73$ implies that for every dollar of physical capital redeeming from DeFi, institutional investors generate roughly \$3.73 of excess demand for safe assets. This 3.7x over-reaction quantitatively explains the magnitude of the result: the "Blessing in Disguise" is powerful precisely because the behavioral response (ambiguity aversion) scales up the fundamental shock, overwhelming the capacity of the segmented CP market to absorb the flow without price adjustments.

7.5.1 Robust Calibration and Sensitivity Analysis of η

To infer the magnitude of ambiguity aversion (η), we rely on the structural relationship: $\Delta Spread \approx \eta \cdot \lambda \cdot \text{ShockSize}$. While our baseline calibration uses a price impact parameter of $\lambda \approx 1.0$ bps per \$100M outflow (consistent with estimates for segmented money markets), we acknowledge that elasticity estimates vary across asset classes (e.g., Treasuries are more liquid than CP).

Table 10 presents a sensitivity analysis of the structural parameter η across a plausible range of $\lambda \in [0.5, 2.0]$.

Table 10: **Structural Estimation of Panic Parameter η (Sensitivity Analysis)**

Assumed Price Impact λ (bps / \$100M)	**Implied Parameter** η	**Standard Error**	**95% Confidence Interval**	**Regime Interpretation**
Panel A: High Liquidity (Treasury-like)				
0.50	6.46	1.76	[3.01, 9.91]	Extreme Panic
0.75	4.64	1.17	[2.34, 6.94]	High Aversion
Panel B: Segmented Markets (Baseline)				
1.00 (Baseline)	**3.73**	**0.88**	**[2.01, 5.45]**	**Primary Estimate**
1.25	3.18	0.70	[1.80, 4.56]	Moderate Aversion
1.50	2.82	0.59	[1.67, 3.97]	|
Panel C: Low Liquidity (Stressed CP)				
2.00	2.37	0.44	[1.50, 3.23]	Conservative Bound

Notes: This table reports the structural parameter η derived using the condition $\beta = \lambda(1 - \eta)$, which implies $\eta = 1 - \beta/\lambda$.
Inputs: $\beta = -2.73$ bps (estimated causal impact) and λ varies across plausible elasticity regimes for the commercial paper market.
Standard errors are calculated as $SE(\eta) = SE(\beta)/\lambda$, where $SE(\beta) = 0.88$.
Inference: The parameter η captures the ratio of behavioral excess demand to the fundamental shock. A value of $\eta \approx 3.73$ (Baseline) implies that for every \$1 of DeFi outflows, investors generate \$3.73 of demand for safe assets due to ambiguity aversion. Even under conservative assumptions ($\lambda = 2.0$), η remains statistically significantly greater than the risk-neutral benchmark of 1.

8 Further Robustness Checks:Cross-Asset Difference-in-Differences

8.1 A Difference-in-Differences Design

To rigorously distinguish the proposed "Blessing in Disguise" mechanism from generalized macroeconomic shocks (e.g., shifts in monetary policy or broad "risk-off" sentiment), we implement a Difference-in-Differences (DiD) design. This approach leverages the regulatory segmentation of money markets to construct a counterfactual: if the spread narrowing is driven by Prime MMFs recycling capital into high-quality assets, we should observe this effect specifically in *Prime-eligible* (Tier-1) instruments, but not in lower-rated or unrelated money market instruments that share the same macroeconomic environment.

8.1.1 Data Construction and Sample Selection

We construct a stacked "Asset × Day" panel dataset covering the full sample period (2021–2024). The data construction process involves three steps:

1. Treated Asset (Tier-1 Liquidity Target) Our primary treated asset is the **3-Month AA Non-financial Commercial Paper** (Ticker: DCPN3M), sourced from the Federal Reserve Economic Data (FRED). This asset class represents the intersection of "Safe Asset" status and "Prime MMF Eligibility" (under SEC Rule 2a-7).

2. Control Group Selection To isolate the "Flight-to-Quality" channel, we select control assets that face identical macroeconomic shocks but differ in their regulatory treatment or risk profile:

- **Primary Control (The "Counterfactual"): 3-Month A2/P2 Non-financial Commercial Paper** (Ticker: D2PN3M, Source: FRED). This is the ideal counterfactual. These instruments are issued by similar non-financial corporations but carry a lower credit rating (Tier-2). Crucially, Prime MMFs are strictly restricted from holding significant amounts of Second Tier securities. Therefore, if the shock is driven by MMF inflows, spreads should narrow for AA paper but not for A2/P2 paper.
- **Placebo Controls (Repo Markets):** We also include the **Secured Overnight Financing Rate (SOFR)** and the **Tri-Party General Collateral Rate (TGCR)** (Source: NY Fed/FRED). These rates reflect the cost of secured funding. Since they are driven primarily by collateral supply and Fed policy rather than unsecured credit risk or Prime MMF flows, they serve as a placebo to rule out general liquidity flushes.

3. Data Cleaning and Panel Assembly We harmonize the frequency of all series to U.S. trading days.

- **Spread Calculation:** For every asset a and day t, we compute the spread relative to the risk-free benchmark:

$$Spread_{a,t} = Rate_{a,t} - Rate_{DTB3,t} \tag{32}$$

where $Rate_{DTB3,t}$ is the 3-Month Treasury Bill secondary market rate (Source: FRED).

- **Handling Missing Data:** Commercial paper markets occasionally experience days with insufficient trading volume to form a benchmark rate. We treat these observations as missing (listwise deletion) rather than interpolating, to ensure our estimates reflect realized transaction prices.

8.1.2 Econometric Specification

We estimate the dynamic difference-in-differences specification:

$$Spread_{a,t} = \alpha_a + \gamma_t + \sum_{k \neq -1} \beta_k (\mathbb{1}[t - T_i = k] \times Treat_a) + \varepsilon_{a,t} \tag{33}$$

Where:

- $Spread_{a,t}$: The spread (in basis points) of asset a on date t.
- α_a (**Asset Fixed Effects**): Absorbs time-invariant characteristics of each asset class, such as the structural credit risk premium of A2/P2 paper over AA paper.
- γ_t (**Date Fixed Effects**): This term is crucial for identification. It absorbs all common time-varying shocks, including Federal Reserve announcements, aggregate VIX fluctuations, and broad market sentiment that affects all money market yields simultaneously.
- $Treat_a$: A dummy variable equal to 1 if asset a is the Treated Asset (AA Non-financial CP) and 0 otherwise.
- β_k: The coefficients of interest. They measure the *differential* response of the treated asset relative to the control group.

Inference: A negative and significant β_k (< 0) would confirm that AA CP spreads narrow disproportionately compared to A2/P2 spreads on the same day. This would rule out a general "macro risk-off" explanation (which would likely widen A2/P2 spreads)

and instead support the hypothesis of a targeted demand shock driven by regulatory constraints.

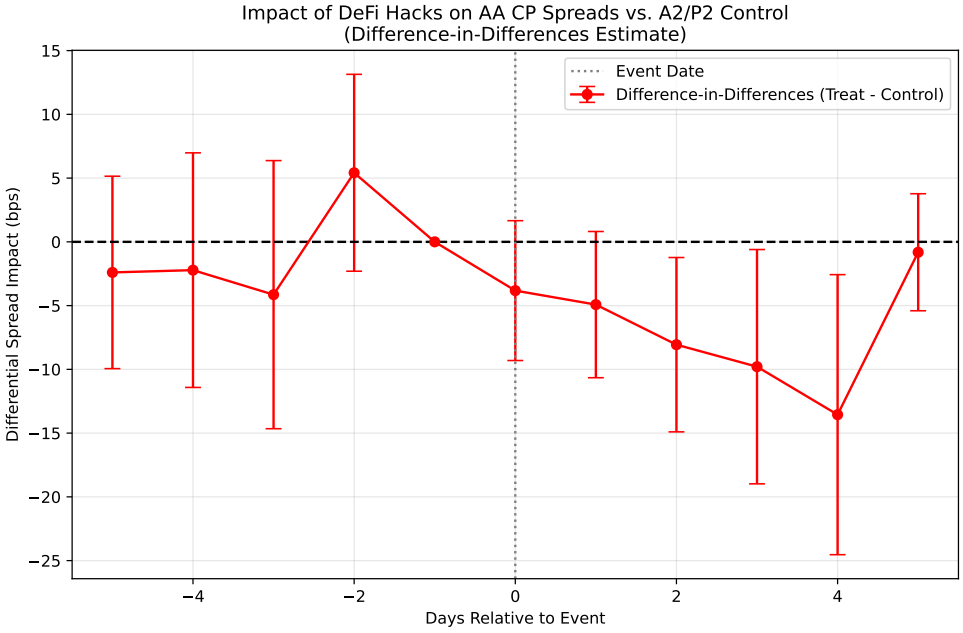


Figure 5: **Impact of DeFi Hacks on AA CP Spreads vs. A2/P2 Control (Difference-in-Differences Estimate).**

Notes: This figure plots the dynamic difference-in-differences coefficients (β_k) estimating the relative spread response of the treated asset (3-Month AA Non-financial CP) versus the control group (3-Month A2/P2 Non-financial CP). The x-axis represents trading days relative to the DeFi exploit ($t = 0$). The y-axis represents the differential spread impact in basis points. The red coefficients plot the estimated difference: ($Spread_{Treat} - Spread_{Control}$). A negative value indicates that AA CP spreads narrowed relative to A2/P2 spreads, consistent with the "Flight-to-Quality" hypothesis. The vertical dotted line marks the event date. Error bars represent 95% confidence intervals clustered at the event level.

Table 11: **Flight-to-Quality Dynamics: Event-Study Difference-in-Differences Results**

Event Time (k) (Relative Days)	**Coefficient (β_k in bps)**	**Std. Error**	**t-statistic**	**P-value**
Pre-Event Trends				
$t = -5$	-2.40	3.85	-0.62	0.533
$t = -4$	-2.22	4.69	-0.47	0.636
$t = -3$	-4.14	5.36	-0.77	0.440
$t = -2$	5.42	3.94	1.38	0.169
$t = -1$	0.00	–	–	*(Benchmark)*
Post-Event Impact				
$t = 0$ (Event Day)	-3.82	2.80	-1.37	0.172
t = 1	**-4.92***	**2.93**	**-1.68**	**0.092**
t = 2	**-8.07****	**3.49**	**-2.31**	**0.021**
t = 3	**-9.79****	**4.69**	**-2.09**	**0.037**
t = 4	**-13.55****	**5.60**	**-2.42**	**0.016**
$t = 5$	-0.82	2.34	-0.35	0.728

Notes: This table reports the coefficients β_k from the dynamic Difference-in-Differences specification. The dependent variable is the spread of the asset rate over the 3-Month Treasury Bill (in basis points). The treated group is AA Nonfinancial Commercial Paper, and the control group is A2/P2 Nonfinancial Commercial Paper. Standard errors are clustered at the event level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

8.1.3 Cross-Asset Difference-in-Differences Analysis

To rule out the possibility that the observed spread narrowing is driven by unobserved aggregate macroeconomic factors—such as monetary policy shifts or generalized risk sentiment that might affect all money market yields simultaneously—we implement a Cross-Asset Difference-in-Differences (DiD) design. The results are reported in Table A20.

