Cross-firm implied volatility spillovers at earnings announcements

[a]

[a], ,

ARTICLE INFO

Keywords:
Implied volatility
Earnings announcements
Information spillovers
Realized volatility
Options market efficiency

ABSTRACT

This study finds that when a firm announces earnings, the implied volatility of industry-peer firms moves with the announcer's implied volatility. Using U.S. equity option data (2010–2024) and GICS peer groupings, spillovers are statistically significant, even when using date-fixed effects to eliminate common variation. Controlling for industry-date averages, spillover operates at the industry level rather than through firm-specific signals. Two tests suggest that spillovers carry information about future uncertainty. First, the spillover persists into the peer's own subsequent earnings announcement, particularly if the peer reports shortly after. Second, the spillover does not dissipate across the earnings season.

1. Introduction

Earnings announcements represent the current financial situation of the reporting firm, but also highlight the market conditions facing its industry as a whole. Prior literature has quantitatively studied how cross-firm information transmits through equity prices. However, there is less evidence on what information propagates through the options market. As implied volatility reflects the market's expectation of future risk, the forward-looking metrics can reveal how market participants update their uncertainty assessments across firms during earnings announcements. Spillovers in returns and volatility are well documented at the aggregate market level (Diebold and Yilmaz, 2009); whether a comparable process operates across firms within the options market at scheduled information events is less understood.

This paper studies cross-firm implied volatility spillovers at earnings announcements. When one firm in an industry reports its quarterly results, the question is whether the change in its implied volatility predicts contemporaneous changes in the implied volatilities of peer firms, and whether these spillovers are consistent with information about future uncertainty or reflect transient noise.

Utilizing a sample of U.S. equity options from 2010 to 2024, this paper measures the 30-day at-the-money implied volatility changes around earnings announcements. Industry peers are identified using GICS classifications. The primary specification uses date fixed effects with two-way clustered standard errors, absorbing common daily variation in aggregate volatility. The spillover effect is robust across less demanding fixed effects, and the coefficients are larger as less time variation is absorbed. The results are statistically significant throughout, including under date fixed effects. Different-sector and random non-peer firms show no significant spillover once date fixed effects are included.

The spillover effect operates at the industry level. When controlling for the sector-date average of announcer implied volatility changes, the individual announcer's contribution drops to zero while the industry-level signal remains significant. Thus, individual earnings announcements serve as conduits for industry-wide uncertainty, rather than only transmitting firm-specific signals.

There are two ways to show that the spillovers contain information and do not simply represent noise. First, there is a persistence test which looks at how long the spillover will last once the peer firm has reported on their own. If the peer firm announces very close together (i.e., soon), then the spillover persists; that is, the implied volatility adjustment received from the first announcer remains when the peer firm's earnings provide clarity. However, if the peer firm announces much farther into the future than the original announcer, eventually this effect decays as subsequent events render the first announcer's signal stale. Second, when examining the order in which firms within an industry release their quarterly results, the spillover does not decrease as more firms within the industry have released their results—the amount of the spillover generated by a late announcer is equivalent to what was produced by an early announcer. While the cumulative number of prior announcers reduces the remaining peer uncertainty directly, the sensitivity to each individual announcement remains roughly constant.

ORCID(s):

This research further analyzes the connection between spillover and realized volatility. Peer implied volatility changes correspond closely to subsequent realized volatility in a direct assessment; however, the cross-firm transmission mechanism is associated with a widening of the variance risk premium: spillovers generate implied volatility without corresponding increases in realized volatility, which is consistent with the long-recognized gap between implied and realized volatility extending into the cross-firm space.

The equity returns show no predictability across different firms when looking beyond a few days. Tests for cross-sectional variation (i.e., firm size, options depth, or earnings surprise magnitude) also fail to demonstrate heterogeneity. However, there are some points that seem to indicate this phenomenon could be more prevalent in smaller and less-liquid firms. Formal tests do not reject symmetry of the spillover with respect to the sign of the shock ($p = 0.105$), though point estimates suggest the effect may be somewhat stronger for volatility increases.

