

ARTICLE TYPE: PERSPECTIVE WITH ORIGINAL DATA SYNTHESIS

Off-Exchange Migration and the Limits of AI Attribution: A Data Synthesis of Venue Composition, Execution Quality, and Enforcement in U.S. Equity Markets, 2019 to 2025

Authors: Z. Bari^{1*}, K. Ragunath^{1*}, S. Tripathi^{1*†}

Affiliation: ¹ Middleton High School, Tampa, Florida, United States

* These authors contributed equally to this work and are listed in alphabetical order by last name.

† **Corresponding author:** S. Tripathi, suyash2101@gmail.com

Manuscript length: 4,767 words in the main text (Sections 1 to 5), excluding the abstract, references, tables, figures, and appendices. The NSRI template sets 1,000 to 2,000 words for an Opinion and 1,500 to 3,000 for a Research Archive but states no limit for a Perspective with Original Data Synthesis, which combines the two. The count reflects that combination: the argument occupies roughly the Opinion allowance and the documented data synthesis the Research Archive allowance. All dataset documentation and the derivation ledger sit in Appendix A and Supplementary Files S1 to S7 rather than in the main text.

Abstract

Background. Off-exchange venues executed a record 51.8% of United States equity share volume in January 2025. A widely repeated interpretation attributes this migration to machine learning embedded in matching and routing infrastructure. This paper tests whether public data can support that attribution, and reports what they support instead.

Methods. Six datasets were compiled from FINRA, SEC, and ESMA publications, from enforcement releases, and from the peer-reviewed and working-paper literature, and are provided in full with a calculation workbook in which every derived statistic is a live formula. A three-tier evidence protocol separates directly measured values from author-derived arithmetic and from interpretive attribution, and ten robustness tests are reported alongside the results they test.

Results. In share terms, non-ATS off-exchange venues account for 86.3% of the 14.6 percentage-point rise in off-exchange share since 2019, a result robust to endpoint and window variation but not to dollar denomination, under which the figure falls to 54.1%. The 15.93 basis-point price-impact differential between retail and institutional execution indicates lower adverse selection in retail flow and is not, as it is often read, an execution penalty borne by retail investors. Sixteen enforcement actions totaling \$330.15 million cluster in 2015 to 2018, and none alleges conduct involving a machine-learning system.

Conclusion. The public record establishes a compositional shift toward bilateral internalization, a segmented market, and a disclosure regime that has not been updated for model-driven venues. It does not establish that machine learning caused any of these. The governance question stands without the attribution, and a difference-in-differences design using audit-trail microdata is specified that would settle it.

Keywords: venue composition; order-flow segmentation; alternative trading systems; wholesale internalization; evidence tiering; regulatory disclosure; reproducible synthesis

1. Introduction

Off-exchange venues executed a record 51.8% of United States equity share volume in January 2025, the fifth consecutive monthly record and the first sustained period in which the majority of consolidated volume printed away from displayed markets.¹ Over the same years, machine-learning systems moved to the center of execution infrastructure, governing order matching, counterparty screening, and routing at update frequencies that the disclosure architecture of Regulation ATS (1998) and Regulation NMS (2007) was not designed to observe.^{2,3} Those two facts are frequently joined into a causal claim. This paper separates them.

Public regulatory reporting does not tag venues by the technology they deploy. No FINRA table, no Form ATS-N filing, and no enforcement order in the record assembled here distinguishes a venue that uses machine learning from one that does not, so a synthesis built on those sources cannot attribute what changed to a technology it never observes. The research question is therefore diagnostic rather than causal: what do the public data establish about off-exchange growth, execution quality, and enforcement, and what would have to be observed to support the machine-learning interpretation?

1.1 What the paper claims

Three claims are advanced, each tied to an evidence tier defined in Section 2.3. First, off-exchange growth between 2019 and January 2025 is concentrated in the non-ATS segment rather than in registered alternative trading systems when measured in share of volume; this is a Tier 2 result, and Section 3.1 shows that it reverses in dollar terms. Second, the execution-quality literature documents a stable difference in how retail and institutional order flow are treated, but the direction of that difference does not support reading it as a per-trade retail cost. Third, the enforcement record documents recurring opacity-enabled misconduct concentrated in 2015 to 2018, with no order alleging anything about a machine-learning system.

The normative argument is correspondingly narrow: the public price-formation process is now supported by a minority of consolidated volume, supervisory visibility into the venues holding the majority is structural rather than behavioral, and this is a governance problem whether or not machine learning is the proximate cause. The argument does not require, and this paper does not make, any causal claim about technology.

1.2 Classification and contribution

The manuscript is a Perspective with Original Data Synthesis: not Original Research, because it estimates no model and identifies no treatment effect, and not a Review, because roughly half its content is original compilation. The contribution is a reproducible assembly of public data plus an argument about what that assembly can bear. The six datasets and the calculation workbook are released as a Research Archive supplement (Supplementary Files S1 to S7), and every derived number carries a ledger entry in Appendix A giving its inputs, formula, and workbook cell.

2. Methods

2.1 Venue taxonomy

Public commentary routinely treats off-exchange trading as a near-synonym for dark pools. In United States market structure these are different objects with different regulatory statuses, economic functions, and counterparties, and conflating them is the single largest source of error in this area. Table 1 sets out the taxonomy used throughout.

Two consequences follow. Bilateral principal execution against a wholesaler is not dark-pool trading in the Regulation ATS sense, even though both print off-exchange and neither contributes pre-trade information to the public book. And a rise in the non-ATS bucket cannot be attributed to any firm, technology, or business model, because the bucket is an aggregate whose components are not separately reported.

2.2 Terminology and causal language

Vocabulary can carry a conclusion before any evidence is offered. A label such as *AI-enhanced bilateral routing* for the non-ATS OTC bucket asserts in a noun phrase precisely what the data cannot establish, and once adopted it propagates into

Category	Regulatory status	Typical operator and counterparty	Pre-trade transparency	FINRA reporting bucket
Lit exchange	National securities exchange, Regulation NMS	Registered exchange; all participants	Full displayed quotations	Exchange volume
ATS, dark pool (block or continuous cross)	Broker-dealer operating an ATS under Regulation ATS; Form ATS-N	Broker-dealer or agency operator; institutional subscribers matched against one another	None pre-trade; post-trade printed to the tape	ATS OTC volume, venue-level weekly detail
ATS, conditional or negotiation venue	Regulation ATS; Form ATS-N	Agency operator; institutional blocks negotiated on invitation	None pre-trade; indications not public	ATS OTC volume
Wholesale internalizer	Broker-dealer, no ATS registration; Rule 605 and 606 reporting	Market-making firm; principal counterparty to segmented retail flow	None; execution against the operator's own book at or inside the NBBO	Non-ATS OTC volume
Single-dealer platform	Broker-dealer, no ATS registration	Bank or dealer; principal counterparty to institutional clients	None; bilateral streaming prices	Non-ATS OTC volume
Other OTC market making	Broker-dealer	Dealer capital commitment, manual and upstairs trades	None	Non-ATS OTC volume

Table 1. Venue taxonomy used in this manuscript. The distinction that matters empirically is the last column: FINRA publishes ATS and non-ATS OTC volume separately, but the non-ATS bucket aggregates wholesale internalizers, single-dealer platforms, and other OTC market making, and identifies none of them individually. Authors' compilation from Regulation ATS, Form ATS-N, and FINRA reporting definitions.^{2,4,5}

every sentence that follows. Table 2 lists the formulations this paper avoids and the wording it uses instead.