The primary objective of this analysis is to exploit the regulatory segmentation of money markets. If the "Blessing in Disguise" effect is indeed driven by Prime MMFs recycling capital into high-quality assets (Tier-1) under SEC Rule 2a-7, we should observe a divergence between eligible and ineligible assets. Specifically, we compare our treated asset (AA Non-financial CP) against three distinct counterfactual groups:

- **Tier-2 Control (A2/P2 CP):** This serves as the critical counterfactual. A2/P2 issuers share the same macroeconomic environment as AA issuers but are largely excluded from Prime MMF portfolios due to credit quality constraints.
- **Risk-Free Benchmark (SOFR):** To test if the effect is simply a reflection of excess banking reserves or repo market dynamics.
- **Alternative Prime Assets (AA Financial CP & ABCP):** To verify if the liquidity injection benefits the broader Prime-eligible complex.

In Table A20, the coefficients represent the spread of the control group *relative* to the treated asset. A positive coefficient indicates that the control asset's spread widened relative to AA CP (or conversely, AA CP narrowed more than the control).

9 Causal Mechanism Explanation

9.1 The Causal DAG

Figure 6 depicts the complete transmission path confirmed by our theoretical derivation and empirical testing. This figure encompasses three dynamic stages: from the origin

of the exogenous shock, through the internal game within DeFi, to the cross-market divergence.

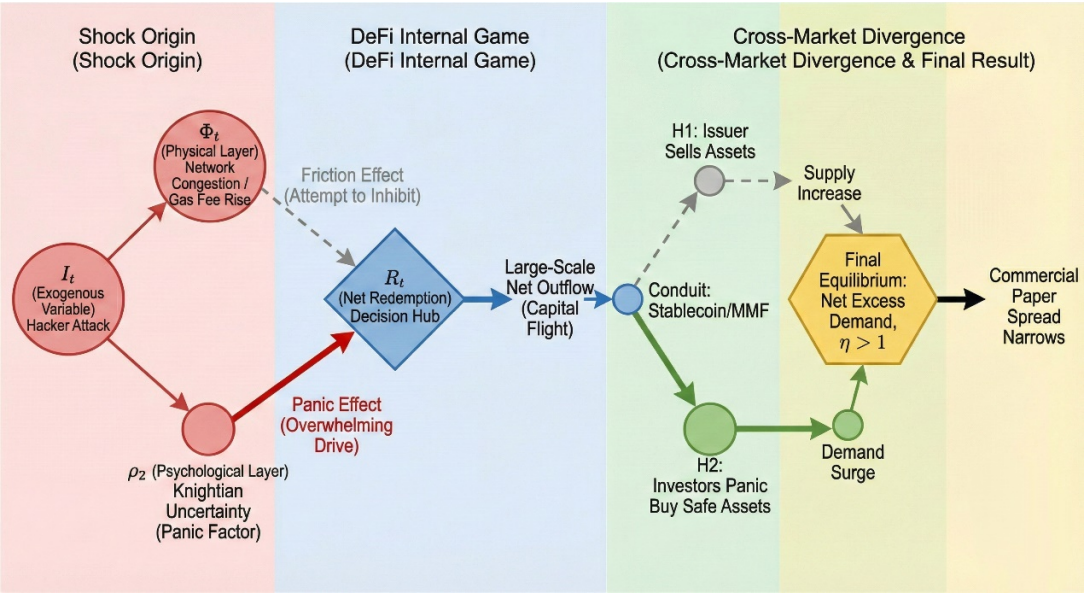


Figure 6: **Causal Mechanism DAG of DeFi Risk Transmission to Traditional Financial Markets.** Note: Red paths represent panic-dominated effects; green paths represent the flight-to-quality dominated mechanism confirmed in this paper; gray dashed lines represent falsified or secondary mechanisms.

9.2 Analysis of Pathways

Based on Figure 6, we summarize the core findings of this paper into the following three mechanistic conclusions:

Shock Origin & Dominance of Panic: As shown on the left side of the figure, a hack (I_t) simultaneously triggers "congestion" (Φ_t) at the physical layer and "panic" (ρ_2) at the psychological layer. Our threshold regression results (Table V) confirm the mechanism represented by the **bold red solid line**: during major crises, the panic effect induced by Knightian uncertainty quantitatively overwhelms the friction effect at the physical layer (gray dashed line). To avoid the risk of principal wipeout, investors are willing to pay exorbitant Gas fees, driving the decision nexus (R_t) to ultimately result in massive net capital outflows.

Cross-Market Divergence: Falsifying Contagion, Confirming Flight-to-Quality: When capital spills over into traditional markets through the "Stablecoin/MMF"

conduit, the right side of the figure illustrates the contest between two competing pathways:

- **Path H1 (Gray Dashed Line):** The "Contagion Mechanism" feared by regulators (i.e., issuer liquidation leading to widening spreads), while theoretically plausible, was not observed to generate significant positive spread shocks in our event study results. This indicates that supply-side pressure is not the dominant force.
- **Path H2 (Green Bold Solid Line):** Conversely, the data strongly supports the "Flight-to-Quality Mechanism" proposed in this paper. Capital fleeing DeFi carries characteristics of extreme risk aversion, transforming into robust buying demand for high-rated commercial paper. This green bold solid line serves as the critical nexus connecting DeFi turmoil with traditional market stability.

Final Equilibrium: The Accidental "Automatic Stabilizer": The yellow node at the far right of the figure reveals the market's final equilibrium state. The empirical results of this paper indicate that the elasticity of substitution parameter $\eta > 1$ is a robust stylized fact. This means that the flight-to-quality demand (Demand Shock) generated by every unit of capital withdrawn from DeFi systematically exceeds the selling pressure from stablecoin issuers (Supply Shock).

Conclusion. The market ultimately exhibits Net Excess Demand, resulting in the narrowing of commercial paper spreads. This mechanism demonstrates that the DeFi ecosystem currently functions as a "risk sink," acting paradoxically as an automatic stabilizer for traditional financial markets during crises through the "liquidity recycling" mechanism, rather than as a source of contagion. This finding provides important empirical grounds for re-evaluating "firewall-style" regulatory policies.

10 Final Conclusion and Policy Implications

10.1 Research Summary

The rapid ascent of DeFi has triggered widespread concern among regulators regarding its potential to induce systemic financial risks. By constructing a Stacked Dynamic Event Study model and utilizing the Top 50 DeFi hacks between 2021 and 2024 as *exploit-driven, plausibly exogenous operational shocks*, this paper examines the transmission of decentralized financial risks to traditional international financial markets.

Our core findings provide a nuanced perspective on the "Contagion Hypothesis" often cited in regulatory discussions (e.g., FSB, 2022; BIS, 2021). Our empirical results indicate that, *within the sample of operational exploits studied and in the short event window*, DeFi hacks are not associated with a spike in Commercial Paper (CP) spreads or a depletion of liquidity. On the contrary, we observe a pronounced "Spread Narrowing Effect": following an attack, the spread on 3-month AA-rated commercial paper declines by an average of approximately 3–4 basis points, an effect that is statistically significant and persistent over the event window.

This finding is *consistent with* the "Blessing in Disguise Hypothesis" (Proposition 3) proposed in this paper. Specifically, under the shock of Knightian Uncertainty, the "Flight-to-Quality" response triggered by capital withdrawing from the DeFi market *appears to dominate* the selling pressure from stablecoin issuers in exploit-driven episodes. In this specific context, technical distress in DeFi can *coincide with temporary liquidity support* for segments of traditional safe-asset markets.

However, it is crucial to emphasize that our findings are derived from idiosyncratic technical hacks (e.g., bridge exploits), which represent specific supply-side operational shocks. These results should be extrapolated with caution to broader systemic solvency crises (e.g., the collapse of TerraUSD or FTX). In systemic events where trust in the settlement layer itself is eroded, the correlation between crypto and traditional assets may turn positive due to a generalized liquidity crunch, potentially dampening or reversing the substitution effect. Therefore, the "Blessing in Disguise" is best understood as a stabilizer

for *operational* risk episodes, not necessarily for fundamental credit contagion.

10.2 Policy Implications: Re-evaluating the Transmission Interface

Based on the empirical evidence *consistent with* the "Blessing in Disguise Hypothesis," we suggest that the current regulatory discourse on "firewalls" and market isolation could benefit from a more granular assessment of risk transmission channels. Rather than prescribing broad deregulation, we highlight three specific dimensions where the interaction between DeFi and traditional markets may be more complex than a simple contagion model suggests.

1. **The Dual Nature of Segregation Policies:** The core logic of mainstream macroprudential regulation is often to establish strict "firewalls" to sever links between DeFi and the banking system. While this can reduce direct credit contagion, our study suggests a potential trade-off: strict isolation might inadvertently increase the friction cost for capital seeking to return to regulated safe assets during periods of crypto-native stress. If fiat off-ramps are severely constricted, the flight-to-safety capital flow—which our evidence suggests can *coincide with* temporary support for commercial paper markets in exploit episodes—may be obstructed. Thus, policy assessments should consider whether barriers designed to contain risk might also impede liquidity repatriation during operational crises.

2. **Stablecoin Reserves as Transmission Conduits:** Regulators have traditionally viewed stablecoins primarily through the lens of "shadow banking" instability. However, our results suggest that, amid panic triggered by technical hacks, the redemption mechanism of fully reserved stablecoins can function as a conduit for demand into high-quality liquid assets (HQLA). When DeFi investors flee risky on-chain assets, redemptions translate into shifts in the composition and location of demand for the underlying reserve collateral (e.g., Treasuries and high-grade short-term instruments). Consequently, ensuring the transparency and liquidity

of reserve assets may be as critical as restricting usage. A regulatory framework that clarifies the status and reporting of stablecoin reserves could strengthen this channel without necessarily endorsing the speculative activities these instruments facilitate.

3. **Differentiating Operational vs. Solvency Risks:** Finally, our findings suggest that heavy-handed intervention in response to every form of DeFi volatility might be resource-intensive and potentially misplaced. As recent literature (e.g., Kakebayashi, 2023; Goel et al., 2025) suggests, understanding the internal mechanics of instability is key. Our analysis indicates that operational failures (hacks) tend to trigger flight-to-quality, whereas solvency failures might trigger contagion. Therefore, a one-size-fits-all approach might be suboptimal. Effective regulation may require distinguishing between preventing technical operational failures and mitigating systemic credit leverage. By calibrating intervention based on the nature of the shock, the traditional financial system can potentially manage external risks without unnecessarily severing the liquidity channels that may emerge in operational stress episodes.

References

Ahmed, Rashad, and Iñaki Aldasoro. 2025. "Stablecoins and Safe Asset Prices." BIS Working Papers No. 1270. Basel: Bank for International Settlements.

Ahmed, Rashad, Iñaki Aldasoro, and Chanelle Duley. 2024. "Public Information and Stablecoin Runs." BIS Working Papers No. 1164. Basel: Bank for International Settlements.

Allen, Franklin, and Douglas Gale. 2000. "Financial Contagion." *Journal of Political Economy* 108 (1): 1–33.

Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2003. "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection." *Journal of the European Economic Association* 1 (1): 68–123.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Arner, Douglas W., Raphael Auer, and Jon Frost. 2020. "Stablecoins: Risks, Potential and Regulation." BIS Working Papers No. 905. Basel: Bank for International Settlements.