This paper contributes to the literature on cross-firm information transmission and options market efficiency. Compared with Hann, Kim and Zheng (2019) (who show that peers are predicted by the first announcer's implied volatility), this analysis demonstrates that the spillover persists across the full sequence of announcers within an earnings season, as well as tests whether these spillovers continue into the subsequent announcement from the peers. Relative to the earnings contagion literature in equity markets (Foster, 1981; Hou, 2007), the results show that the options market and equity market process cross-firm information differently: equity returns show no predictability with time fixed effects, while implied volatility exhibits a significant and persistent spillover.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data and sample construction. Section 4 presents the empirical strategy. Section 5 reports the results, and Section 6 concludes.

2. Related literature

This paper relates to four strands of the finance and accounting literatures: earnings contagion in equity markets, the information content of implied volatility, options market efficiency around information events, and the specific study of cross-firm implied volatility dynamics.

The earnings contagion literature finds that a single firm's earnings announcement has information relevant to other firms. Foster (1981) presents some of the earliest evidence that intra-industry information transfer is connected to the earnings announcements of individual firms, showing that the stock price of a non-announcing firm will react to the earnings of an announcing peer within its industry. Han and Wild (1990) confirm the Foster result with further evidence of intra-industry return transfer, and Kim, Lacina and Park (2008) show that intra-industry transfers from management forecasts can be positive for nonrival firms and negative for rivals, reflecting whether industry commonalities or competitive-shift effects dominate. Ramnath (2002) demonstrates that the first announcer's earnings contain information about the earnings of subsequent announcers in the same industry, but that both analysts and investors underreact to this cross-firm signal, producing predictable returns for later announcers. Hou (2007) provides evidence that slow dissemination of industry-level information leads to lead-lag effects in the stock returns of firms. Beyond industry classifications, Moskowitz and Grinblatt (1999) document industry-level momentum, Cohen and Frazzini (2008) show that returns are predictable along customer-supplier links, and Gleason, Jenkins and Johnson (2008) find analogous contagion in response to accounting restatements. Thomas and Zhang (2008) study whether intra-industry earnings return spillovers are efficient and find evidence of overreaction that subsequently reverses. These studies focus exclusively on equity returns and analyst forecasts.

Implied volatility contains information about future realized risk. Patell and Wolfson (1979) provide the earliest evidence that call option prices anticipate scheduled earnings announcements, and Patell and Wolfson (1981) extend this result. Christensen and Prabhala (1998) demonstrate that implied volatility outperforms historical volatility as a forecast of future realized volatility, subsuming its information content. Dubinsky, Johannes, Kaeck and Seeger (2019) introduce reduced-form option-pricing models that separate earnings-announcement uncertainty from day-to-day volatility and show that the anticipated earnings-event component is large and predicts future return volatility. Rogers, Skinner and Van Buskirk (2009) document that implied volatility responds to management forecasts, and Barth and So (2014) show that earnings-announcement volatility risk is priced in option returns. These studies focus on the own-firm relationship between implied and realized volatility.

Option prices respond to information events in a way that points to the relative efficiency of the options market compared with the underlying equity market. An, Ang, Bali and Cakici (2014) study the joint cross section of stocks and options and show that option-implied information helps explain the cross section of expected returns. Chakravarty, Gulen and Mayhew (2004) provide evidence that a portion of price discovery happens in the options market. Relatedly, Pan and Poteshman (2006) show that option volume predicts stock returns and Cremers and Weinbaum (2010) show that deviations from put-call parity contain information about future equity returns. That price discovery across markets happens first in the options market, on average, in line with these papers.

Other literature documents comovement in option-implied quantities across firms. Driessen, Maenhout and Vilkov (2009) show that correlation risk across single-stock options is priced. Herskovic, Kelly, Lustig and Van Nieuwerburgh (2016) show that there is a common factor in firm-level idiosyncratic volatility, with important quantitative asset-pricing implications. These all illustrate that option-implied uncertainty comoves across firms in systematic ways, but none isolates the channel of earnings announcements.