Formulation avoided here	Wording used instead	Reason
AI-enhanced bilateral routing	Non-ATS off-exchange volume (wholesalers, single-dealer platforms, other OTC)	The published series identifies venue registration status, not technology. The old term asserted the paper's conclusion in its vocabulary.
AI has structurally accelerated fragmentation	Off-exchange share rose over a period in which firms documented machine-learning deployment	States the observed co-movement without asserting a direction of causation that the data cannot support.
Price-impact differential as an execution-cost gap	Price-impact differential as a measure of order-flow segmentation	A lower retail price impact indicates less adverse selection in that flow rather than a cost borne by retail investors; Section 3.2 explains why.
Enforcement clusters on the AI deployment wave	Enforcement clusters in 2015 to 2018, a period that also follows a supervisory initiative on dark pools	No order in the compiled record alleges conduct involving a machine-learning system, and the alternative explanation is at least as consistent with the timing. ⁶
Self-reinforcing informational asymmetry as a finding	Self-reinforcing informational asymmetry as a hypothesis, with an identifying test	The mechanism is plausible and testable, but nothing in the six datasets measures whether an operator's informational advantage grows with the flow it processes.

Table 2. Formulations avoided in this paper and the wording used instead. Authors' compilation.

2.3 Evidence tiering

To keep measurement separate from interpretation, every quantitative statement carries one of the three tiers defined in Table 3, applied in the text, in every table, and in the released datasets. Tier 3 statements never appear as numbers.

Tier	Definition	Example from this paper	Where it may appear
T1	Directly measured and published by a named regulator, exchange, or peer-reviewed study; transcribed without alteration	Off-exchange share of 51.8% in January 2025	Results, tables, figures
T2	Derived by the present authors by arithmetic on T1 values; formula and inputs disclosed in Appendix A and reproducible in Supplementary File S7	Non-ATS volume accounts for 86.3% of the 14.6 pp rise in off-exchange share	Results, tables, figures, always labeled
T3	Interpretive attribution linking observed change to an unobserved cause, including any statement connecting venue behavior to machine learning	The composition shift is consistent with routing optimization	Discussion only, in prose, never as a number

Table 3. Evidence-tiering protocol. Without an explicit tier, a Tier 3 attribution travels unnoticed inside a Tier 1 or Tier 2 sentence, which is the failure the protocol is designed to prevent. Authors' compilation.

2.4 Data sources, compiled datasets, and release

All base data are public. No proprietary trading data, non-public filing, or licensed commercial dataset was used, and none was available to the authors. Six datasets were compiled by reconciling figures across regulatory publications and the peer-reviewed and working-paper literature, and each is released as a machine-readable CSV file with a calculation workbook in which every derived statistic is a live formula. Table 4 is the manifest.

File	Dataset	Coverage	Primary sources	Rows
S1	Off-exchange volume and venue composition	2019 to Jan 2025	FINRA 2024 Industry Snapshot Tables 3.1.1.1 to 3.1.1.2; Nasdaq Market Intelligence; Rosenblatt Securities ^{1,4,7}	7
S2	Execution-quality measures	2019 to 2023	Dyhrberg et al.; Brown et al.; SEC Rel. 34-96495; Schwarz et al. ⁸⁻¹¹	7
S3	Enforcement actions	2011 to 2025	SEC, FINRA, and NYAG public releases ^{12,13}	16
S4	European venue shares, MiFID II episode	2017 to 2020	ESMA TRV 2/2018; ESMA ASR 2020; DLA Piper ¹⁴⁻¹⁶	8
S5	Deployment and rulemaking chronology	1998 to 2026	Federal Register; SEC releases; firm announcements ^{5,17}	11
S6	Compound scenario model output	30-year horizon	Original model, this paper	31
S7	Calculation workbook (7 sheets, 266 formulas)	All of the above	Derives every T2 statistic from its T1 inputs; includes robustness grids	n/a

Table 4. Supplementary data package. Files S1 to S6 are CSV, each carrying a tier and a source per record; S7 is a spreadsheet workbook whose Ledger sheet reproduces Appendix A with live formulas. A README documents every column, source URL, and replication step.

Two execution-quality sources are working papers that have not completed peer review.^{8,9} They are the best available evidence on retail execution quality, are labeled as preprints wherever used, and no claim rests on either alone.¹¹

2.5 Derivation ledger and robustness strategy

Every Tier 2 number is documented at source level. Appendix A lists, for each statistic, its identifier in the text, its inputs and their sources, the formula, the workbook cell in which it is evaluated, and the resulting value; a reader who disagrees with an input can change it in Supplementary File S7 and observe the consequence for every dependent figure.

Five robustness families are reported in Section 3 alongside the results they test, rather than collected where they can be skipped: endpoint, metric, window, attribution, and parameter sensitivity. Table 10 collects all ten tests.

2.6 What this design cannot establish

Identification statement

All six datasets are observational, and none contains a variable that distinguishes venues by the technology they deploy. The design cannot identify a causal effect of machine-learning adoption on fragmentation, execution quality, or enforcement incidence; it cannot rank the confounds listed below; and it cannot measure the incidence of misconduct, only its detected and settled fraction. Any sentence connecting an observed change to machine learning is a Tier 3 interpretation offered as a hypothesis, and Section 4.7 states the design that would test it.

Four alternative explanations for the migration fit the same series, and the design cannot separate them: the elimination of retail commissions from October 2019, which expanded retail participation and therefore wholesaler flow independently of technology; the retail trading surge of 2020 and 2021, a discrete demand shock with a documented start and decay; best-execution routing incentives created by Regulation NMS, which predate machine-learning deployment and already favored off-exchange execution for marketable retail orders;³ and payment-for-order-flow economics, which changed over the same window for reasons unrelated to matching technology. Section 4 argues that these account for a substantial share of the migration and that no part of the present evidence isolates a technology-specific increment.

3. Results

3.1 Venue composition of off-exchange growth

Table 5 reports the annual series and Figure 1 presents the decomposition. Total off-exchange share rose from 37.2% in 2019 to 51.8% in January 2025, an increase of 14.6 percentage points [T2, V1]. Over the same window the ATS share rose from about 14.0% to about 16.0%, an increase of 2.0 points [T2, V2], so the non-ATS share rose by 12.6 points [T2, V3]. Non-ATS volume therefore accounts for 86.3% of the increase and ATS volume for 13.7% [T2, V4 and V5], and the ratio of non-ATS to ATS share moved from 1.66 in 2019 to 2.24 in January 2025 [T2, V6].

Period	Total off-ex. [T1]	ATS share [T1]	Non-ATS share [T2]	Non-ATS \$/day [T1]	ATS \$/day [T1]	Ratio [T2]	Contemporaneous context
2019	37.2%	14.0%	23.2%	n/r	n/r	1.66×	Baseline; pre-zero-commission
2020	41.5%	14.2%	27.3%	\$130.4B	\$54.4B	1.92×	COVID-19 retail participation surge
2021	43.7%	14.0%	29.7%	\$153.9B	n/r	2.12×	Meme-stock episode (Q1)
2022	42.0%	13.5%	28.5%	\$181.0B	n/r	2.11×	Fed tightening cycle; non-ATS \$ peak
2023	44.0%	14.9%	29.1%	\$147.4B	\$68.8B	1.95×	Annual record off-exchange share
2024	50.0%	15.4%	34.6%	n/r	n/r	2.25×	First month above 50% (Nov 2024)
Jan 2025	51.8%	16.0%	35.8%	n/r	n/r	2.24×	Fifth consecutive monthly record