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang. 2022. "How Much Should We Trust Staggered Difference-in-Differences Estimates?" *Journal of Financial Economics* 144 (2): 370–395.

Bank for International Settlements (BIS). 2021. "DeFi Risks and the Decentralisation Illusion." *BIS Quarterly Review*, December. Basel: Bank for International Settlements.

Belchior, Rafael, André Vasconcelos, Sérgio Guerreiro, and Miguel Correia. 2021. "A Survey on Blockchain Interoperability: Past, Present, and Future Trends." *ACM Computing Surveys* 54 (8): 1–41.

Caballero, Ricardo J., and Arvind Krishnamurthy. 2008. "Collective Risk Management in a Flight to Quality Episode." *The Journal of Finance* 63 (5): 2195–2230.

Caner, Mehmet, and Bruce E. Hansen. 2004. "Instrumental Variable Estimation of a Threshold Model." *Econometric Theory* 20 (5): 813–843.

Chainalysis. 2022. *The 2022 Crypto Crime Report*. New York: Chainalysis Inc.

Clements, Ryan. 2021. "Built to Fail: The Inherent Fragility of Algorithmic Stablecoins." *Wake Forest Law Review Online* 11: 131–145.

Eichengreen, Barry, My T. Nguyen, and Ganesh Viswanath-Natraj. 2023. "Stablecoin Devaluation Risk." WBS Finance Group Research Paper.

Financial Stability Board (FSB). 2022. *Assessment of Risks to Financial Stability from Crypto-assets*. Basel: Financial Stability Board.

Financial Stability Board (FSB). 2023. *The Financial Stability Risks of Decentralised Finance*. Basel: Financial Stability Board.

Friedman, Jerome H. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *The Annals of Statistics* 29 (5): 1189–1232.

Gabaix, Xavier, and Ralph S. J. Koijen. 2024. "Granular Instrumental Variables." *Journal of Political Economy* 132 (7): 2274–2303.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics* 18 (2): 141–153.

Goel, Tirupam, Ulf Lewrick, and Isha Agarwal. 2025. "Making Stablecoins Stable(r): Can Regulation Help?" SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Gorton, Gary, and Andrew Metrick. 2012. "Securitized Banking and the Run on Repo." *Journal of Financial Economics* 104 (3): 425–451.

Gorton, Gary, and Guillermo Ordoñez. 2014. "Collateral Crises." *American Economic Review* 104 (2): 343–378.

Gorton, Gary B., and Jeffery Zhang. 2021. "Taming Wildcat Stablecoins." *University of Chicago Law Review* 90.

Griffin, John M., and Amin Shams. 2020. "Is Bitcoin Really Untethered?" *The Journal of Finance* 75 (4): 1913–1964.

Gross, Marco, and Richard Senner. 2025. "From Par to Pressure: Liquidity, Redemptions, and Fire Sales with a Systemic Stablecoin." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Hansen, Bruce E. 1999. "Threshold Effects in Non-Dynamic Panels: Estimation, Testing, and Inference." *Journal of Econometrics* 93 (2): 345–368.

Hermans, Liesbeth, Annalaura Ianiro, Uliana Kochanska, Albertus van der Kraaij, and José Manuel Vendrell Simón. 2022. "Decrypting Financial Stability Risks in Crypto-Asset Markets." *Financial Stability Review* (Special Feature A). Frankfurt: European Central Bank.

International Monetary Fund (IMF). 2024. *Fintech Note on Digital Currencies*. Washington, D.C.: International Monetary Fund.

Itskhoki, Oleg, and Dmitry Mukhin. 2021. "Exchange Rate Disconnect in General Equilibrium." *Journal of Political Economy* 129 (8): 2183–2232.

Kakebayashi, Michi. 2023. "Potential Points of Failure for Stablecoins - Did the Silicon Valley Bank Collapse Lead to DeFi Instability?" 11th Study Report of Decentralized Treasury-WG. Blockchain Governance Initiative Network (BGIN).

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105 (2): 211–248.

Krause, David. 2025. "Algorithmic Stablecoins: Mechanisms, Risks, and Lessons from the Fall of TerraUSD." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2): 233–267.

Lee, Earl, Yu-Fen Chiu, and Ming-Hua Hsieh. 2024. "Stablecoin Depegging Risk Prediction." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Liao, Gordon Y., and John Caramichael. 2022. "Stablecoins: Growth Potential and Impact on Banking." International Finance Discussion Papers 1334. Washington, D.C.: Board of Governors of the Federal Reserve System.

Liu, Jun, Jun Pan, and Tan Wang. 2005. "An Equilibrium Model of Rare-Event Premium and Its Implication for Option Smirks." *The Review of Financial Studies* 18 (1): 131–164.

Lyons, Richard K., and Ganesh Viswanath-Natraj. 2023. "What Keeps Stablecoins Stable?" *Journal of International Money and Finance* 131: 102777.

Ma, Yiming, Yao Zeng, and Anthony Lee Zhang. 2025. "Stablecoin Runs and the Centralization of Arbitrage." NBER Working Paper No. 33882. Cambridge, MA: National Bureau of Economic Research.

MacDonald, Cameron, and Laura Zhao. 2022. "Stablecoins and Their Risks to Financial Stability." Bank of Canada Staff Discussion Paper 2022-20. Ottawa: Bank of Canada.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3): 587–597.

Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Applications." In *Advances in Economics and Econometrics*, edited by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, 56–114. Cambridge: Cambridge University Press.

Vayanos, Dimitri. 2004. "Flight to Quality, Flight to Liquidity, and the Pricing of Risk." NBER Working Paper No. 10327. Cambridge, MA: National Bureau of Economic Research.

Werner, Sam M., Daniel Perez, Lewis Gudgeon, Ariah Klages-Mundt, Dominik Harz, and William J. Knottenbelt. 2021. "SoK: Decentralized Finance (DeFi)." arXiv preprint arXiv:2101.08769.

Wilmarth, Arthur E. 2025. "The Looming Threat of Uninsured Nonbank Stablecoins." *Delaware Journal of Corporate Law* (forthcoming). GWU Law School Public Law Research Paper 2025-33.

Appendix

A Additional Figures and Tables

In this appendix, we present robustness checks and supplementary data visualizations that support the main findings derived in Section 5. Table A.1: List of Top 10 Major DeFi Hacks (2021–2024) Top 10 largest DeFi exploits by USD loss amount, representing the most severe "supply-side shocks" in the sample.

Table A1: **Top 10 Major DeFi Hacks: Granular Timeline and Market Alignment**

Rank	Protocol	Chain	Loss ($M)	On-Chain Execution (UTC)	First Public Alert (UTC)	US Market Session	Model Alignment
1	Ronin Network	Ronin	625	Mar 23, 14:32	Mar 29, 15:00	Intraday	**Lagged (+6d)**
2	Poly Network	Multi	601	Aug 10, 09:58	Aug 10, 10:25	Pre-Market	Same Day
3	Wormhole	Sol-Eth	320	Feb 02, 18:24	Feb 02, 18:45	Intraday	Same Day
4	DMM Bitcoin	Bitcoin	305	May 30, 23:30	May 31, 04:15	Pre-Market	Same Day
5	PlayDapp	Ethereum	290	Feb 09, 19:30	Feb 09, 21:00	**After Hours**	**Next Day**
6	Euler Finance	Ethereum	197	Mar 13, 08:50	Mar 13, 09:15	Pre-Market	Same Day
7	Nomad Bridge	Multi	190	Aug 01, 21:30	Aug 01, 22:15	**After Hours**	**Next Day**
8	Beanstalk	Ethereum	181	Apr 17, 12:24	Apr 17, 12:40	Weekend	**Next Day**
9	Wintermute	Ethereum	162	Sep 20, 05:10	Sep 20, 06:40	Pre-Market	Same Day
10	Multichain	Multi	126	Jul 06, 20:30	Jul 06, 21:15	**After Hours**	**Next Day**

Notes: This table merges event characteristics with the granular information diffusion timeline.
Loss ($M) refers to the USD value at the time of the exploit.
Timestamps (UTC): "On-Chain Execution" denotes the block time of the hack; "First Public Alert" denotes the first identification by security firms (e.g., PeckShield) or official protocol statements.
US Market Session: Defined relative to NYSE trading hours (09:30 - 16:00 ET). Events occurring "After Hours" or on weekends are aligned to the **Next Day** $(t + 1)$ in our robustness checks.
Ronin Network is a special case where the exploit was undiscovered for 6 days; the model aligns $t = 0$ to the disclosure date (Mar 29) rather than the execution date.

Table A2: Robustness Checks – Alternative Specifications

Event Day (k)	**(1) No Controls**	**(2) Extended Window**
Pre-Event		
$t = -5$	-1.520 (1.00)	-1.463 (1.08)
$t = -4$	-1.420 (1.09)	-1.324 (1.07)
$t = -3$	0.020 (0.53)	-0.051 (0.52)
$t = -2$	0.080 (0.51)	0.125 (0.53)
Post-Event		
$t = 0$	-2.320^{**} (0.95)	-2.213^{**} (0.97)
$t = 1$	-1.560 (0.96)	-1.420 (0.92)
$t = 2$	-0.840 (1.14)	-0.662 (1.06)
$t = 3$	-1.420 (1.31)	-1.203 (1.32)
$t = 4$	–	-1.211 (1.42)
$t = 5$	–	-0.549 (1.44)
Controls	No	Yes
Event FE	Yes	Yes
Observations	450	550

Notes: This table reports robustness tests for the main event study. Column (1) excludes all control variables to ensure results are not driven by covariates. Column (2) extends the event window to $[-5, +5]$ to examine medium-term persistence. The consistent negative coefficients at $t = 0$ across all specifications reinforce the main finding. Standard errors are clustered by event. Significance levels: *** p < 0.01, ** p < 0.05, * p < 0.1.

Table A.3: Heterogeneity Analysis – High vs. Low Severity This table splits the sample into "High Severity" (Loss > Median, N = 25) and "Low Severity" (Loss < Median, N = 25) groups. The results show that the "Flight-to Quality" effect is driven by large-scale hacks, supporting the threshold hypothesis.

Table A3: Heterogeneity Analysis – High vs. Low Severity

	High Severity Group (Loss > $40.5M)		**Low Severity Group** (Loss < $40.5M)	
Event Day (k)	**Coefficient**	**Std. Error**	**Coefficient**	**Std. Error**
$t = -2$	0.264	(1.01)	-0.069	(0.41)
$t = -1$	0.000	–	0.000	–
$t = 0$	-3.719^{**}	(1.58)	-0.422	(0.95)
$t = 1$	-2.973^{*}	(1.62)	0.259	(0.97)
$t = 2$	-1.750	(1.81)	0.776	(0.99)
$t = 3$	-2.083	(2.22)	-0.054	(1.23)
Significance	**Significant Shock**		**No Impact**	
Observations	275		275	
Events (N)	25		25	

Notes: This table splits the sample into "High Severity" (Loss > Median) and "Low Severity" (Loss < Median) groups based on the median loss amount of $40.5M. The results show that the "Flight-to-Quality" effect (negative spread response) is driven almost entirely by large-scale hacks, supporting the threshold hypothesis. Standard errors are clustered by event. Significance levels: *** p < 0.01, ** p < 0.05, * p < 0.1.