Hann et al. (2019) demonstrate intra-industry information flow using changes in implied volatility around quarterly earnings announcements. They show that the change in the first announcer's implied volatility predicts changes in implied volatility among other peers in the same industry and quarter, and that this effect is larger when the first announcer is more informative. This paper differs in several aspects. First, this paper studies how cumulative learning occurs as all firms announce within an industry-quarter, not just the first announcing firm. Second, this paper tests whether the spillover persists even after the peer firm has subsequently announced its own quarterly earnings, directly testing the information-versus-noise distinction. Additionally, to address concerns about noise versus information in the spillover, a series of specifications is used to absorb as much common variation as possible, and the spillover remains robust even under the most aggressive time fixed effects specifications. Finally, the sample extends far beyond the S&P 500, including approximately 1,600 firms across 15 years. Hao (2025) argues that implied volatility spillovers are simply indicative of overreactions that eventually reverse. The empirical evidence presented here offers a contrary view—the spillover persists at the peer's own subsequent announcement for peers that report soon after, directly countering the prediction of reversal. A direct calibration test shows that peer implied volatility changes are closely calibrated to subsequent realized volatility ($\gamma \approx 1$) in aggregate. However, when decomposing the peer's ΔIV into two components (the cross-firm spillover component, and the residual), the spillover channel increases implied volatility without generating a proportional increase in realized volatility—thereby widening the variance risk premium through the cross-firm channel.

The gap addressed here is the absence of direct tests of whether cross-firm implied volatility spillovers at earnings announcements carry information about future uncertainty. Prior work has documented the existence of the spillover and debated whether it represents information or overreaction, but has not tested persistence at the peer's own announcement, cumulative learning across the earnings season, or robustness across a ladder of fixed effects specifications. This paper provides these tests and finds evidence consistent with the spillover reflecting industry-level information that persists when the peer firm subsequently reports.

3. Data and sample construction

3.1. Data Sources and Sample Construction

The sample contains U.S. listed firms with traded equity options over the period 2010—2024. Daily option chain data were downloaded from FirstRate Data. This data includes closing prices, strikes, expirations, and implied volatilities for listed equity option trading on U.S. exchanges. Interpolation is used to compute thirty (30)-day at-the-money (ATM) Black–Scholes implied volatility (Black and Scholes, 1973), which allows for a constant maturity measure. Stock prices were also downloaded on a daily basis from FirstRate Data. Earnings announcements were found through publically available earnings calendars.

Industry peers are defined by 6-digit GICS industry classification. GICS is preferred over SIC- or NAICS-based schemes because Bhojraj, Lee and Oler (2003) show that GICS outperforms these alternatives in explaining return comovement and peer-firm characteristics.[1] All same-industry firms that did not announce earnings on the same day as the announcing firm are designated as peers. Announcing firms and firms without sufficient option data are excluded from the analysis. A series of robustness checks use 2-digit GICS sector designations; a series of placebo tests use firms located within a different GICS sector and randomly selected non-peer firms.

Percentage changes in implied volatility are computed for both the announcing firm and its peer group, based on a pre-announcement baseline (the average level over trading days $[-40, -11]$ relative to the announcement) and a

[1] An alternative is the Text-based Network Industry Classification of Hoberg and Phillips (2016), which defines peers from 10-K product-description similarity. GICS is retained here for broader coverage of firms with traded options.

Table 1
Sample Construction

Step	Count	Unit
Firms with daily price data (2010–2024)	2,828	Firms
Firms with valid 30-day ATM IV coverage	2,792	Firms
Unique earnings events (after deduplication)	6,238	Events
Announcing firms with ≥1 industry peer	942	Firms
Peer-event observations	66,982	Obs.
of which: with return and RV data	66,896	Obs.

Notes: Sample construction. The initial universe consists of all U.S. listed firms with equity options data from FirstRate Data. Firms are matched to earnings announcement calendars, and industry peers are identified using 6-digit GICS industry classifications. The sample reduction from 1,601 announcing firms to 942 reflects that 41% of announcing firms have no industry peers with sufficient option data to compute implied volatility changes in the event window.

post-announcement level (the average over trading days $[0, +3]$, i.e., the announcement day through three trading days after). Both variables are winsorized at the 1st and 99th percentiles.

Sample construction data is listed in Table 1. After data quality filtering, the original 2,828 firms narrow to 1,372 unique firms, with 66,982 peer-event observations that span 1,143 announcement dates. Of the 1,365 unique peer firms, 430 are only used as peers and not as individual announcers. The remaining 935 firms serve in both roles at different times.

The sample data spans 2010 through 2024, but coverage of options data varies per year due to increased data availability and growth of the U.S. options market. There were fewer than a thousand total observations for the first six years (2010–2016), and the major samples fall within the years of 2017-2024. The key findings were verified not to be driven by single-year outliers using year fixed effects and confirming coefficient stability across subperiods.