Table 5. Off-exchange volume and venue composition, 2019 to January 2025. ATS shares are published to approximately one decimal place and are treated as approximate throughout; the sensitivity of the decomposition to that approximation is reported below. Here n/r denotes not reported in the source table. The 2024 entry is an annual figure and the January 2025 entry is a single month. Sources: FINRA 2024 Industry Snapshot; Nasdaq Market Intelligence; Rosenblatt Securities.^{1,4,7}

Robustness

R1 and R2, endpoint and window sensitivity. The published ATS shares are approximate. Varying the 2019 endpoint over 13.5% to 14.5% and the January 2025 endpoint over 15.5% to 16.5% moves the non-ATS contribution across a band of 79.5% to 93.2%, so the qualitative conclusion survives the full grid but the figure of 86% should not be quoted to the




Figure 1. Composition of off-exchange volume and decomposition of its increase. Panel (a) splits total off-exchange share into its ATS and non-ATS components; the line tracks the total, and the label inside each ATS segment gives the non-ATS to ATS ratio for that period. Panel (b) decomposes the 14.6 percentage-point rise and shows the endpoint-sensitivity band described in the text. Authors' compilation from Supplementary File S1.

decimal. Restricting the comparison to full years, 2019 to 2023, gives a 6.8-point rise of which non-ATS accounts for 86.8%, so the result is not an artifact of the January 2025 endpoint.

R3, metric sensitivity. This is the check that most qualifies the earlier draft. Asked in dollar rather than share terms, the same question gives a different answer. Between 2020 and 2023 non-ATS OTC volume grew 13.0% in dollar terms while ATS volume grew 26.5% [T2, V7 and V9], compound annual rates of about 4.2% and 8.1% respectively [T2, V8 and V10], so non-ATS activity accounts for 54.1% of the dollar increase rather than 86%. Two features explain the divergence: the share and dollar series come from different FINRA tables covering different security universes and are not fully reconcilable, and the dollar window is shorter because FINRA reports a dollar breakdown for fewer years. The compositional claim therefore holds in share terms over the full window and fails in dollar terms over the shorter window for which dollar data exist.

Interpretation [T3]

The result is consistent with a routing story, in which marketable retail orders are increasingly internalized bilaterally rather than crossed in an ATS, and equally consistent with a participation story, in which more retail accounts trading more often mechanically enlarges the bucket that receives retail flow. Nothing in Table 5 distinguishes them, and the divergence between the share and dollar results is what one would expect if retail order count grew faster than retail notional. Neither reading is evidence of technology-driven routing optimization.

3.2 Execution quality and the provenance of each estimate

The request to separate original from prior estimates is answered directly: this paper produces no original execution-cost estimate. It has no trade-level data, no brokerage accounts, and no proprietary feed. Every level in Table 6 is transcribed from a cited study and the only original operation is subtraction. Figure 2 plots the two studies reporting comparable venue-type levels.

The 15.93 basis-point price-impact differential invites a natural but incorrect reading as a retail execution penalty, and the sign runs the other way. Price impact measures how far the market moves against a trade after execution, so a lower value indicates that the counterparty faces less adverse selection. Retail flow routed to wholesalers has a lower price impact precisely because it is uninformative, which is what makes it attractive to internalize and why wholesalers can quote inside the spread. The gap therefore measures order-flow segmentation, not a cost borne by retail investors, and Brown et al. point the same way with an adverse-selection component of 6% off-exchange against 73% on-exchange.⁹

Two measures do run against retail. Effective spreads are wider in the retail sample by 2.08 basis points, and the execution-

Measure	Retail or off-exchange	Institutional or on-exchange	Difference	Origin of the level	Origin of the difference	Tier
Effective spread (bps)	49.06	46.98	2.08, retail higher	Dyhrberg et al. (2023), preprint ⁸	Computed here	T1 / T2
Five-minute price impact (bps)	36.80	52.73	15.93, retail lower	Dyhrberg et al. (2023), preprint ⁸	Computed here	T1 / T2
Execution-to-quote ratio	0.76	0.97	0.21	Dyhrberg et al. (2023), preprint ⁸	Computed here	T1 / T2
All-in trading cost (bps)	1.8	4.6	2.8, retail lower	Brown et al. (2025), preprint ⁹	Computed here	T1 / T2
Adverse-selection share of cost	6%	73%	67 pp	Brown et al. (2025), preprint ⁹	Reported by source	T1
Marketable-order fill rate	69.0%	27.3%	41.7 pp	SEC Rel. 34-96495 ¹⁰	Reported by source	T1
Account-level round-trip cost	0.07% to 0.46% across brokers		0.39 pp dispersion	Schwarz et al. (2025), J. Finance ¹¹	Computed here	T1 / T2

Table 6. Provenance of every execution-quality estimate used in this paper. No level in this table is original to the present authors. Preprint status is stated where it applies.




Figure 2. Execution-quality measures by venue type. Panel (a) shows the two Dyhrberg et al. measures that move in opposite directions: retail effective spreads are slightly higher, retail price impact is substantially lower. Panel (b) decomposes the Brown et al. all-in cost into its adverse-selection and residual components, with each component drawn in the same color across both venue types. Authors' compilation from Supplementary File S2.

to-quote ratio of 0.76 against 0.97 indicates that retail orders capture a smaller fraction of the quoted spread than venue-level averages suggest. The most credible measure of what retail investors actually pay comes from Schwarz et al., who placed 85,000 simultaneous market orders across six accounts at five brokers and found mean account-level round-trip costs of 0.07% to 0.46% excluding commissions, the dispersion driven by wholesalers giving systematically different prices for identical trades routed by different brokers.¹¹ That 39 basis points per round trip between the best and worst broker is an order of magnitude larger than any venue-type average in Table 6, and it is attributable to broker routing arrangements rather than to venue technology.

Interpretation [T3]

Order-flow segmentation is the mechanism common to all four studies: uninformative flow is separated from informative flow and priced accordingly, producing better measured per-trade outcomes for the segmented flow and worse ones for what remains on the exchange. Whether machine learning intensifies that separation is exactly what the public data cannot answer, since segmentation

predates modern machine-learning execution by at least a decade¹⁸ and none of these studies observes venue technology.

3.3 Enforcement record

No public source aggregates dark-pool enforcement actions, so this study compiled one. Sixteen actions brought between 2011 and 2025 are recorded in Table 7 with a combined settlement value of \$330.15 million [T2, F1], and Figure 3 shows their distribution. Misrepresentation of venue operation is the stated violation in 11 of the 16 actions, unauthorized access for high-frequency participants in 7, misuse of subscriber order information in 7, and conflicts of interest in 4, most actions carrying more than one tag. Four actions totaling \$4.6 million fall between 2011 and 2014; eleven totaling \$320.55 million, or 97.1% of all settlement dollars, fall between 2015 and 2018; and one \$5.0 million action falls after 2018.