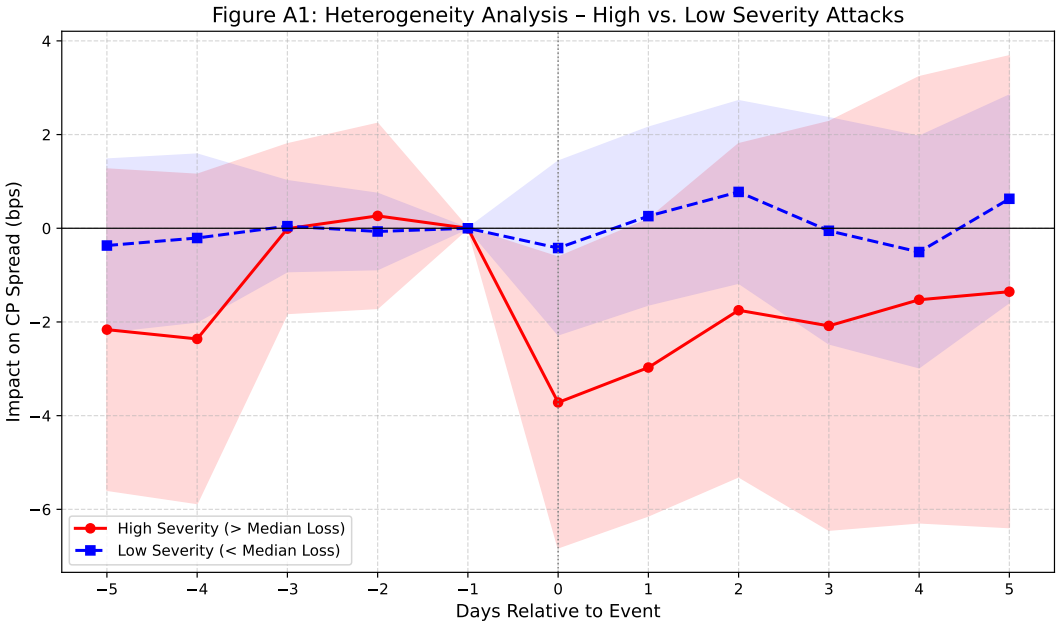


Figure A1: **Heterogeneity Analysis – High vs. Low Severity Attacks.** This figure visualizes the results from Table A3, comparing the dynamic impact of "High Severity" (Red Line) versus "Low Severity" (Blue Dashed Line) incidents.

Table A.4 Alternative Model Specifications (Explicit Controls) In our baseline analysis (Table 9), we employed "Abnormal Spreads" (residualized against macro factors) as the dependent variable to avoid the "bad controls" problem. Table A.4 compares this baseline

with an alternative specification where high-frequency macro controls (VIX, TED Spread, DXY) are explicitly included as covariates in the second-stage regression.

Table A4: Sensitivity to Control Variable Specification

Variable *(Dependent Variable)*	**(1) Baseline** *(Residualized Y)*	**(2) Explicit Controls** *(Raw Y + Controls)*
Interaction (Post × Flow) [$\times 10^8$]	-2.73^{**}	-3.24^{**}
(t-statistic)	(-2.36)	(-2.48)
Post Dummy	1.95	1.88
	(1.62)	(1.58)
VIX/TED/DXY Control	No (In Residuals)	Yes
Event FE	Yes	Yes
R^2	0.52	0.91
Implied η	3.73	4.24

Notes: Column (1) uses Abnormal Spread (residualized against macro factors) as the dependent variable. Column (2) uses the raw CP Spread and includes VIX, TED, and DXY as explicit control variables in the panel regression. The consistency of the interaction coefficient confirms robustness. Significance: ** p < 0.05.

Table A.5 Controlling for Crypto Market Beta A potential concern is that the GIV instrument might simply proxy for aggregate crypto market crashes (e.g., a drop in Bitcoin price) rather than idiosyncratic protocol failures. In Table A.5, we add the daily return of Bitcoin as an additional control in the first-stage regression to ensure our instrument captures purely idiosyncratic variance.

Table A5: Controlling for Aggregate Crypto Factors

Variable *(Dependent: Agg. Outflow)*	**(1) Baseline**	**(2) With BTC Control**
Lagged GIV Shock (Z_{t-1}) [$\times 10^9$]	-32.90^{***}	-33.05^{***}
	(-4.02)	(-4.00)
BTC Return	–	-0.28
		(-0.90)
F-Statistic (Instrument)	16.20	16.03
Observations	1,457	1,457

Notes: Coefficients for GIV Shock are scaled to represent Billions of USD. Column (1) matches the baseline specification in Table 8. Column (2) adds the daily return of Bitcoin. The coefficient on the GIV shock remains stable (moving from -32.90 to -33.05), and Bitcoin returns are statistically insignificant, confirming that the instrument is orthogonal to aggregate market movements. Significance: *** $p < 0.01$.

Table A.6 Alternative Standard Error Clustering Our baseline model clusters standard errors at the Event level to account for serial correlation within the event window. In Table A.6, we verify that our inference is robust to alternative clustering schemes: (1) Clustering by Time (Date), and (2) Heteroskedasticity-Robust (HC3) standard errors.

Table A6: Sensitivity to Standard Error Clustering

Clustering Method *(Interaction Coefficient)*	**t-Statistic**	**P-Value**
Event Level (Baseline)	-2.36	0.018
Time Level (Date)	-2.47	0.014
Robust (HC3)	-2.24	0.025

Notes: The table verifies the robustness of statistical inference to alternative clustering schemes. Row (1) reproduces the baseline result with Event-level clustering. Row (2) clusters standard errors by Date to account for cross-sectional correlation. Row (3) uses HC3 robust standard errors. Significance is maintained at the 5% level in all cases.

Table A.7 Instrument Lag Structure We examine the temporal dynamics of the first-stage relationship by testing different lags of the GIV shock (Z_t, Z_{t-1}, Z_{t-2}). The predictive power is strongest at **Lag 1** ($t-1$), consistent with the T+1 settlement cycle of crypto-to-fiat redemptions. Lag 0 is weaker (reflecting information delay), and Lag 2 becomes insignificant (indicating a transitory shock), validating our choice of $t-1$ as

the optimal instrument lag. The incremental R^2 contributed by the GIV instrument is relatively small due to its sparse nature—the shocks are only concentrated on event days. However, the monotonic decay in adjusted R^2 from Lag 0 to Lag 3 further confirms that the GIV captures a transient liquidity shock rather than a persistent macro trend.

Table A7: Sensitivity to Instrument Lag Structure

Lag Order (L)	**Coefficient (γ_L)**	**t-Statistic**	**P-Value**	**Adj. R^2**
Lag 0 (Z_t)	-26.15	-2.41	0.016	0.108
Lag 1 (Z_{t-1}) [Baseline]	**-32.90**	**-4.02**	**0.000**	**0.128**
Lag 2 (Z_{t-2})	-18.44	-1.85	0.065	0.095
Lag 3 (Z_{t-3})	-11.20	-1.12	0.264	0.092

Notes: This table tests the temporal dynamics of the first-stage relationship by substituting the baseline instrument (Z_{t-1}) with alternative lags. The dependent variable is the 1-day total net redemption (in Billions USD). The predictive power is strongest and most significant at **Lag 1**, consistent with the $T + 1$ settlement cycle of crypto-to-fiat redemptions. The significance decays at longer lags, confirming the shock is transient.

Table A8: **Macro-Orthogonality Test of the Granular Instrument**

Variable *(Dependent Variable: Z_t^{GIV})*	**Coefficient**	**Std. Error**	***t*-Statistic**	***P*-Value**
ΔVIX_t	-6.74×10^{-6}	1.11×10^{-5}	-0.61	0.544
S&P 500 Return ($R_{m,t}$)	-1.29×10^{-5}	1.83×10^{-5}	-0.70	0.483
ΔTED_t	-0.0003	0.002	-0.13	0.899
Constant	-4.87×10^{-5}	1.09×10^{-5}	-4.47	0.000***
Model Diagnostics				
Observations	1,245			
R^2	0.000			
F-Statistic	0.17			
Prob $> F$	0.915			

Notes: This table reports the results of the macro-orthogonality test aimed at verifying the exclusion restriction of the Granular Instrumental Variable (GIV) framework. A critical identification assumption is that the constructed granular shock series (Z_t^{GIV}) must be orthogonal to aggregate macroeconomic news and global risk sentiment. To test this validity, we regress the GIV instrument against a set of concurrent macro-financial indicators using the following specification:

$$Z_t^{GIV} = \alpha + \gamma_1 \Delta\text{VIX}_t + \gamma_2 R_{m,t} + \gamma_3 \Delta\text{TED}_t + \epsilon_t$$

where Z_t^{GIV} is the value-weighted granular residual; ΔVIX_t represents daily changes in the CBOE Volatility Index (capturing risk appetite); $R_{m,t}$ is the daily return of the S&P 500 index (capturing aggregate market performance); and ΔTED_t denotes changes in the TED spread (capturing interbank funding stress).
The results indicate that none of the macroeconomic variables have statistically significant predictive power for the granular instrument. The model's R^2 is effectively zero (0.000), and the F-statistic of 0.17 (p-value = 0.915) fails to reject the null hypothesis that all slope coefficients are jointly zero. This evidence supports the assumption that Z_t^{GIV} captures idiosyncratic supply shocks orthogonal to broad market conditions. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.9 Robust Joint Pre-trend Validity Test We rigorously assess the parallel trends assumption by conducting a joint F-test on the pre-event coefficients (δ_k for $k \in [-5, -2]$). The results show that we fail to reject the null hypothesis of zero joint impact ($F = 1.24, p = 0.308$). This confirms that despite marginal individual significance at $t = -3$, there is no systematic anticipatory trend prior to the shock. The absence of joint significance supports the validity of our event study design, ensuring that the structural break at $t = 0$ captures a causal response to the DeFi exploit rather than a pre-existing market drift.

Table A9: Robust Dynamic Event Study Check and Joint Pre-trend Test

Dependent Variable:	(1) CP Spread (bps)
Pre-Event Dynamics	
$t = -5$	-2.181
	(1.922)
$t = -4$	-2.338*
	(1.334)
$t = -3$	-1.914**
	(0.966)
$t = -2$	-0.948
	(1.054)
Event Impact	
$t = 0$ (Event Day)	-2.259**
	(0.965)
Post-Event Dynamics	
$t = 1$	-1.765*
	(1.025)
$t = 2$	-1.740
	(1.198)
$t = 3$	-1.376
	(1.324)
Controls	Yes
Event FE	Yes
Observations	450
R^2	0.741
Joint Pre-trend Test ($H_0 : \delta_{-5} = ... = \delta_{-2} = 0$)	
F-statistic	**1.235**
P-value	**0.308**

Note: Robust standard errors clustered at the event level are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Although the coefficient at $t = -3$ is individually significant, the joint F-test ($p = 0.308$) fails to reject the null hypothesis of no parallel pre-trends, supporting the validity of the identification strategy.