3.2. Variable Definitions

The two most important variables are the announcing firm's percentage change in implied volatility (ΔIV^{ann}) and the peer firm's percentage change in implied volatility (ΔIV^{peer}), both defined as:

$$\Delta IV = \frac{IV_{post} - IV_{pre}}{IV_{pre}} \tag{1}$$

where IV_{pre} is the average 30-day ATM implied volatility in the pre-announcement window ($[-40, -11]$ trading days) and IV_{post} is the average implied volatility during the post-announcement window ($[0, +3]$ trading days).

To test if peer implied volatility changes forecast future realized volatility, the percent change in post-earnings-date realized volatility is computed against a baseline. The baseline is the standard deviation of daily return data for a time interval of $[-40, -11]$, 30 trading days prior to an earnings announcement. This interval is chosen to avoid overlap with the implied volatility measure, occurring during days $[0, +3]$. Realized volatility is then computed on three different post-announcement intervals: $[+4, +10]$, $[+4, +20]$, and $[+4, +30]$, again so that they do not overlap with the implied volatility measurement period. The four-day gap eliminates any mechanical relationship between the two measures.

To study persistence, the change in implied volatility that a peer firm experiences at its own subsequent earnings announcement is related to the spillover received from the prior announcer. The gap between the two announcements ranges from one to ninety days.

Cumulative abnormal returns (CARs) are defined as buy-and-hold returns over windows of 1, 3, 5, 10, and 20 trading days following the announcing firm's earnings date, adjusted for the market return over the same horizon.

3.3. Summary Statistics

Table 2 highlights the descriptive statistics. Consistent with volatility theory after earnings, the median announcing firm sees a volatility decrease of 9.1% relative to its pre-announcement baseline. Peer firms' implied volatility changes are smaller and centered near zero (mean 0.3%, median −2.8%).

Table 2
Summary Statistics

	N	Mean	Median	Std. Dev.	P5	P95
Panel A: Main Variables (Winsorized)						
Announcer ΔIV (%)	66,982	−8.01	−9.13	15.68	−30.05	17.47
Peer ΔIV (%)	66,982	0.25	−2.82	22.61	−30.46	40.25
Peer baseline IV	66,982	0.440	0.377	0.230		
Post-event days	66,982	3.86	4.00	0.35		
|EPS surprise| (%)	66,982	—	15.72	—		

Notes: Summary statistics on the main variables in the dataset. Both announcer ΔIV and peer ΔIV measure the percentage change in 30-day ATM implied volatility from a pre-announcement baseline (average of trading days $[-40, -11]$) to the post-announcement window (trading days $[0, +3]$). Data are winsorized at both the 1st and 99th percentiles. Peer baseline IV is the level of 30-day ATM implied volatility on the day before the announcement. Post-event days is the number of trading days in the post-announcement measurement window. |EPS surprise| is the absolute percentage deviation of reported earnings per share from the consensus estimate; the mean and standard deviation are omitted because the distribution is highly right-skewed with outliers.

4. Empirical strategy

4.1. Core Specification

The overall model for assessing cross-firm implied volatility spillovers from earnings announcements is described as follows:

$$\Delta IV_{i,t}^{peer} = \alpha_t + \beta\, \Delta IV_{j,t}^{ann} + \varepsilon_{i,t} \tag{2}$$

In the above equation, $\Delta IV_{i,t}^{peer}$ is the percentage change in 30-day ATM implied volatility for firm i's peer firms around firm j's earnings announcement on trading day t. In addition, $\Delta IV_{j,t}^{ann}$ is the equivalent percentage change in implied volatility for firm j when it announces earnings. Finally, α_t is a fixed effect representing each date. The coefficient β measures the degree to which implied volatility increases or decreases for peer firms around an individual firm's earnings announcement, while the date fixed effect absorbs all common variation in aggregate volatility affecting firms on the same trading day.

Following the methods from Cameron, Gelbach and Miller (2011) and Petersen (2009), two-way robust standard error clustering is applied to the data. This addresses potential issues related to the fact that some peer firms may appear multiple times in the sample because they are observed across multiple earnings events during the time period. Clustering by firm corrects for serial correlation at the firm level. Additionally, clustering by date corrects for cross-sectionally correlated residual values that occur on the same trading day. Because the model includes fixed effects, a degrees-of-freedom adjustment is also applied. The adjustment adds about one percent to the standard errors but has no impact on inference.