Year	Operator and venue	Settlement (\$M)	Primary stated violation	Agency	Venue type	ML cited
2011	Pipeline Trading / Pipeline ATS	1.00	Misrepresentation; conflict of interest	SEC	ATS (dark pool)	No
2012	eBX / LeveL ATS	0.80	Subscriber information misuse	SEC	ATS (dark pool)	No
2014	Liquidnet / Liquidnet ATS	2.00	Subscriber information misuse	SEC	ATS (dark pool)	No
2014	Goldman Sachs / Sigma X	0.80	HFT access; best execution	FINRA	ATS (dark pool)	No
2015	UBS Securities / UBS ATS	14.40	Misrepresentation; HFT access; information misuse	SEC	ATS (dark pool)	No
2015	ITG / AlterNet / POSIT	20.30	Conflict of interest; information misuse	SEC	ATS (dark pool)	No
2015	Goldman Sachs / Sigma X (OATS)	1.80	Misrepresentation (reporting)	FINRA	ATS (dark pool)	No
2016	Barclays Capital / LX	70.00	Misrepresentation; HFT access	SEC + NYAG	ATS (dark pool)	No
2016	Credit Suisse / Crossfinder	84.30	Misrepresentation; HFT access; information misuse	SEC + NYAG	ATS (dark pool)	No
2016	Deutsche Bank / SuperX+	37.00	Misrepresentation; conflict of interest	SEC + NYAG	ATS (dark pool)	No
2016	Deutsche Bank / SuperX (FINRA)	3.25	Disclosure failures	FINRA	ATS (dark pool)	No
2017	Citadel Securities	22.60	Misrepresentation	SEC	Non-ATS (wholesaler)	No
2018	Merrill Lynch / Instinct X	42.00	Misrepresentation; HFT access; conflict	SEC + NYAG	ATS (dark pool)	No
2018	ITG / POSIT (second action)	12.00	Misrepresentation; HFT access; information misuse	SEC	ATS (dark pool)	No
2018	Citigroup / CORE / Citi Match	12.90	Misrepresentation; HFT access	SEC	ATS (dark pool)	No
2025	Liquidnet / Liquidnet ATS	5.00	Information misuse; misrepresentation	SEC	ATS (dark pool)	No

Table 7. Compiled dark-pool and off-exchange enforcement record, 2011 to 2025. The Credit Suisse total comprises a \$30M SEC penalty, \$24.3M in disgorgement and interest, and a \$30M NYAG penalty; the Barclays settlement was divided equally between the SEC and the NYAG; ITG in 2015 was the first ATS operator to admit wrongdoing. The final column records the result of reading each order for any allegation involving a machine-learning or artificial-intelligence system. Authors' original dataset, Supplementary File S3.^{12,13}

Robustness

R4 to R7, attribution sensitivity. The \$330.15 million headline is convention dependent. Counting civil penalties only, and so excluding the \$24.3 million of disgorgement and interest in the Credit Suisse settlement, gives \$305.85 million; excluding the separately disclosed state components, \$35 million for Barclays and \$30 million for Credit Suisse, gives a federal-only total of \$265.15 million, the Deutsche Bank and Merrill Lynch splits not being separately disclosed; and excluding the single wholesaler action, which is not an ATS matter, gives \$307.55 million for registered dark pools alone. Moving the era boundary to 2016 still leaves 86.0% of settlement dollars in the cluster. The clustering is robust; the headline dollar figure should carry its convention.




Figure 3. Enforcement penalties by year and stated violation, with attribution sensitivity. Panel (a) stacks each year's actions by primary stated violation on a continuous year axis, so that years with no action appear as gaps; the shaded band marks the 2015 to 2018 cluster. Panel (b) shows how the headline total moves under four defensible accounting conventions. Authors' compilation from Supplementary File S3.

Two limits deserve emphasis. The sixteen actions are the detected and settled fraction of an unknown distribution, so the compilation is evidence of structural vulnerability rather than a measure of incidence. And the cluster follows the Securities and Exchange Commission's own 2014 to 2015 supervisory initiative on dark pools,⁶ which makes heightened enforcement attention at least as plausible an explanation for the timing as heightened misconduct.

Interpretation [T3]

The cluster is sometimes read as evidence that enforcement tracks machine-learning deployment, and the compiled record does not support that reading. No order alleges conduct involving a machine-learning system; the violations concern how venues described themselves to subscribers, who was allowed to trade in them, and what was done with subscriber order information. These are governance and disclosure failures available to an operator running entirely rule-based software.

3.4 The MiFID II displacement episode

The European Union's second Markets in Financial Instruments Directive, effective 3 January 2018, capped dark trading at 4% of any single instrument at any single venue and 8% across all dark venues over a rolling six-month window.¹⁹ The resulting episode is the clearest natural experiment available on what venue-type caps do. Dark-pool volume in suspended instruments collapsed from roughly 8% to 9% to approximately 0.15% by May 2018, did not reappear on lit venues, and returned above 5% within weeks of the first suspensions being lifted in September 2018, while reported OTC volume rose across the same window. Table 8 gives the series and Figure 4 plots it.

R8, reporting-regime sensitivity. The raw OTC rise from December 2017 to August 2018 is 7 percentage points, but MiFID II simultaneously widened the set of transactions subject to post-trade OTC reporting, so an unknown part is mechanical. If the reporting change accounts for anything between 0 and 4 points of the OTC level, the displacement-attributable rise falls to a band of 3 to 7 points; the dark-pool fall of 7.5 points is not subject to the same distortion. Only the direction survives, that capped volume moved to other opaque venues rather than to lit ones. No specific displacement magnitude, including the tripling of OTC volume sometimes cited from this episode, is separable from the reporting change.

The comparison with the United States is indicative rather than exact, because the two regimes classify venues differently. European authorities acted when aggregate dark trading reached roughly 8% to 9% of volume; United States off-exchange volume, comprising both registered ATS and non-ATS bilateral execution, passed 50% in November 2024 with no comparable cap in force.¹

Period	Phase	Dark pool [T1]	OTC [T1]	Source
Dec 2017	Pre-MiFID II baseline	8.5%	15.0%	ESMA TRV 2/2018; DLA Piper (2018)
Jan 2018	Regime live; DVC data delayed	8.5%	15.5%	ESMA TRV 2/2018
Mar 2018	First DVC suspensions (618 ISINs)	3.5%	22.5%	ESMA DVC register
May 2018	Peak OTC print (932 ISINs suspended)	0.15%	24.0%	DLA Piper (2018)
Aug 2018	Caps fully binding	1.0%	22.0%	ESMA ASR 2020
Sep 2018	Cohort 1 suspensions lifted	5.2%	21.5%	ESMA ASR 2020
Dec 2019	Post-episode normalization	8.0%	26.0%	ESMA ASR 2020
Mar 2020	COVID-19 volatility peak	9.0%	25.0%	ESMA ASR 2020

Table 8. European venue shares across the double volume cap episode. The OTC series from January 2018 onward reflects the widened post-trade reporting obligations introduced with MiFID II and is not directly comparable with the December 2017 baseline. Sources: ESMA TRV 2/2018; ESMA Annual Statistical Report 2020; DLA Piper (2018).¹⁴⁻¹⁶




Figure 4. MiFID II venue displacement. Panel (a) tracks dark-pool and OTC shares across all eight observations in Table 8, with the double volume cap period shaded. Panel (b) decomposes the December 2017 to August 2018 change and marks the portion of the OTC rise that cannot be separated from the widened reporting regime. Authors' compilation from Supplementary File S4.

3.5 Price-discovery thresholds in the published literature

Independent lines of evidence converge on the conclusion that dark trading degrades price discovery beyond identifiable thresholds. Comerton-Forde and Putnins find that informational efficiency deteriorates once dark trading exceeds roughly 10% of dollar volume, the move from 10% to 20% raising informational inefficiency by 10% to 15% of a standard deviation under least squares and 19% to 26% under instrumental variables.²⁰ The CFA Institute, examining 450 United States stocks, locates a second and higher threshold once more than half of trading occurs in undisplaced venues.²¹ A Bank for International Settlements working paper supplies a mechanism, finding that a material fraction of dark trading occurs at stale reference prices with the fastest participants capturing the benefit,²² consistent with the latency-race estimates of Aquilina, Budish, and O'Neill.²³ Degryse, de Jong, and van Kervel find from European data that dark trading impairs liquidity at the traditional market even where visible fragmentation improves aggregate liquidity,²⁴ and Wu reports a similar asymmetry in more recent United States data.²⁵

One qualification governs how much weight this convergence can bear. The thresholds are estimated on dark venue shares as those authors defined them, which in the United States means registered ATS volume, not the combined off-exchange total; applying a 10% ATS threshold to a 51.8% off-exchange figure conflates the two categories that Section 2.1 separates. On the narrower measure the United States sits at about 16%, above the first threshold and well below the second; on the

broader measure it sits above both. Both readings are reported rather than the more dramatic one alone.

