Proceedings of the 5th International Conference on Financial Technology and Business Analysis (ICFTBA 2026)
DOI: 10.54254/2754-1169/2026.39251

Repetition, Weighting, and the Pricing of Technical Event Risk: Evidence from Commercial Space Options

Chengze Song[1,*], Zequn Song[2]

[1]*The Village School, Houston, USA*
[2]*Princeton University, Princeton, USA*
**Corresponding Author. Email: zesong999x@gmail.com*

Abstract. A failed space mission can wipe out much of the issuer's market value within two trading days. After removing independent news, the absolute abnormal return is 13.2% for failures and 5.9% for successes. Pre-mission implied volatility, or AIV, uses days (-10, -1) for the disclosure screen and (0, +1) for the actual stock reaction. Missions of the same issuer within 20 trading days is treated separately. Under the AIV-aligned screen, the standardized task-weighted AIV is +0.040, and an independent reconstruction with a maximum δ error of 0.10 gives +0.028. Across four issuers, the smallest exact sign p-value is 0.125, and no main test family passes the Holm correction. Schedule-verified non-overlapping missions show a higher premium, +0.123 standardized and +0.188 raw, but this fades under the disclosure screen. The evidence support a conditional, not universal, pre-mission pricing mechanism. Launch cycles and overlaps matter.

Keywords: equity options, implied volatility, scheduled event risk, repeated events, commercial space

1. Introduction

Regular major events let option prices reflect how the market values uncertainty before information disclosure, like earnings announcements and elections, and contracts can cover the event or not . But many studies ignore a design problem: some firms face only one key event, while others face dozens of similar operational events, so the pooled analysis are dominated by the latter. Even when research focus on typical issuers, event-weighted and issuer-weighted estimators can diverge for economic reasons [1]. But what economic unit the estimator represents, and whether this signal survive empirical information screens. Public information auditing is the key step. We keep (-5, +3) as a diagnostic window, and define (-10, -1) for the AIV screen plus (0, +1) for the reaction screen. Official SEC filings and issuer announcements are linked at event level.

2. Conceptual framework and predictions

Option prices embed risk-neutral expectations and risk premia, not physical probabilities . A scheduled event can add a discrete variance component to any contract that spans it [2]. Following the event-variance logic developed for earnings announcements , let total implied variance satisfy

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.39251

W(T) = vT + J, where v is background variance per year and J is a discrete event component common to the 30- and 91-day maturities. Solving the two-maturity system gives

$$J_{30,91}=[91\cdot W(30)-30\cdot W(91)]/61, W(T)=IV(T)^2\cdot T/365. \quad (1)$$

Let y_ie be a pre-event risk measure for event e at issuer i. The mission-weighted estimator averages all events, Σ_i Σ_e y_ie / Σ_i n_i,The issuer-equal estimator just averages the issuer means, I^{-1} Σ_i (n_i^{-1} Σ_e y_ie). The first one tells us what happens before the average mission in the observed event population, while the second tells what happens before the average issuer's mission. They coincide only under restrictive balance or homogeneity assumptions, which rarely holds in real data. When cadence itself changes risk, through operational learning, investor habituation, smaller incremental cash-flow news, or overlapping option windows, their divergence becomes economically meaningful. Hiding it behind a single pooled coefficient simply discards the finding.

3. Data and event-level audit

ATM IV averages the qualifying call and put with the closest delta, only when both absolute delta errors are within 0.10. For each event, the audit records pre-mission information sources, first US reaction date, three window classifications, and materiality judgments with written notes. EDGAR mappings stay fixed for the earlier (-5, +3) window, and the new t = −10 to t = −6 period are built from SEC and issuer sources.