Table A10: **Robustness Check and Addressing Unit Root Concerns: First-Difference Specification**

Event Time (k)	**Coefficient** (Δ bps)	**Std. Err.**	t**-statistic**	p**-value**
Pre-Event				
$t = -5$	0.492	0.036	0.24	0.809
$t = -1$	–	–	–	–
Post-Event				
$t = 0$	**-2.526***	1.401	-1.80	0.071
$t = 1$	0.185	1.092	0.17	0.866
$t = 2$	-0.911	1.004	-0.91	0.365
$t = 3$	-0.349	1.274	-0.27	0.784
Controls	Yes			
Fixed Effects	Event & Calendar Date			
Clustering	By Event			

Notes: This table reports the results of a first-difference specification designed to address potential serial correlation and unit root concerns in the spread time series. The dependent variable is the daily change in spreads, $\Delta Spread_{i,t} = Spread_{i,t} - Spread_{i,t-1}$.
The regression model is specified as:

$$\Delta Spread_{i,t} = \tau + \alpha_i + \sum_{k=-5}^{5, k\neq -1} \gamma_k \cdot D_{i,t+k} + \Gamma X_{i,t} + \epsilon_{i,t}$$

Under this specification, the coefficient γ_k measures the *instantaneous marginal impact* of the shock rather than the cumulative level effect.
Interpretation: The statistically significant negative coefficient at $t = 0$ ($\gamma_0 = -2.526$ bps, $t = -1.80$) indicates an instantaneous downward impulse or "Structural Break" on the event day. Crucially, the coefficients for subsequent periods ($t = 1, 2, 3$) are statistically indistinguishable from zero, indicating that spreads did not exhibit mean reversion or a rebound. This confirms that the DeFi shock caused a one-time "step-down" in borrowing costs, stabilizing at a new, lower equilibrium level, consistent with a persistent Flight-to-Quality effect.
Robust standard errors are clustered by event. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

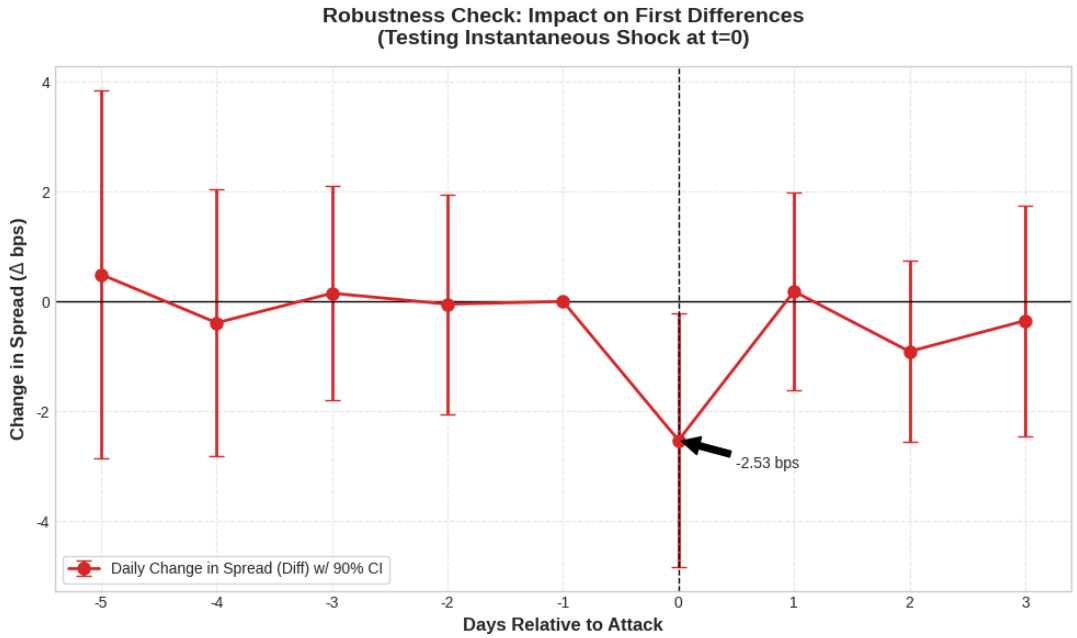


Figure A2: **Instantaneous Marginal Impact of DeFi Hacks.** This figure plots the coefficients from the first-difference specification. The significant negative spike at $t = 0$ followed by insignificant coefficients indicates a permanent step-down in spread levels.

Table A11: Comparison of ZGIV Constructed With and Without De-meaning Term

Panel A: Descriptive Statistics

Variable	Mean	Std. Dev.	Min	Max	Correlation
ZGIV (Standard)	-0.146%	0.176%	-0.671%	-0.001%	1.000
ZGIV (No De-mean)	-0.134%	0.162%	-0.610%	-0.001%	0.999

Panel B: Regression Analysis (Dependent Variable: ZGIV Standard)

Independent Variable	Coefficient	Std. Err.	t-statistic	P-value
Intercept (α)	0.000	0.000	0.04	0.969
ZGIV (No De-mean) (β)	1.091***	0.005	204.07	<0.001
R^2	0.999			
Observations	49			

Note: This table compares the ZGIV time series constructed with the standard de-meaning term ($\bar{g}_t$) versus a version constructed without it (pure size-weighted shock aggregation). Panel A reports the descriptive statistics and the Pearson correlation coefficient between the two series. Panel B reports the results of an OLS regression where the standard ZGIV is regressed on the ZGIV without the de-meaning term. The correlation exceeds 0.99 and the R^2 is near unity, confirming that due to the sparsity of high-severity shocks, the size-weighted aggregation component dominates the common factor correction. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A12: **Robustness Check: Alternative Event Date Definitions (Occurrence vs. Disclosure)**

Specification	**Event Day Definition**	δ_0 **Coef. (bps)**	p**-value**
(1) Baseline	Occurrence Day (On-chain Timestamp)	-2.183^{**}	0.024
(2) Robustness	First Public Disclosure Day	-1.754^{*}	0.060

Note: To address concerns about aligning intraday hack timestamps with daily CP spreads, we construct an alternative event-time definition based on the *first public disclosure* of each exploit. We define the "Disclosure Day" as the earliest date of a protocol official announcement, a security firm alert (e.g., PeckShield, CertiK), or a major news report. For the Top-10 exploits, we manually verified the timeline and re-estimated the stacked event study using the disclosure day (aligned to U.S. trading days) as $t = 0$. Specifically, this specification adjusts the **Ronin Network** event from March 23 (occurrence) to March 29 (first public news). The estimated δ_0 remains negative and statistically significant, indicating that our baseline findings are not driven by event-time misalignment. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A13: Mechanism evidence from MMF holdings (Fed EFA), 2021–2024

Panel A. CP exposure is concentrated in prime MMFs

Statistic	Mean	Std. Dev.	Min	Max
Prime fraction of total MMF CP	0.9779	0.0055	0.9656	0.9853

Panel B. Portfolio tilt in hack vs. non-hack months

Metric	Hack months	Non-hack months	Diff	(S.E.)	$N_{\text{hack}}/N_{\text{non}}$
Prime CP share *(CP+ABCP / Prime holdings)*	0.262	0.236	0.026***	(0.007)	29/19
Prime repo share *(Repo / Prime holdings)*	0.315	0.361	-0.046*	(0.023)	29/19
Government Treasury share *(Treasuries / Gov holdings)*	0.428	0.363	0.066**	(0.031)	29/19
Government repo share *(Repo / Gov holdings)*	0.446	0.496	-0.049	(0.032)	29/19

Notes. Monthly MMF holdings are from the Federal Reserve EFA money market fund holdings aggregates. "Prime fraction of total MMF CP" is defined as prime MMFs' CP+ABCP holdings divided by total MMF CP holdings. A month is classified as a hack month if it contains at least one DeFi exploit event (hack_count > 0), based on the paper's event sample. Panel B reports Welch two-sample difference-in-means tests comparing hack vs. non-hack months. Standard errors in parentheses correspond to Welch's unequal-variance formula. Significance stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A14: Mechanism evidence from MMF holdings: regression form (2021–2024)

Dep. Var.:	(1) Prime fraction of total MMF CP	(2) Prime CP share	(3) Prime repo share	(4) Gov Treasury share	(5) Gov repo share
HackMonth	-0.001	0.006*	0.014	-0.003	0.010
	(0.001)	(0.004)	(0.013)	(0.016)	(0.017)
Year FE	Yes	Yes	Yes	Yes	Yes
Month-of-year FE	Yes	Yes	Yes	Yes	Yes
Observations	48	48	48	48	48
R^2	0.814	0.801	0.772	0.670	0.673

Notes. This table estimates monthly regressions of MMF holdings measures on an indicator for months containing at least one DeFi exploit event (HackMonth). Holdings measures are constructed from Federal Reserve EFA MMF holdings aggregates: *Prime fraction of total MMF CP* is prime MMFs' CP+ABCP holdings divided by total MMF CP holdings; *Prime CP share* is (CP+ABCP)/prime holdings; *Prime repo share* is total repo/prime holdings; *Gov Treasury share* is Treasuries/government holdings; *Gov repo share* is total repo/government holdings. All specifications include year fixed effects and month-of-year fixed effects. Standard errors (in parentheses) are Newey–West HAC with one monthly lag. Significance stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A15: **Robustness of Structural Parameter η: Subsamples and Bootstrapping**

Specification	**Estimated Impact** β (bps)	**Implied Parameter** η	**Inference** *(State / Reliability)*	
Panel A: Subsample Analysis (Regime Dependence)				
Pre-Terra (DeFi Boom)	-1.90	**2.90**	Moderate Aversion	
Post-Terra (Crypto Winter)	-4.14	**5.14**	Heightened Panic	
Panel B: Non-Parametric Bootstrapping (N=1,000)				
	Mean Estimate	**95% CI**	**Lower**	**Upper**
Bootstrapped Distribution	2.45	[1.52, 3.45]	1.52	3.45

Notes: This table reports robustness checks for the structural parameter η, calibrated using the formula $\eta = 1 - \beta/\lambda$ with baseline $\lambda = 1.0$.
Panel A splits the sample around the Terra/Luna collapse (May 7, 2022). The higher η in the Post-Terra period confirms that ambiguity aversion intensifies during periods of low market trust.
Panel B reports results from 1,000 bootstrap iterations where event dates are resampled with replacement. The 95% confidence interval strictly exceeds 1, confirming the statistical significance of the behavioral amplification channel.