3.6 Deployment and rulemaking chronology

Firms have documented machine-learning execution deployments publicly since at least 2014: Goldman Sachs Marquee in 2014, JPMorgan LOXM and Goldman Sachs Sigma X2 in 2017, UBS routing in 2018, signal generation from 2022, and agentic workflows in 2025, six in all.^{26,27} Over the same period a single ATS-specific rule took effect, Form ATS-N in 2018, twenty years after Regulation ATS.^{2,5} Four market-structure rules proposed in 2022 were among the fourteen withdrawn in June 2025, and the scope of the Consolidated Audit Trail came under review in 2026.¹⁷ Figure 5 places the two series side by side.




Figure 5. Documented firm deployments and enacted ATS-specific rulemaking, 2010 to 2026. Deployment dates are firm announcements and trade-press reports; rulemaking dates are from SEC releases and the Federal Register. The chronology records co-occurrence in time and carries no causal content. Authors' compilation from Supplementary File S5.

This chronology is weaker evidence than it appears. Deployment announcements are marketing documents as well as disclosures, so their dates record when a firm chose to publicize a capability, not when it began using one. Counting rules is a poor measure of regulatory activity, since it omits examinations, no-action positions, and enforcement, all of which respond faster than rulemaking. And six deployments against one rule is not a meaningful ratio, because the two series measure different things. What the chronology does establish is narrower: the last enacted rule specific to alternative trading system operation predates every deployment in the right half of the figure, and it requires disclosure of order-interaction rules rather than of the systems that now implement them.

3.7 An illustrative distributional scenario

Compound wealth models in this literature are usually anchored to a measured execution gap, and Section 3.2 explains why the 15.93 basis-point price-impact differential cannot serve that purpose. The model below is therefore a transparent calculator answering a conditional question: if a retail investor were to bear a sustained annual execution disadvantage of a given size, what would it cost over a retirement horizon? The drag is an input to be varied, not a quantity this paper claims to have measured. Table 9 reports four scenarios and Figure 6 their sensitivity.

R9 and R10, parameter sensitivity. The dollar shortfall is highly sensitive to the base return and the horizon; the percentage of terminal wealth forgone is not. A 0.15% drag costs 1.39% of terminal wealth over 10 years, 4.12% over 30 years, and 5.46% over 40 years, and at 30 years that percentage moves only from 4.20% to 4.08% as the base return varies from 5% to 8%. The dollar figures over the same range run from \$41,405 to \$96,517, so a headline such as the \$3,137 shortfall in Scenario C is a statement about the assumed return rather than about execution quality. Percentages are therefore reported first throughout.

Scenario	Annual drag	Basis for the parameter	Terminal value	Shortfall	Wealth forgone
A	0.00%	No drag; 7.00% base return, S&P 500 geometric mean 1993 to 2023	\$76,123	n/a	n/a
B	0.05%	Illustrative low case	\$75,063	\$1,060	1.39%
C	0.15%	Illustrative high case	\$72,985	\$3,137	4.12%
D	0.39%	Broker dispersion in round-trip cost, Schwarz et al. (2025), applied at one turnover per year ¹¹	\$68,224	\$7,898	10.38%

Table 9. Compound scenarios on a \$10,000 initial investment over 30 years, computed as terminal value equals principal times one plus base return less drag, raised to the horizon. Scenarios B and C are illustrative parameters, not estimates. Scenario D is the only scenario anchored to a measured quantity, and that quantity is dispersion across brokers rather than a level attributable to any venue type. Authors' original model, Supplementary File S6.




Figure 6. Compound scenarios and their sensitivity. Panel (a) plots the four scenarios of Table 9, with an inset covering years 25 to 30 where the paths separate visibly. Panel (b) reports terminal wealth forgone as a function of the assumed drag at four horizons; the relationship is close to linear in the drag and nearly invariant to the base return, which is why the percentage measure is more informative than the dollar shortfall. Authors' original model.

What this model is and is not

It is a compounding identity with an assumed parameter, containing no market data, no estimated coefficient, and no behavioral assumption beyond a constant drag and a constant return. It cannot establish that any retail investor bears any drag, and Section 3.2 points the opposite way for the average marketable retail order. Its role is to bound the stakes of the policy question, not to answer it.

4. Discussion

4.1 What the synthesis supports

Three conclusions survive the robustness program. The public price-formation process has operated on a minority of consolidated United States equity volume since November 2024. The market is segmented, with uninformative retail flow separated from institutional flow and priced accordingly, a structure documented across four independent studies. And the venues holding the majority of volume are subject to a disclosure regime whose last substantive amendment took effect in 2018 and which requires descriptions of order-interaction rules rather than of the systems that implement them.

A fourth conclusion carries the qualification established in Section 3.1. Measured in share terms, off-exchange growth

ID	Robustness test	Headline result	Result under the test	Verdict
R1	ATS endpoints varied by ± 0.5 pp	Non-ATS is 86.3% of the rise	79.5% to 93.2%	Survives; do not quote to the decimal
R2	Annual-only window, 2019 to 2023	Non-ATS is 86.3% of the rise	86.8%	Survives
R3	Dollar rather than share denomination	Non-ATS dominates the growth	54.1% of the dollar increase; ATS dollar CAGR 8.1% against 4.2%	Fails; claim restricted to share terms
R4	Era boundary moved to 2016	97.1% of dollars in 2015 to 2018	86.0% in 2016 to 2018	Survives
R5	Civil penalties only	\$330.15M total	\$305.85M	Convention dependent
R6	Federal component only	\$330.15M total	\$265.15M	Convention dependent
R7	Registered ATS venues only	\$330.15M total	\$307.55M	Survives
R8	Reporting-regime adjustment	EU OTC rose about 7 pp	3 to 7 pp attributable to displacement; 0 to 4 pp to the reporting change	Direction survives; magnitude not identified
R9	Drag and horizon varied	\$3,137 lost over 30 years	0.47% to 13.59% of terminal wealth	Parameter dependent by construction
R10	Base return varied 5% to 8%	\$3,137 lost over 30 years	\$1,814 to \$4,109 at the same 0.15% drag; 4.20% to 4.08% of terminal wealth	Dollar figure not informative; use percentages

Table 10. Consolidated robustness results. Every test is reproducible in the Sensitivity sheet of Supplementary File S7. CAGR denotes compound annual growth rate.

between 2019 and January 2025 has come predominantly from non-ATS venues, principally wholesale internalizers acting as principal against segmented retail flow rather than registered dark pools crossing institutional blocks. This distinction is more useful than the dark-versus-lit binary that dominates public discussion, because the two categories serve different clients, carry different obligations, and would respond differently to any intervention.

4.2 What the synthesis does not support

Three attributions commonly drawn from this material do not follow from it. That machine learning drove the composition shift: the data identify venue registration status, not technology, and the four confounds in Section 2.7 remain unranked. That the price-impact differential measures a retail execution penalty: it measures the informational content of retail flow, and the sign of that effect is favorable to retail at the per-trade level. And that the enforcement record tracks machine-learning deployment: it tracks a supervisory initiative and a category of governance failure that requires no machine learning to occur.