Table 1. Issuer-level event and usable-option counts

Issuer	Events	Adverse	Schedule verified	AIV-aligned clean	Non-overlap	Std AIV N	Strict raw AIV N
ASTR	5	3	5	3	5	5	2
ASTS	2	0	2	0	2	2	2
LUNR	1	1	1	1	1	0	0
MNTS	3	1	3	1	3	1	0
RKLB	36	1	34	22	5	36	24
SPCE	11	1	11	7	7	10	10
VORB	3	1	3	1	3	2	1
Total	61	8	59	35	26	56	39

As shown in Table 1 .OptionMetrics IvyDB US provides contract-level daily closing quotes, volume, open interest, implied volatility, Greeks, expiry, strike, and standardized volatility surface data. Security name history gives permanent SECID mappings across SPAC predecessors, ticker changes, reverse splits, and delistings. CRSP provides daily stock returns and value-weighted market returns [3]. Coverage runs to January 31, 2025, so late-2024 missions keep their post-event observations, and licensed raw files cannot be redistributed. Raw quotes must satisfy bid > 0, ask > bid, midprice > 0.05, relative spread ≤ 50%, days to expiry between 7 and 180, IV between 0.05 and 5.00, contract size equal to 100, zero special settlement, and open interest ≥ 10 or volume ≥ 1.

DOI: 10.54254/2754-1169/2026.39251

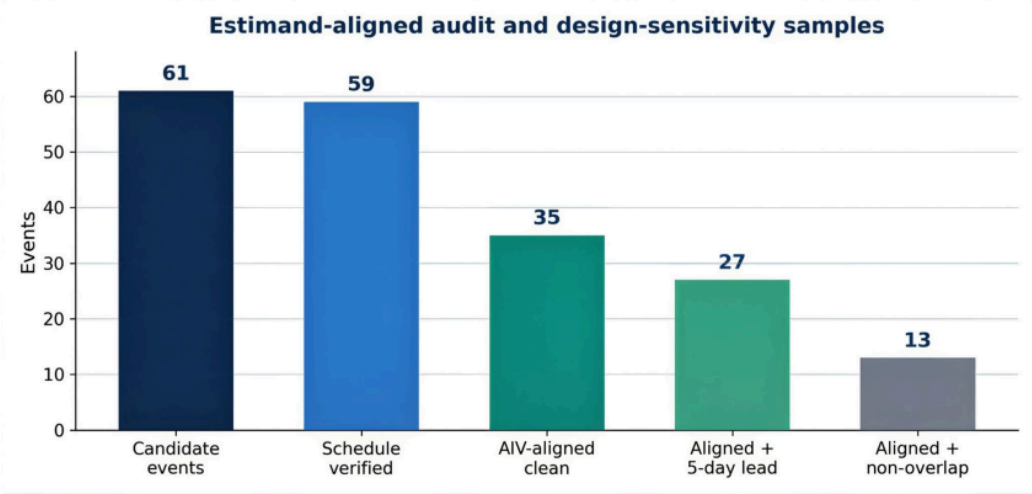

Figure 1. Estimand-aligned audit and design-sensitivity sample counts. Lead time and non-overlap are separate design restrictions, not folded into one preferred sample

As shown in the figure 1 .Among the 34 events classified as "clean" in the earlier wide screen, 10 were announced less than 10 calendar days before the mission, and 4 were announced less than 5 days before. On the issuer trading calendar, 42 verified missions contain at least 5 known trading days, and 27 contain at least 10 trading days. The disclosure screen consistent with the AIV standard keeps 35 events, while adding the five-day lead rule leaves 27 events [4]. Also, 35 of the 61 missions have other events of the same issuer within 20 trading days.

4. Estimation results and design sensitivities

As shown in Table 2 .The full sample preserves the weighting contrast: standardized issuer-equal AIV is +0.132, and the corresponding strict raw estimate is +0.165. Attribution then chan-ges the reading. In the AIV-aligned sample the standardized estimates are +0.040 mission weigh-ted and +0.065 issuer equal; the strict raw estimates are +0.028 and +0.137. The standardized iss-uerequal cluster p-value is 0.030. Its exact sign p-value, however, is 0.125, which is also the sm-allest value four issuers can produce.