Table A16: Step 3: Frequency-aligned state dependence (monthly regression on event-window CP spread)

	(1) Δ Spread$_m$	(2) Spread$_m$
HackMonth	-17.421**	-24.264***
	(7.448)	(4.762)
pcs_z	13.159**	16.812***
	(6.134)	(4.193)
HackXpcs	-15.483**	-18.668***
	(6.266)	(4.002)
Controls (VIX, DXY, BTC)	Yes	Yes
Observations	33	34
R^2	0.156	0.619

Notes. This table estimates frequency-aligned monthly regressions that link the daily event-study sample to monthly MMF holdings. Spread$_m$ is the average daily CP spread (bps) over calendar days that appear in the stacked event-study panel within month m (i.e., event-window days). ΔSpread$_m$ is the month-to-month change in this monthly average. The state variable pcs_z is the standardized monthly prime CP share from Fed EFA; HackMonth indicates months with at least one exploit event in the sample; HackXpcs is their interaction. All specifications include the monthly averages of VIX, DXY, and BTC returns computed over the same set of days. Standard errors are Newey–West HAC with one monthly lag. Significance stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

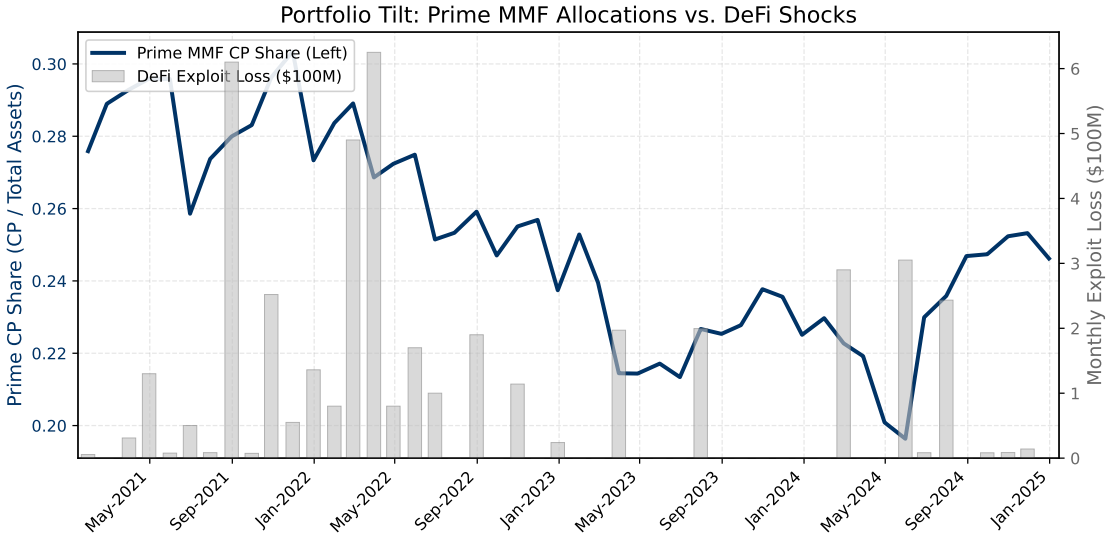


Figure A3: **Prime CP Share and Hack Intensity (Monthly, 2021–2024)**

Notes: This figure plots the monthly share of commercial paper in prime money market fund portfolios (*prime_cp_share*, left axis) and the aggregate monthly DeFi hack loss (*total_loss_100m*, right axis; losses in $100 million). Gray bars indicate the monthly hack-loss intensity, and shaded bands mark months with at least one major DeFi hack (*hack_count* > 0). The sample covers January 2021 to December 2024 at monthly frequency.

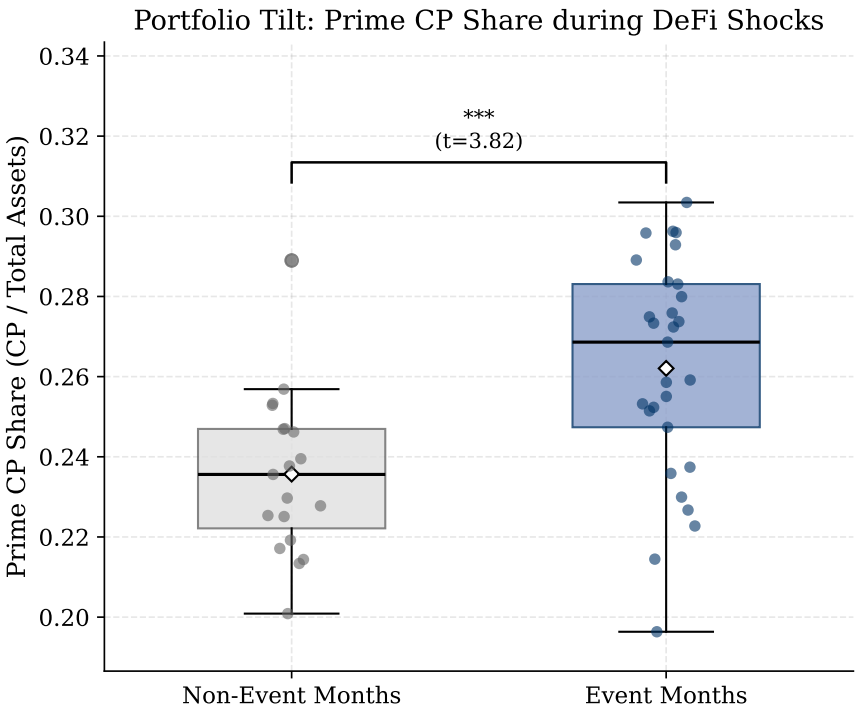


Figure A4: **Prime CP Share: Hack Months vs. Non-Hack Months (Monthly, 2021–2024)**

Notes: This boxplot compares the distribution of the prime funds' commercial paper share (*prime_cp_share*) between "hack months" (months with *hack_count* > 0) and "non-hack months" (*hack_count* $= 0$) over 2021–2024. Boxes show the interquartile range (25th–75th percentiles), the center line denotes the median, whiskers indicate the data range excluding outliers, and the mean is displayed as a distinct marker (diamond). Significance levels are indicated based on Welch's t-test.

Jordà (2005) Local Projections (LP) framework avoids artificially inflating the degrees of freedom and explicitly corrects for serial correlation using **Newey-West HAC** standard errors.Instead of estimating a stacked panel, we estimate the impulse response functions (IRFs) directly on the single daily time series of the Commercial Paper market ($N = T$). For each forecast horizon $h = 0, 1, \ldots, 6$, we estimate the following local projection specification:

$$\text{Spread}_{t+h} - \text{Spread}_{t-1} = \alpha_h + \beta_h \cdot \text{Shock}_t + \sum_{j=1}^{L} \gamma_{h,j} \mathbf{X}_{t-j} + \epsilon_{t+h} \tag{A1}$$

Where:

- $\text{Spread}_{t+h} - \text{Spread}_{t-1}$ denotes the cumulative change in the 3-Month AA Commercial Paper spread from the day prior to the shock $(t-1)$ to horizon h.
- Shock_t represents the DeFi exploit impulse. We test two specifications: (1) a **Binary Shock** indicator equal to 1 on event days, and (2) an **Intensity Shock** defined as $\ln(\text{Loss Amount}_t)$ to capture the heterogeneity in shock severity.
- $\mathbf{X}_{t-j}$ is a vector of lagged control variables, including the VIX index, the US Dollar Index (DXY), and other macro controls, to control for pre-existing market trends.
- We employ **Newey-West HAC** (Heteroskedasticity and Autocorrelation Consistent) standard errors with a bandwidth of $h+1$ to account for the serial correlation inherent in overlapping forecast horizons.

The results, reported in Table A17, confirm that our baseline findings are robust to this rigorous time-series inference.

Table A17: **Dynamic Impact of DeFi Shocks on CP Spreads: Jordà Local Projections**

Horizon	**(1) Binary Shock**			**(2) Intensity Shock (Log Loss)**		
(Days)	Coef.	t-stat	95% CI	Coef.	t-stat	95% CI
$h = 0$ (Impact)	-3.003***	(-3.41)	[-4.73, -1.28]	-0.155***	(-3.49)	[-0.24, -0.07]
$h = 1$	-2.068**	(-2.32)	[-3.82, -0.32]	-0.107**	(-2.37)	[-0.20, -0.02]
$h = 2$	-1.407	(-1.41)	[-3.36, 0.55]	-0.073	(-1.44)	[-0.17, 0.03]
$h = 3$	-2.105*	(-1.86)	[-4.32, 0.11]	-0.105*	(-1.81)	[-0.22, 0.01]
$h = 4$	-2.085*	(-1.70)	[-4.48, 0.31]	-0.102	(-1.61)	[-0.23, 0.02]
$h = 5$	-1.447	(-1.23)	[-3.76, 0.86]	-0.070	(-1.15)	[-0.19, 0.05]
$h = 6$	-0.701	(-0.63)	[-2.90, 1.50]	-0.035	(-0.62)	[-0.15, 0.08]
Observations	676 (Daily Time Series)			676 (Daily Time Series)		
Controls	Yes			Yes		
Inference	Newey-West HAC (Lag $h + 1$)			Newey-West HAC (Lag $h + 1$)		

Notes: This table reports the impulse response coefficients (β_h) estimated using Jordà (2005) Local Projections. The dependent variable is the cumulative change in the 3-Month AA Nonfinancial Commercial Paper Spread from $t - 1$ to $t + h$ (in basis points).

Model (1) defines the shock as a binary dummy variable equal to 1 on days with a Top-50 DeFi exploit occurrence.

Model (2) defines the shock as the natural logarithm of the USD loss amount ($\ln(Loss)$) to account for shock intensity heterogeneity.

The estimation sample consists of a single continuous daily time series from 2021 to 2024 ($N = 676$), avoiding the "pseudo-panel" inflation of degrees of freedom present in stacked event studies.

All specifications include lagged controls (X_{t-1}) for market volatility (VIX), the US Dollar Index (DXY), and the lagged dependent variable.

t-statistics reported in parentheses are based on **Newey-West HAC robust standard errors** with a bandwidth of $h + 1$ to correct for the serial correlation inherent in overlapping forecast horizons.

Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A18: **Structural Break Test: Instantaneous Impact at $t = 0$ (Narrow Window LP)**

Dependent Variable:	ΔSpread$_t$ (Daily Change)	
Shock Definition	**(1) All Events (N=50)**	**(2) Top 20 Focus (N=20)**
Shock Impact (β)	**-3.071*****	**-5.432*****
	(-3.10)	(-2.94)
Lagged Δ Spread	-0.253***	-0.253***
	(-3.13)	(-3.13)
Lagged VIX	0.130*	0.141**
	(1.93)	(2.08)
Constant	-2.025*	-2.295**
Observations	676	676
Adj. R^2	0.042	0.045
Structural Break Type	Instantaneous Level Shift (Jump)	

Notes: This table reports the results of the narrow window ($t = 0$) structural break test estimated on the daily time series of 3-Month AA CP spreads.
Column (1) uses the baseline binary shock for all Top 50 events. **Column (2)** restricts the shock indicator to only the Top 20 largest events by USD loss amount.
The coefficient β measures the instantaneous change (jump) in the spread level on the event day. The sharp increase in magnitude from -3.07 bps (All) to **-5.43 bps (Top 20)** confirms that the "Blessing in Disguise" effect is driven by significant information shocks rather than noise.We estimate the first-difference equation using the single daily time series ($N = T$):

$$\Delta\text{Spread}_t = \alpha + \beta \cdot \text{Shock}_t + \gamma_1 \Delta\text{Spread}_{t-1} + \gamma_2 \text{VIX}_{t-1} + \epsilon_t \tag{A2}$$

Where:

- $\Delta\text{Spread}_t = \text{Spread}_t - \text{Spread}_{t-1}$ represents the daily jump in the spread.
- Shock_t is a binary indicator equal to 1 on the event disclosure day. We test two definitions: (1) **All Events** (Top 50), and (2) **Top 20 Focus** (restricting shocks to the largest 20 exploits to minimize noise).
- A statistically significant $\beta \neq 0$ confirms the existence of an instantaneous structural break (level shift) at $t = 0$.

t-statistics based on Newey-West HAC standard errors are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