A fourth proposition is offered only as a hypothesis. An endogenous, self-reinforcing informational asymmetry would be a theoretical contribution: an advantage that grows with the flow it processes, stays invisible by design, and updates faster than any reporting regime can track. It is coherent and has a clear empirical signature, namely that execution quality at a venue should improve with cumulative flow conditional on order characteristics, but nothing in the six datasets tests it. The framework it would extend is worth naming: Akerlof, and Stiglitz and Weiss, both assume an informational advantage that is stable and individually exploitable.^{28,29} An advantage that compounds with volume would be a different object, and demonstrating that one exists remains work to be done.

4.3 The positive and the normative claim

The strongest objection to any distributional argument here is that retail investors currently receive better per-trade execution than they would on a lit exchange, and it is correct on its own terms: effective spreads at wholesalers are competitive, fill rates on marketable orders are substantially higher, and adverse selection borne by retail flow is low.⁸⁻¹⁰

The normative claim is separate. Segmentation delivers a per-trade benefit in exchange for structural exclusion from the process that sets the reference price against which those trades settle: retail marketable orders receive prices derived from a national best bid and offer formed on venues holding a shrinking minority of volume, while the flow that would otherwise

contribute to that formation is internalized. Whether that exchange is a good one is a question about the value of price formation as a public good, not about basis points, and the execution-quality literature cannot settle it in either direction. Reasonable analysts weigh it differently and this paper does not claim the question is closed. The narrower point the evidence does support is that retail investors have limited visibility into how their orders are segmented and internalized, even where the measured per-trade outcome is favorable to them.

4.4 Why the current disclosure architecture is a poor fit

Regulation ATS and Form ATS-N are built on the premise that operators can describe their venues in prose that remains accurate between filings.^{2,5} That premise holds for static matching rules and fails for systems whose parameters are re-estimated continuously. Form ATS-N requires an operator to state how orders interact and who may access the venue. It does not require disclosure of what models are used, what data they consume, how often they update, or how the operator defines toxic order flow, which is the definition that determines who is excluded and on what basis. This is a gap of kind rather than of degree.

The wholesaler segment sits in a weaker position still. It falls outside Regulation ATS entirely, so no ATS-N filing exists for it, and its execution quality is observed through venue-level Rule 605 aggregates that the Schwarz et al. results show can mask large differences between brokers routing to the same wholesaler.¹¹ The 2024 amendments to Rule 605 improve granularity from 2025 onward and are the single most useful development for future work here.³⁰ The MiFID II episode supplies the constraint any remedy must respect: caps imposed on a venue type relocate opacity rather than removing it.

4.5 Policy implications

The recommendations below follow from the Tier 1 and Tier 2 findings, not from the retired attributions. None would eliminate off-exchange trading, which serves documented institutional needs, including the reduction in institutional execution costs O'Hara and Ye estimated at 20% to 30% and that Zhu's model attributes to the migration of uninformed flow.^{31,32} The aim is supervisory visibility proportionate to the share of price formation now conducted out of view.

1. **Extend systems disclosure to cover models, not only rules.** Form ATS-N should require the type, general purpose, update cadence, and known limitations of any automated system that determines order interaction or participant access, on the risk-tiered model of the European Union Artificial Intelligence Act and without requiring model weights or training data.³³ The threshold should be whether a system materially determines who trades with whom, not whether it meets a technical definition of machine learning, which is the definitional problem this paper met in its own data.
2. **Extend the same obligation to non-ATS bilateral execution.** Section 3.1 implies that a rule reaching only registered ATS venues would address a shrinking share of off-exchange volume. Wholesale internalizers should file an equivalent disclosure covering order handling, price improvement methodology, and the basis on which flow is accepted or rejected.
3. **Prioritize granular execution-quality reporting over new venue caps.** The amended Rule 605 regime, with its order-size breakdowns, will support the broker-level comparisons that Schwarz et al. showed matter more than venue-level averages, without creating the displacement incentive the MiFID II experience documents.^{11,30} Related tick-size and access-fee amendments should be assessed on the same evidentiary standard, which industry commenters argued was not met in the parallel best-execution proposal.^{34,35} Supervisors should also retain the Consolidated Audit Trail, the only dataset capable of answering the question in Section 4.7, and implement its research-access provisions.
4. **Coordinate disclosure standards internationally.** The MiFID II episode showed that unilateral venue-type restrictions produce cross-border and cross-venue arbitrage. The International Organization of Securities Commissions' work on artificial intelligence in capital markets is the natural home for a common disclosure taxonomy that would at minimum make venue technology observable to researchers and supervisors, which it currently is not.³⁶

4.6 Limitations

Beyond the identification limit in Section 2.7, five limitations constrain the conclusions. The analysis has no order-level data and observes venue behavior only through published aggregates, so within-venue heterogeneity is invisible. The execution-quality evidence rests substantially on two preprints, one of whose results this paper reinterprets rather than replicates.

The enforcement compilation observes settled actions only. The European comparison spans regimes that classify venues differently, and the reporting change makes the OTC series discontinuous. And the compound model's central parameter is an assumption, which is why its output is a conditional percentage rather than a dollar loss.

4.7 A design that would identify the mechanism

The following design is stated entirely in the conditional; it has not been carried out. A stacked difference-in-differences study using Consolidated Audit Trail microdata would compare execution-quality outcomes at venues operated by firms with dated, publicly documented machine-learning deployments against venues with no such announcement in the same window, stacking three events with disclosed dates: Goldman Sachs Sigma X2 in May 2017, JPMorgan LOXM in July 2017, and UBS routing in 2018. Outcomes would be venue-level effective spreads, price impact, fill rates, and the execution-to-quote ratio, with stock and venue fixed effects, a parallel-trends test on twelve pre-event months, and placebo tests on false deployment dates.

Two threats should be stated in advance. Announcement dates are endogenous, since firms publicize capabilities when it suits them and may have deployed earlier, so the estimand is the effect of announced deployment, which bounds rather than equals the effect of deployment. And control venues are unlikely to be untreated, because execution technology diffuses through vendors and staff mobility, biasing any estimated effect toward zero. A null result would therefore be weak evidence against the mechanism, while a positive result surviving the placebo tests would be strong evidence for it. A related question is whether the execution-quality gap breaks structurally around the 2017 to 2019 window or is continuous with pre-existing segmentation trends.³⁷

5. Conclusion

The majority of United States equity volume now executes in venues that publish no pre-trade information, and in share terms the growth has come predominantly from bilateral internalization rather than from the registered dark pools that dominate public discussion. That distinction is this paper's central empirical contribution, and its limits are documented rather than asserted: it holds in share terms across every endpoint and window tested, and fails in dollar terms over the shorter period for which dollar data exist.

What the paper declines to claim matters as much as what it does. Machine learning is present throughout modern execution infrastructure and is not identifiable in any public dataset used here, so attributing a measured compositional shift to an unmeasured technology is the error the analysis is built to avoid. The argument is stronger without it: the governance question raised by opaque majority-share venues is the same whether the opacity is produced by a neural network or a lookup table, and it is not answered by disclosure rules written for venues whose operating logic could be described in a paragraph and stay accurate for a decade.

The evidence needed to close the gap exists. Consolidated Audit Trail microdata would support the design in Section 4.7, and the amended Rule 605 regime will support broker-level and order-size-level comparisons from 2025 onward. Whether those data reach researchers, and whether the audit trail is retained, are policy decisions currently under consideration.¹⁷ They will determine whether the question this paper leaves open can be answered by anyone outside the firms that already know the answer.