Table 2. Primary abnormal-IV family under estimand-aligned attribution

Surface	Weighting	Sample	Estimate	Cluster p	Exact p	Min exact p	Holm p	N	G
Standardized	Issuer equal	Full	+0.132	0.037	0.0312	0.0312	—	6	6
Raw, δ≤0.10	Mission wt.	AIV-aligned	+0.028	0.591	0.3750	0.1250	0.591	23	4
Raw, δ≤0.10	Issuer equal	Full	+0.165	0.043	0.1250	0.0625	—	5	5
Raw, δ≤0.10	Issuer equal	AIV-aligned	+0.137	0.202	0.2500	0.1250	0.403	4	4

The remaining adjusted p-values are larger still. The primary family thus offers partial positive evidence that does not stand up to surface construction, small-G exact inference, and multiplicity all at once.

DOI: 10.54254/2754-1169/2026.39251

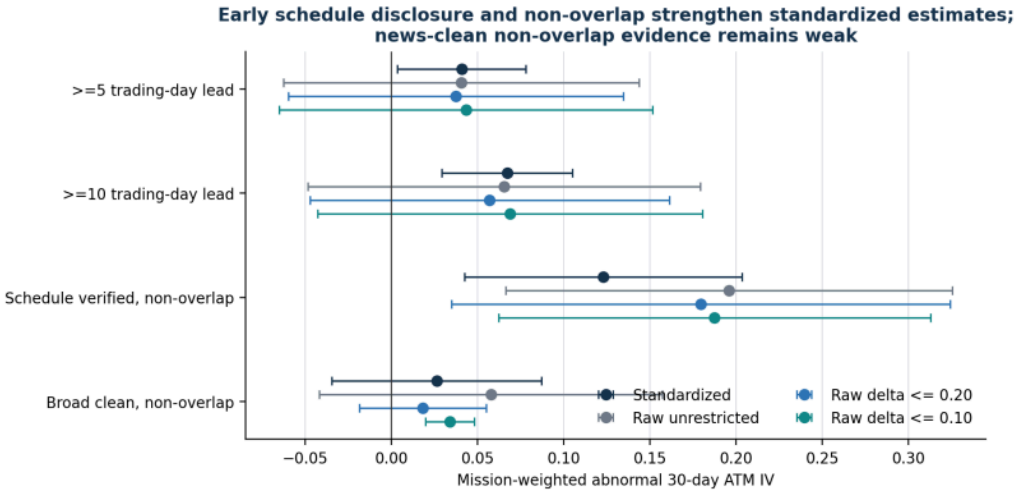

Figure 2. Lead-time, overlap, and raw-chain delta sensitivities for mission-weighted AIV. The broad clean non-overlap row contains only two strict-raw issuers

As shown in the figure 2 .Both restrictions move the estimates, though in different ways once the disclosure screen enters， gives standardized mission-weighted AIV of +0.052 (cluster p = 0.068; exact p = 0.125; N = 26) and strict-raw AIV of +0.037 (cluster p = 0.692; exact p = 0.750; N = 17). Among missions known at least ten trading days in advance, the full-window standardized estimate is +0.067 (cluster p = 0.018); the strict raw estimate is +0.069 with p = 0.293 and an exact p-value of 0.125, and the aligned ≥10-day row of Table 3 shows +0.073 against +0.030 on strict raw chains [5]. Schedule-verified non-overlapping missions produce the largest estimates, +0.123 standardized and +0.188 strict raw, with cluster p-values of 0.030 and 0.043. Yet the broad-screen clean non-overlapping estimates collapse to +0.026 and +0.034, and exact inference there is weak.