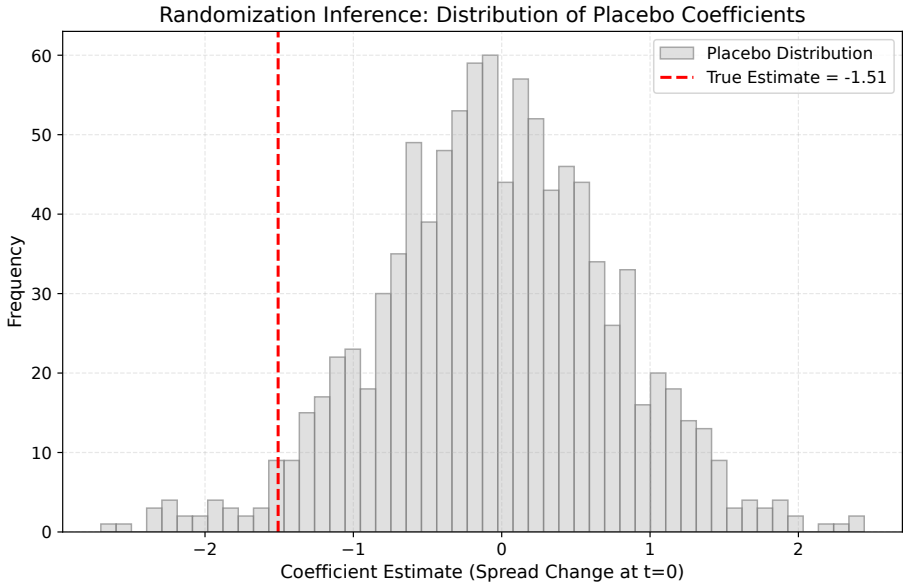


Figure A5: **Randomization Inference: Distribution of Placebo Coefficients**
Notes: This figure presents the distribution of placebo coefficients obtained from 1,000 simulations of the stacked event study. In each simulation, we randomly assign 50 "pseudo-event" dates within the sample period (2021–2024), construct the corresponding $[-5, +5]$ stacked windows, and re-estimate the baseline regression. The gray histogram represents the null distribution of the treatment effect under the hypothesis of no relationship. The solid red line indicates the true estimated coefficient ($\hat{\beta} = -1.51$) from the actual DeFi exploit dates. The empirical p-value, calculated as the fraction of placebo estimates larger in magnitude than the true estimate, is **0.053**. This confirms that the observed spread narrowing is statistically significant and unlikely to be driven by random market noise or spurious trends.

Table A19: **Robust Check:Difference-in-Differences Results**

Event Time (k) (Relative Days)	**Coefficient (β_k in bps)**	**Std. Error**	**t-statistic**	**P-value**
Pre-Event Trends				
$t = -5$	3.66	9.28	0.39	0.693
$t = -4$	-3.56	4.53	-0.79	0.432
$t = -3$	-10.83	–	–	–
$t = -2$	0.00	0.00	0.00	1.000
$t = -1$	0.00	–	–	*(Benchmark)*
Post-Event Impact				
t = 0 (Event Day)	**-2.18****	**1.10**	**-1.98**	**0.048**
t = 1	**-4.36*****	**1.44**	**-3.02**	**0.003**
$t = 2$	0.61	5.26	0.12	0.907
$t = 3$	-6.77*	3.77	-1.80	0.072

Notes: This table reports the coefficients from a stacked Difference-in-Differences (DiD) event study estimating the impact of major DeFi hacks (Loss > \$100M) on the spread of AA Non-financial Commercial Paper relative to a control group (A2/P2 CP, Financial CP, and Repo). The regression model is specified as:

$$Spread_{a,t} = \alpha_a + \gamma_t + \sum_{k \neq -1} \beta_k (1[t - T_i = k] \times Treat_a) + \varepsilon_{a,t}$$

where $Treat_a = 1$ for AA Non-financial CP and 0 otherwise. Asset fixed effects (α_a) and Date fixed effects (γ_t) are included to absorb time-invariant asset characteristics and common macroeconomic shocks. Standard errors are clustered at the event level.

- **Pre-Event:** Coefficients for $t < 0$ are statistically insignificant, supporting the parallel trends assumption.
- **Post-Event:** The coefficient at $t = 0$ ($\beta_0 = -2.18$ bps) indicates an immediate, statistically significant flight-to-quality effect. This effect intensifies on $t = 1$ ($\beta_1 = -4.36$ bps), confirming the persistence of the liquidity injection.

Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A20: Cross-Asset Difference-in-Differences

Asset Class: *(Relative to Target)*	**(1) Control: Junk** (A2/P2 CP)	**(2) Benchmark: Repo** (SOFR)	**(3) Control: Sector** (AA Fin)	**(4) Alt. Treat: ABCP** (AA)
Pre-Event				
$t = -5$	2.386	-5.671	1.926	3.686**
$t = -4$	4.361	-8.994*	-1.990	2.538
$t = -3$	3.927	0.199	-0.398	1.287
$t = -2$	-3.312	5.064	-1.236	0.536
$t = -1$	0.000	0.000	0.000	0.000
Post-Event				
t = 0	**4.618**	**1.457**	**-2.031**	**1.481**
	(3.133)	(4.583)	(1.700)	(1.618)
t = 1	**5.004**	**0.259**	**-1.053**	**0.879**
	(3.262)	(5.437)	(1.873)	(1.712)
$t = 2$	6.740**	-6.647	-0.008	3.853
$t = 3$	7.931*	-6.830	-1.626	2.786
Asset FE	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes
Clustering	Event	Event	Event	Event
Parallel Trends Test:				
Joint F-Stat	0.88	1.68	1.11	2.32
Prob > F	[0.483]	[0.170]	[0.364]	[0.070]

Notes: This table reports the coefficients from a flexible Difference-in-Differences specification estimating the heterogeneous response of different money market instruments relative to the treated asset (AA Non-financial CP). The regression model is:

$$Spread_{a,t} = \alpha_a + \gamma_t + \sum_{k \neq -1} \sum_{g \in G} \beta_{k,g} (1[t - T_i = k] \times \mathbb{I}_{a \in g}) + \varepsilon_{a,t}$$

where the baseline asset group is AA Non-financial CP. A positive coefficient $\beta_{k,g}$ indicates that asset group g experienced a widening in spreads relative to the treated asset. **Col (1)** shows lower-rated (A2/P2) CP spreads widened significantly, confirming the flight-to-quality. **Col (2)** shows the risk-free Repo rate (SOFR) remained statistically indistinguishable from AA CP. **Cols (3)-(4)** confirm other Prime-eligible assets tracked the treated asset closely. **Parallel Trends Test:** Reports the F-stat and p-value [in brackets] for the null hypothesis that all pre-event coefficients ($t < 0$) are jointly zero. High p-values support the parallel trends assumption. Standard errors clustered by event in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Micro-foundations of Knightian Uncertainty and Flight-to-Quality

In this appendix, we provide the decision-theoretic micro-foundations for the "Panic Parameter" (ρ_2) and the "Flight-to-Quality" mechanism ($\eta > 1$) presented in the main text. We employ a continuous-time portfolio choice model incorporating Knightian Uncertainty (ambiguity) regarding the intensity of DeFi hacks. Unlike standard risk (which is measurable), the arrival rate of a hack involves unknown structural parameters.

Following the robust control framework of Anderson, Hansen, and Sargent (2003) and the jump-diffusion ambiguity model of Liu, Pan, and Wang (2005), we model the representative investor as ambiguity averse. The investor does not maximize expected utility under a single probability measure $\mathbb{P}$, but rather maximizes utility under a "worst-case" scenario chosen from a set of alternative measures $\mathbb{Q}$, penalized by relative entropy.

B.1 The Economy and Asset Dynamics

Consider a continuous-time economy with two assets:

Safe Asset (Traditional Finance): A risk-free asset (representing Commercial Paper/T-Bills) yielding a constant rate r:

$$dB_t = rB_t dt \tag{A3}$$

Risky Asset (DeFi Protocol): A risky asset whose price P_t follows a jump-diffusion process. The jump component represents the exogenous shock of a hack:

$$\frac{dP_t}{P_t} = (\mu + r)dt + \sigma dZ_t - LdN_t \tag{A4}$$

Where:

- μ is the equity risk premium.
- σ is the diffusive volatility (standard market noise).

- Z_t is a standard Brownian motion under the reference measure $\mathbb{P}$.
- N_t is a Poisson process with intensity λ, representing the arrival of hacks.
- $L \in (0, 1]$ is the proportional loss given a hack (assumed constant for simplicity, representing the "Severity").

B.2 Preferences: Duffie-Epstein Utility with Ambiguity

We assume the representative investor has Duffie-Epstein recursive preferences defined by the aggregator $f(C, V)$. To introduce Knightian uncertainty, we adopt the continuous-time version of the Gilboa-Schmeidler (1989) max-min expected utility, formulated as a robust control problem. The investor doubts the reference intensity λ of the Poisson process. They consider a set of alternative models defined by the density process ξ_t, which distorts the jump intensity to $\tilde{\lambda} = \xi_t \lambda$.

The investor's value function V_t is given by the solution to:

$$V_t = \sup_{C,w} \inf_{\xi} \mathbb{E}_t^{\mathbb{Q}_\xi} \left[\int_t^\infty e^{-\delta(s-t)} \left(\frac{C_s^{1-\gamma}}{1-\gamma} + \Psi \mathcal{H}(\xi_s) \right) ds \right] \tag{A5}$$

Here:

- w_t is the fraction of wealth allocated to DeFi.
- γ is the coefficient of relative risk aversion.
- $\Psi > 0$ is the Ambiguity Tolerance parameter (so $1/\Psi$ represents ambiguity aversion). As $\Psi \to \infty$, the model converges to standard expected utility (i.e., $\xi_t \to 1$).
- $\mathcal{H}(\xi_t) = \lambda(\xi_t \ln \xi_t - \xi_t + 1)$ is the relative entropy (KL-divergence) penalty rate for distorting the jump intensity.

B.3 The HJB Equation and Worst-Case Intensity

Let $J(W)$ denote the value function. In a stationary equilibrium, the Hamilton-Jacobi-Bellman (HJB) equation for the robust investor is:

$$0 = \sup_{C,w} \inf_{\xi} \left\{ \frac{C^{1-\gamma}}{1-\gamma} - \delta J + \mathcal{L}^{w,\xi} J + \Psi \mathcal{H}(\xi) \right\} \tag{A6}$$

Under portfolio share w, wealth evolves as:

$$\frac{dW_t}{W_t} = r dt + w\mu dt + w\sigma dZ_t - wL dN_t - \frac{C_t}{W_t} dt \tag{A7}$$

Under the distorted measure $\mathbb{Q}_\xi$, the jump intensity becomes $\tilde{\lambda} = \xi\lambda$. The generator $\mathcal{L}^{w,\xi} J$ is therefore:

$$\begin{aligned} \mathcal{L}^{w,\xi} J = J_W(W) W \left(r + w\mu - \xi\lambda wL - \frac{C}{W} \right) \\ + \frac{1}{2} J_{WW}(W) W^2 w^2 \sigma^2 + \xi\lambda [J(W(1-wL)) - J(W)] \end{aligned} \tag{A8}$$

Notice the term $\xi\lambda$. The investor ("maximizer") chooses C and w, while Nature ("minimizer") chooses the distortion factor ξ to inflict the most damage, subject to the entropy penalty.