BACK MATTER

Source Reliability

Every citation was verified against the published record. Three sources whose authors have an interest in the regulatory outcome are used only for factual matters such as launch dates: trade-press deployment reporting, a market-maker publication, and the Rosenblatt venue series.^{7,26,27} The CFA Institute threshold analysis is dated to its original publication and predates the market structure it is used here to evaluate.²¹

Abbreviations

ATS	Alternative trading system	ML	Machine learning
bps	Basis points	MTF	Multilateral trading facility
CAGR	Compound annual growth rate	NBBO	National best bid and offer
CAT	Consolidated Audit Trail	NMS	National Market System
DVC	Double volume cap (MiFID II)	NYAG	New York Attorney General
ESMA	European Securities and Markets Authority	OTC	Over the counter
FINRA	Financial Industry Regulatory Authority	pp	Percentage points
HFT	High-frequency trading	SDP	Single-dealer platform
IOSCO	International Organization of Securities Commissions	SEC	United States Securities and Exchange Commission
ISIN	International securities identification number	SI	Systematic internalizer
MiFID II	Markets in Financial Instruments Directive II		

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not for profit sectors.

Conflicts of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics Statement

Ethics approval was not required. The study involved no human participants, no private identifiable data, and no animal subjects, and used only data that were already publicly available at the time of collection.

Data Availability Statement

The compiled datasets and calculation workbook are provided as supplementary files with this submission. Supplementary Files S1 to S6 contain the six compiled datasets as CSV files with a tier and a source identified for every record. Supplementary File S7 is the calculation workbook in which every derived statistic reported in the manuscript is evaluated as a live formula from its inputs, together with the robustness grids of Table 10 and a Ledger sheet reproducing Appendix A. A README accompanies the package and documents every column, source URL, and replication step. All underlying sources are public. On acceptance the package will be deposited in a public repository under CC BY 4.0 and the permanent identifier added to this statement.

AI Use Statement

The authors used Claude (Anthropic) for drafting assistance, editorial revision, figure and workbook code, and the reformatting of references into Vancouver style. All AI-generated content was reviewed, edited, and verified by the authors, who take full responsibility for the accuracy and integrity of the manuscript. Numerical results were independently recomputed by the authors from the compiled datasets.

Authors' Contributions

Z. Bari contributed to conceptualization, dataset compilation for the enforcement and European venue records, and writing of the original draft. K. Ragunath contributed to methodology, the evidence-tiering protocol, robustness analysis, and visualization. S. Tripathi contributed to conceptualization, the volume and execution-quality datasets, the compound scenario model, project management, and writing, review, and editing. All authors reviewed and approved the final manuscript and the supplementary data package.

Acknowledgements

The authors thank the reviewing editor, whose observation that the framing attributed market changes to artificial intelligence more directly than the underlying datasets allow reshaped this paper.

References

1. Nasdaq Market Intelligence. Off-exchange trading increases across all types of stocks. Published February 2025. Accessed August 2026. <https://www.nasdaq.com/articles/exchange-trading-increases-across-all-types-stocks>
2. US Securities and Exchange Commission. Regulation of Exchanges and Alternative Trading Systems (Regulation ATS). Release No. 34-40760; 1998.
3. US Securities and Exchange Commission. Regulation NMS. Release No. 34-55160; 2007.
4. Financial Industry Regulatory Authority. 2024 Industry Snapshot: market activity tables 3.1.1.1-3.1.1.2. FINRA; 2024.
5. US Securities and Exchange Commission. Regulation of NMS Stock Alternative Trading Systems (Form ATS-N). Release No. 34-83663; 2018.
6. White MJ. Shedding light on dark pools. Speech, US Securities and Exchange Commission; June 5, 2015. <https://www.sec.gov/newsroom/speeches-statements/shedding-light-dark-pools>
7. Rosenblatt Securities. Let There Be Light: US edition, issues 47 and 54. Rosenblatt Securities; 2023.
8. Dyhrberg AH, Shkilko A, Werner IM. The retail execution quality landscape. SSRN Working Paper No. 4313095; 2023. Preprint, not peer reviewed. <https://ssrn.com/abstract=4313095>
9. Brown D, Johnson TL, Kothari SP, So EC. Anatomy of trading costs for retail investors: savings from off-exchange execution. SSRN Working Paper No. 3976300; 2025. Preprint, not peer reviewed. <https://ssrn.com/abstract=3976300>
10. US Securities and Exchange Commission. Order Competition Rule. Release No. 34-96495; 2022.
11. Schwarz C, Barber BM, Huang X, Jorion P, Odean T. The “actual retail price” of equity trades. J Finance. 2025;80(4):2507-2541. doi:10.1111/jofi.13467
12. US Securities and Exchange Commission. SEC charges Barclays and Credit Suisse for dark pool violations. Press Release 2016-12; 2016. <https://www.sec.gov/newsroom/press-releases/2016-12>
13. US Securities and Exchange Commission. SEC charges ITG with misleading dark pool subscribers. Press Release 2018-256; 2018. <https://www.sec.gov/newsroom/press-releases/2018-256>
14. European Securities and Markets Authority. Trends, Risks and Vulnerabilities Report No. 2/2018. ESMA; 2018.
15. European Securities and Markets Authority. Annual Statistical Report on EU Securities Markets 2020. ESMA50-165-1355; 2020.
16. DLA Piper Intelligence. The impact of MiFID II on dark pools so far. DLA Piper; November 2018.
17. US Securities and Exchange Commission. Withdrawal of 14 pending proposed rules. Press release; June 12, 2025. <https://www.sec.gov>
18. Hatheway F, Kwan A, Zheng H. An empirical analysis of market segmentation on US equity markets. J Financ Quant Anal. 2017;52(6):2399-2427. doi:10.1017/S0022109017000849
19. Oxford Business Law Blog. Post MiFID II, dark trading should return to basics. University of Oxford Faculty of Law; January 2018. <https://blogs.law.ox.ac.uk/business-law-blog/blog/2018/01/post-mifid-ii-dark-trading-should-return-basics>
20. Comerton-Forde C, Putnins TJ. Dark trading and price discovery. J Financ Econ. 2015;118(1):70-92. doi:10.1016/j.jfineco.2015.06.013
21. CFA Institute. Dark Pools, Internalization, and Equity Market Quality. CFA Institute; 2012.
22. Bank for International Settlements. Sharks in the dark: quantifying HFT dark pool latency arbitrage. BIS Working Papers No. 1115; 2023. <https://www.bis.org/publ/work1115.htm>
23. Aquilina M, Budish E, O'Neill P. Quantifying the high-frequency trading arms race. Q J Econ. 2022;137(1):493-564. doi:10.1093/qje/qjab032
24. Degryse H, de Jong F, van Kervel V. The impact of dark trading and visible fragmentation on market quality. Rev Finance. 2015;19(4):1587-1622. doi:10.1093/rof/rfu027
25. Wu WS. Understanding the impacts of dark pools on price discovery. J Financ Mark. 2024;68:100800. doi:10.1016/j.finmar.2023.100800
26. Institutional Investor. AI jumps into dark pools. Institutional Investor; December 7, 2017.
27. Citadel Securities. Market Lens. Citadel Securities; September 2022.
28. Akerlof GA. The market for lemons: quality uncertainty and the market mechanism. Q J Econ. 1970;84(3):488-500. doi:10.2307/1879431
29. Stiglitz JE, Weiss A. Credit rationing in markets with imperfect information. Am Econ Rev. 1981;71(3):393-410.