5. Robustness, measurement, and interpretation

Table 3. Measurement agreement between unrestricted raw-chain and standardized surfaces

Sample	Metric	Pairs	Pearson	Spearman	Mean raw − standardized
Full	Abnormal IV30	52	0.769	0.865	−0.004
Full	Event variance	52	0.533	0.574	0.000
Identification clean	Abnormal IV30	28	0.537	0.808	−0.030
Identification clean	Event variance	28	0.197	0.505	−0.006
Strict clean	Abnormal IV30	26	0.582	0.846	−0.045
Strict clean	Event variance	26	0.257	0.477	−0.010

As shown in Table 3 .Full-sample AIV agrees well across surfaces: Pearson correlation 0.769, Spearman 0.865, over 52 paired events. Identification-clean correlations fall to 0.537 and 0.808. Event-variance agreement is weaker still—Pearson 0.533 in the full sample and 0.197 after the audit. This is why no clean conclusion in this paper rests on one surface. Standardization suppresses sparse-chain noise but embeds vendor interpolation; raw reconstruction is transparent but exposed to

DOI: 10.54254/2754-1169/2026.39251

the available strikes, expiries, and bid-ask spreads [6]. Their disagreement carries real information about measurement uncertainty, and the conclusion has to carry it too.

The unrestricted raw reconstruction shows why a clear delta rule is necessary. Among 26,794 expiry-day data points, only 19,213 meet the ±0.10 band around the 0.50 delta target for both calls and puts, while the ±0.20 band raises this to 22,843. The strict ±0.10 rule yields 2,666 valid IV30 surfaces, and the ±0.20 rule yields 3,263. Our confirmatory raw chain analysis uses ±0.10 as the standard and ±0.20 as a sensitivity check, with unrestricted matching serving only as a diagnostic. Without this rule, deep ITM or OTM contracts are misclassified as ATM simply because they are the closest quote. Downside skewness shows no significant rise across surfaces and samples, so the story that investors buy extra downside protection before launches are not supported. The full event set contains only eight adverse events, and both option and information screens remove even more. Multi-covariate classifiers mostly fit issuer identity and rare event noise.

6. Conclusion

Commercial-space missions deliver large realized stock-price news; the pre-mission option premium, by contrast, is conditional on the design that defines it. Under a disclosure screen aligned to the (-10,-1) AIV window, standardized mission-weighted AIV is +0.040 and the strict raw-chain estimate is +0.028. Issuer-equal estimates are larger, yet none of the four confirmatory tests survives exact issuer-sign inference and Holm adjustment; with four issuers, indeed, the exact test could not have reached 5% significance in any case. Missions announced at least ten trading days ahead, and schedule-verified missions with no same-issuer event within 20 trading days, show stronger positive estimates. Those patterns still fall short of a universal premium: they weaken under strict raw measurement, and again once the broader disclosure screen is combined with non-overlap.

What survives every screen is the realized side. After excluding independent material news on reaction days 0 and +1, adverse missions average 13.2% in absolute two-day abnormal return against 5.9% for successful ones. That contrast—large realized consequences, specification-sensitive ex-ante signals—is the finding. The methodological lesson is that repeated technical events change the estimand. An event study in this setting should align each information screen to its outcome window, measure how long the event date was public, disclose observation reuse, report event- and issuer-weighted effects, and demand agreement across plausible market-measurement pipelines before calling a premium general.

References

[1] Andrew, D., Michael, J., Andreas, K., et al. (2019). Option pricing of earnings announcement risks. *The Review of Financial Studies, 32*(2), 646-687. DOI:10.1093/rfs/hhy060.

[2] Ayala, L. F., Carsten, W., Vitali, B. (2022). Analysis of space launch vehicle failures and post-mission disposal statistics. *Aerotecnica Missili & Spazio, 101*(3), 243-256. DOI:10.1007/S42496-022-00118-5.

[3] Mackinnon, G. J., Webb, D. M. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics, 32*(2), 233-254. DOI:10.1002/jae.2508.

[4] Megginson, L. W., Zhou, Y. X., Gholson, L. R. (2024). The case against a US sovereign wealth fund. *Financial Review, 60*(1), 5-12. DOI:10.1111/FIRE.12422.

[5] Gary, Solon, Haider, S., et al. (2015). What are we weighting for? *Journal of Human Resources, 50*(2), 301-316.

[6] Yan, S. (2010). Jump risk, stock returns, and slope of implied volatility smile. *Journal of Financial Economics, 99*(1), 216-233. DOI:10.1016/j.jfineco.2010.08.011.