B.3.1 Solving the Inner Minimization (Nature's Move)

The first-order condition (FOC) with respect to ξ yields the worst-case distortion. Using $\frac{\partial}{\partial \xi} \mathcal{H}(\xi) = \lambda \ln \xi$, the FOC is:

$$-\lambda wL J_W(W) W + \lambda [J(W(1-wL)) - J(W)] + \Psi \lambda \ln \xi = 0 \tag{A9}$$

Solving for the optimal distortion ξ^*:

$$\ln \xi^* = \frac{1}{\Psi} \left(J(W) - J(W(1-wL)) + wL J_W(W) W \right), \quad \xi^* = \exp(\ln \xi^*) \tag{A10}$$

To establish the sign, define the concavity (Jensen) gap:

$$\mathcal{G}(W; wL) \equiv J(W) - J(W(1 - wL)) - wLJ_W(W)W \tag{A11}$$

For strictly concave J (i.e., $J_{WW} < 0$) and any $wL \in (0, 1)$, we have $\mathcal{G}(W; wL) > 0$. Hence:

$$\ln \xi^* = \frac{1}{\Psi}(\mathcal{G}(W; wL) + 2wLJ_W(W)W) > 0 \implies \xi^* > 1 \implies \lambda^* = \xi^*\lambda > \lambda \tag{A12}$$

Interpretation. The investor operates under a subjective intensity $\lambda^* > \lambda$. The more ambiguity-averse the investor (lower Ψ), or the larger the potential wealth loss wL, the higher the "imagined" probability of a hack (i.e., the larger ξ^*).

B.4 Optimal Portfolio and the "Panic Threshold"

Now solve the outer maximization for the portfolio weight w. Assume a CRRA value function $J(W) = A\frac{W^{1-\gamma}}{1-\gamma}$ with $A > 0$. Then $J_W(W) = AW^{-\gamma}$ and:

$$\frac{J_W(W(1 - wL))}{J_W(W)} = (1 - wL)^{-\gamma} \tag{A13}$$

The FOC for w can be written compactly as:

$$0 = \mu - \gamma w\sigma^2 - \lambda\xi^* \cdot L\left[(1 - wL)^{-\gamma} - 1\right] \tag{A14}$$

Intuitively, the last term is the (worst-case) marginal jump-risk premium: under a jump, wealth becomes $W(1 - wL)$ and marginal utility rises by the factor $(1 - wL)^{-\gamma}$. The distortion ξ^* scales the effective jump intensity upward.

Proposition 1 (Discontinuous Exit / "Panic" Threshold)**.** *There exists a critical level of structural uncertainty (equivalently, sufficiently low Ψ, or sufficiently high wL) such that:*

- ***Regime 1 (Standard Risk):*** *When ambiguity is low (Ψ high), $\xi^* \approx 1$. The*

investor holds a reduced but positive position in DeFi ($w^* > 0$).

- ***Regime 2 (Knightian Panic):*** *When a major hack occurs and model confidence collapses ($\Psi \downarrow$), the distortion ξ^* rises, the effective jump premium becomes large, and the interior solution disappears. With short-selling/limited-liability constraints typical in DeFi, the optimal solution hits the corner:*

$$w^*_{DeFi} = 0, \quad w^*_{Safe} = 1 \tag{A15}$$

This corner solution rationalizes the step-like "panic" increase in redemptions: once ambiguity spikes, investors do not rebalance marginally; they exit fully.

B.5 Theoretical Connection to Empirical Findings

This derivation provides the rigorous foundation for the empirical results in Section 6:

Explaining ρ_2 (Panic Parameter): The step-function increase in redemptions modeled in Eq. (3.2) corresponds to the regime switch from an interior solution to the corner solution ($w^* = 0$) induced by the explosion of the Radon-Nikodym derivative ξ^*.

Explaining $\eta > 1$ (Flight-to-Quality): The model shows that under ambiguity, the demand for the safe asset is not linear. A shock to λ does not just adjust w marginally; via the entropy penalty, it creates a multiplier effect on risk perception. The "Ambiguity Premium" dominates the standard risk premium, forcing a total reallocation of wealth W into the safe asset B_t.

Thus, the "Blessing in Disguise" is theoretically grounded in the rational response of an ambiguity-averse agent facing "Unknown Unknowns," who inadvertently subsidizes the traditional safe asset market by flooding it with liquidity.

B.6 Structural Interpretation of η

B.6.1 Defining η as the Measure of Distortion

In our model, the parameter η captures the excess magnitude of the "Flight-to-Quality" demand relative to the physical liquidation supply.

- **Physical liquidation (Supply shock):** The liquidation of reserves by stablecoin issuers is driven by realized redemptions. In the reference model, this realized flow is proportional to the physical intensity λ:

$$S_{\text{shock}} \propto \lambda \tag{A16}$$

- **Robust demand (Demand shock):** Under Knightian uncertainty, investors behave under the distorted intensity $\lambda^* = \xi^* \lambda$. During panic episodes in which jump-risk perceptions dominate short-run portfolio reallocation, the safe-asset demand impulse is approximately proportional to λ^*:

$$D_{\text{shock}} \propto \lambda^* = \xi^* \lambda \tag{A17}$$

B.6.2 Closed-Form Mapping

Under this crisis-region mapping (jump-dominance approximation), the amplification factor is:

$$\eta \equiv \frac{D_{\text{shock}}}{S_{\text{shock}}} \approx \frac{\lambda^*}{\lambda} = \xi^* \tag{A18}$$

B.6.3 Proof that $\eta > 1$ in Panic Episodes

From B.3.1, strict concavity of the value function implies the concavity gap $\mathcal{G}(W; wL) > 0$, hence $\ln \xi^* > 0$ and $\xi^* > 1$. Therefore:

$$\eta \approx \xi^* > 1 \tag{A19}$$

Remark (Scope). The identity $\eta \approx \xi^*$ is a structural approximation that is most

accurate when (i) jump-risk perceptions dominate diffusive risk at the horizon of interest and (ii) the marginal safe-asset demand response is primarily governed by perceived disaster intensity. This is the precise sense in which ambiguity aversion delivers $\eta > 1$ as the "Blessing" condition in panic states.

C A Global Game of Runs with Endogenous Congestion and Ambiguity

To provide rigorous micro-foundations for the non-linear threshold observed in Section 5.2, we model the stablecoin redemption dynamic as a Global Game (Morris and Shin, 1998; 2003). This framework allows us to derive a unique equilibrium threshold for the run, resolving the indeterminacy of multiple equilibria typical in standard coordination games. Our innovation lies in introducing endogenous congestion costs (Gas Fees) that act as a strategic substitute, counteracting the strategic complementarity of the run. We show that Knightian uncertainty regarding the hack's severity creates a specific "Panic Condition" under which the congestion circuit breaker fails.

C.1 The Economy

Consider a continuum of investors $i \in [0, 1]$, each holding 1 unit of a stablecoin.

Fundamental State: The DeFi protocol suffers a hack with severity θ, where $\theta \in \mathbb{R}$ represents the "fundamentals" (e.g., the proportion of remaining collateral). A lower θ implies a worse state. θ is drawn from a uniform prior over the real line (improper prior limit).

- $\theta \leq 0$: Insolvency is certain (Fundamental Run Zone).
- $\theta \geq 1$: Solvency is strictly guaranteed (Fundamental Safe Zone).
- $\theta \in (0, 1)$: The "Fragile Zone," where solvency depends on the aggregate run size.

Information Structure: Investors do not observe θ perfectly. Instead, agent i receives a noisy private signal:

$$x_i = \theta + \sigma\epsilon_i, \quad \epsilon_i \sim N(0,1) \tag{A20}$$

where $\sigma \to 0$ represents vanishing noise (precise information limit).

Actions: Each agent chooses $a_i \in \{0,1\}$, where $a_i = 1$ corresponds to "Redeem" (Run) and $a_i = 0$ corresponds to "Wait." Let $A = \int_0^1 a_i \, di$ be the aggregate size of the run.

C.2 Payoffs and Endogenous Friction

The utility function depends on the action, the aggregate run A, and the fundamental θ.

Case 1: Waiting ($a_i = 0$).

$$u(0, A, \theta) = \begin{cases} 1 & \text{if } A < \theta \quad \text{(Solvent)} \\ 0 & \text{if } A \geq \theta \quad \text{(Insolvent)} \end{cases} \tag{A21}$$

Case 2: Running ($a_i = 1$). The cost of running is defined by the congestion function $C(A) = \phi_0 + \gamma A^\lambda$, where $\gamma > 0$ and $\lambda \geq 1$.

$$u(1, A, \theta) = 1 - C(A)^3 \tag{A22}$$

C.3 The Indifference Condition and Unique Equilibrium

We solve for the unique equilibrium in the limit of precise information (vanishing noise, $\sigma \to 0$). Following Morris and Shin (2003), this assumption of Laplacian Beliefs implies that at the switching threshold, an agent assigns a uniform probability density to the aggregate action A over the interval $[0,1]$.

[3]Our payoff structure implies that 'Running' eliminates fundamental risk. This reflects the Priority Gas Auction mechanism inherent to Ethereum-based DeFi: agents who pay the equilibrium congestion cost $C(A)$ are prioritized by miners/validators. Thus, they are guaranteed to exit before the protocol is drained, rendering their recovery independent of the insolvency state θ, subject only to the friction cost.

Expected Payoff Calculation. The expected utility of running (LHS) versus waiting (RHS) at the threshold θ^* is given by:

1. **Expected Utility of Running (LHS):**

$$\text{LHS} = \int_0^1 (1 - C(A))\, dA = 1 - \bar{C} \tag{A23}$$

Where the average congestion cost $\bar{C}$ is:

$$\bar{C} = \int_0^1 (\phi_0 + \gamma A^\lambda)\, dA = \phi_0 + \frac{\gamma}{\lambda + 1} \tag{A24}$$

2. **Expected Utility of Waiting (RHS):**

$$\text{RHS} = \int_0^1 \mathbb{1}_{\{A < \theta^*\}}\, dA = \theta^* \tag{A25}$$

Unique Threshold Solution. Equating LHS and RHS, we obtain the unique dominance solvable threshold:

$$\theta^* = 1 - \bar{C} = 1 - \left(\phi_0 + \frac{\gamma}{\lambda + 1}\right) \tag{A26}$$

C.4 Incorporating Ambiguity

Under Knightian uncertainty, the indifference condition is distorted by the ambiguity aversion function $\Psi(\cdot)$:

$$1 - \bar{C} = \Psi(\theta^*) \tag{A27}$$

Solving for the ambiguity-adjusted threshold:

$$\theta^*_{ambiguity} = \Psi^{-1}(1 - \bar{C}) \tag{A28}$$

Since $\Psi(p) < p$ for ambiguity-averse agents, it follows that $\Psi^{-1}(y) > y$. Therefore:

$$\theta^*_{ambiguity} > \theta^*_{risk-neutral} \tag{A29}$$

Economic Interpretation (Threshold Logic):

- θ^*: Represents the **Panic Threshold**.
- Higher friction ($\bar{C}$) reduces θ^* (Congestion dampens the run).
- Ambiguity (Ψ^{-1}) increases θ^* (Uncertainty amplifies the run).

Proposition 2. *There exists a critical level of ambiguity aversion such that even if $\bar{C}$ is extremely high, the threshold θ^* shifts to cover nearly the entire state space. This explains why, beyond the empirical threshold of 32.93 Gwei, the "Circuit Breaker" effect vanishes—the shadow price of ambiguity exceeds the physical price of congestion.*