30. US Securities and Exchange Commission. Disclosure of Order Execution Information (amended Rule 605). Release No. 34-99679; 2024.
31. O'Hara M, Ye M. Is market fragmentation harming market quality? J Financ Econ. 2011;100(3):459-474. doi:10.1016/j.jfineco.2011.02.006
32. Zhu H. Do dark pools harm price discovery? Rev Financ Stud. 2014;27(3):747-789. doi:10.1093/rfs/hht078
33. European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Off J Eur Union; 2024.
34. US Securities and Exchange Commission. Regulation NMS: minimum pricing increments, access fees, and transparency of better-priced orders. Rule 610 and 612 amendments; 2024.
35. Securities Industry and Financial Markets Association. Rethinking the economic analysis in the SEC's best execution proposal. SIFMA; August 2024.
36. International Organization of Securities Commissions. Artificial intelligence in capital markets: use cases, risks, and challenges. IOSCO; 2024.
37. International Organization of Securities Commissions. Findings and observations during COVID-19-induced market stresses. IOSCO; August 2021.

APPENDIX A. DERIVATION LEDGER

Table A1 lists every Tier 2 statistic reported in the manuscript with its inputs, formula, evaluating cell in Supplementary File S7, and value. Inputs are Tier 1 unless otherwise noted.

ID	Statistic as reported in the text	Inputs and their sources	Formula	Workbook cell	Value
V1	Off-exchange share rose 14.6 pp, 2019 to Jan 2025	51.8%; 37.2% (FINRA / Nasdaq)	end less start	Volume!B12	14.60 pp
V2	ATS share rose 2.0 pp	16.0%; 14.0% (FINRA)	end less start	Volume!B13	2.00 pp
V3	Non-ATS share rose 12.6 pp	V1; V2	V1 less V2	Volume!B14	12.60 pp
V4	Non-ATS is 86.3% of the increase	V3; V1	100 times V3 divided by V1	Volume!B15	86.30%
V5	ATS is 13.7% of the increase	V4	100 less V4	Volume!B16	13.70%
V6	Non-ATS to ATS ratio, 1.66 / 1.95 / 2.24	non-ATS share; ATS share	non-ATS divided by ATS	Volume!G5, G9, G11	1.66; 1.95; 2.24
V7	Non-ATS dollar volume grew 13.0%, 2020 to 2023	\$147.4B; \$130.4B per day	100 times (end divided by base, less 1)	Volume!B17	13.04%
V8	Non-ATS dollar CAGR 4.2%, 2020 to 2023	\$147.4B; \$130.4B per day	100 times ((end divided by base) to the power one third, less 1)	Volume!B18	4.17%
V9	ATS dollar volume grew 26.5%, 2020 to 2023	\$68.8B; \$54.4B per day	100 times (end divided by base, less 1)	Volume!B19	26.47%
V10	ATS dollar CAGR 8.1%, 2020 to 2023	\$68.8B; \$54.4B per day	100 times ((end divided by base) to the power one third, less 1)	Volume!B20	8.14%
R1	Endpoint sensitivity band for V4	ATS endpoints 13.5 to 14.5 and 15.5 to 16.5	100 times (V1 less assumed ATS change) divided by V1	Sensitivity!B5:D7	79.5% to 93.2%
R2	Annual-window check on V4	2019 and 2023 shares	as V4 on the 2019 to 2023 window	Volume!B21	86.76%
R3	Dollar-weighted check on V4	FINRA dollar series, 2020 and 2023	100 times non-ATS dollar increase divided by total dollar increase	Volume!B22	54.14%
E1	Price-impact differential of 15.93 bps	52.73; 36.80 bps (Dyhrberg et al.)	institutional less retail	Execution!C13	15.93 bps
E2	Retail effective spread higher by 2.08 bps	49.06; 46.98 bps (Dyhrberg et al.)	retail less institutional	Execution!F5	2.08 bps
E3	Broker dispersion of 0.39 pp per round trip	0.46%; 0.07% (Schwarz et al.)	high less low	Execution!F11	0.39 pp
F1	16 actions totaling \$330.15M	16 settlement values (SEC, FINRA, NYAG)	sum of settlements	Enforcement!B22	\$330.15M
F2	2015 to 2018 subtotal of \$320.55M, 97.1%	settlement values by year	SUMIFS on year, then share of F1	Enforcement!B25, B26	\$320.55M; 97.09%
R4	Era boundary moved to 2016 to 2018	settlement values by year	SUMIFS on year	Enforcement!B28	\$284.05M; 86.04%
R5	Civil penalties only	F1; \$24.3M disgorgement	F1 less disgorgement	Enforcement!B29	\$305.85M
R6	Federal component only	F1; disclosed NYAG amounts	F1 less known NYAG components	Enforcement!B30	\$265.15M
R7	Registered ATS venues only	settlement values by venue type	SUMIFS on venue type	Enforcement!B31	\$307.55M
M1	EU OTC share rose 7.0 pp, Dec 2017 to Aug 2018	22.0%; 15.0% (ESMA, DLA Piper)	Aug 2018 less Dec 2017	MiFID!C14	7.00 pp
M2	EU dark share fell 7.5 pp over the same window	8.5%; 1.0% (ESMA)	Dec 2017 less Aug 2018	MiFID!C15	7.50 pp
R8	Reporting-adjusted OTC rise	M1; assumed reporting effect of up to 4 pp	M1 less assumed reporting effect	MiFID!C17	0 to 4 pp
C1	Scenario A terminal value	\$10,000; 7.00%; 30 years	principal times (1 plus r) to the power n	Compound_Model!B4	\$76,122.55
C2	Scenario C terminal value	\$10,000; 7.00% less 0.15%; 30 years	principal times (1 plus r less d) to the power n	Compound_Model!D43	\$72,985.37
C3	Scenario C shortfall	C1; C2	C1 less C2	Compound_Model!D4	\$3,137.18
C4	Scenario C wealth forgone	C1; C2	100 times (1 less C2 divided by C1)	Compound_Model!D46	4.12%
R9	Drag by horizon grid	drag 0.05 to 0.39%; horizons 10 to 40 years	100 times (1 less ((1 plus r less d) divided by (1 plus r)) to the power h)	Sensitivity!B13:E17	0.47% to 13.59%
R10	Drag by base-return grid	base 5 to 8%; drag 0.05 to 0.39%	principal times (1 plus r less d) to the power 30	Sensitivity!B21:E23	\$38,654 to \$99,238; 4.20% to 4.08% forgone at 0.15% drag

Table A1. Derivation ledger. Reproduced with live formulas in the Ledger sheet of Supplementary File S7.

Supplementary Material

Table A2 lists the files accompanying the manuscript, cited in the text as Supplementary Files S1 to S7.

File	Format	Contents
S1_offexchange_volume.csv	CSV	Volume and venue composition, 2019 to January 2025, with tier and source per record
S2_execution_quality.csv	CSV	Execution-quality measures with the origin of each level and each difference
S3_enforcement_actions.csv	CSV	Sixteen enforcement actions with venue type, agency, NYAG component, and ML-reference field
S4_mifid_venue_shares.csv	CSV	European venue shares across the double volume cap episode, with a comparability caveat per record
S5_deployment_and_rulemaking_timeline.csv	CSV	Deployment and rulemaking chronology with an evidence class per record
S6_compound_model.csv	CSV	Year-by-year output of the four compound scenarios
S7_calculation_workbook.xlsx	XLSX	Seven sheets, 243 live formulas: Volume, Execution, Enforcement, MiFID, Compound_Model, Sensitivity, Ledger

Table A2. *Supplementary file manifest.*