β measures the sensitivity of implied volatility changes for peer firms to the announced firm's implied volatility change, given the specific trading day.

4.2. Choice of Fixed Effects

One of the main decisions regarding methodology is the time fixed effect level of detail. The most-preferred choice utilizes date-level fixed effects; these absorb all common daily variation, including changes in the VIX, macroeconomic news, and overall market-wide risk perception. The date-level fixed effect also absorbs the cross-sectionally averaged dependent variable for each date across all of its respective peers. The identification of β is based upon within-day variations in the announcer's ΔIV across peer groups that have been exposed to different announcers. On days when there were multiple announcers representing different industries, this variation was significant. When there was one announcer per day, however, identification relied upon the comparison between the announcer's signal and the date mean, where the date mean utilized data from peer groups in other industries. Therefore, the single-announcer subsample analysis shown in Table 4 represents the most direct examination of the spillover effect.

A "specification ladder" is presented that shows how the coefficient responds to decreasing levels of restrictions imposed by various time-based fixed effects: (i) date fixed effects, (ii) week fixed effects (both with and without inclusion of VIX), (iii) month fixed effects (both with and without inclusion of VIX), (iv) VIX included as a control variable, and (v) no controls. Using month fixed effects yields larger coefficients than date fixed effects; with only 174 year-month groups compared to 1,143 date groups, month fixed effects absorb less common variation, preserving cross-date variance that inflates the estimated spillover. By providing this specification ladder, the reader can evaluate whether a portion of the raw co-movements between peers reflect common time-series patterns or whether they represent within-day, peer-specific spillovers.

4.3. Identification

The variation in Equation (2) which identifies β comes from within-date differences in announcer implied volatility changes—different announcers on the same day have different ΔIV, and that cross-sectional variation identifies the spillover. Date fixed effects remove all common daily shocks, including changes in the VIX and macroeconomic news. Therefore, the confounding factor becomes specific to the industry rather than aggregate.

There can be industry-level shocks specific to a particular day which may cause both the announcer and peer to have a similar change in their implied volatility. An alternative approach to consider is to add industry-by-date fixed effects to account for these shocks. However, there is no valid way to do this in this study. For example, if three companies from the same industry announce on the same day, every peer of those announcing companies will receive exposure to all of the announcements at the same time. Therefore, when estimating the within-industry-date model, it attempts to isolate the effect of a single announcer's signal by comparing peers that are all exposed to the exact same set of announcements; therefore, it creates a mechanically biased estimate.

The industry-level identification challenge is addressed through two complementary approaches. First, the sample is restricted to days where only one company in an industry has announced its earnings. On these "single-announcer" days the announcer's signal is identical to the industry signal by design, thereby eliminating the multi-exposure problem. Second, the sector-date average of announcer implied volatility changes is included as an additional control variable. This test determines whether individual announcers provide new information relative to what is contained in the industry signal.

4.4. Persistence at Own Announcement

A study of the informational content will investigate if the spillover continues to persist after the peer firm has released its own quarterly earnings. The reason for this is that, as indicated in the prior discussion, if the spillover reflected an increase in general uncertainty about the specific industry in which both firms operate, the implied volatility adjustment should be consistent with the peer's own-announcement volatility change. Conversely, if the spillover was simply caused by a transitory event or noise, we would expect a reversal. This specification is estimated:

$$\Delta IV_i^{peer,own} = \alpha_t + \beta\,\Delta IV_i^{peer,spillover} + \varepsilon_i \tag{3}$$

where $\Delta IV_i^{peer,own}$ is the peer firm's change in implied volatility at the time of their subsequent reporting of quarterly earnings, and $\Delta IV_i^{peer,spillover}$ is the spillover effect on the peer firm's implied volatility that resulted from the earlier announcer. If $\beta > 0$, then the results indicate persistence (the spillover contained information) since it continued to have an impact. On the other hand, if $\beta < 0$, then the results indicate a reversal.

4.5. Cumulative Learning

To see if the spillover has changed because there are additional firms in the same industry reporting their earnings, the announcer's implied volatility change has been interacted with the number of previous announcers from the same industry that reported their earnings during the same quarter. As such, n^{prior} is included as both a main effect (β_0) and an interaction (β_2):

$$\Delta IV_{i,t}^{peer} = \alpha_t + \beta_0\, n_{j,t}^{prior} + \beta_1\,\Delta IV_{j,t}^{ann} + \beta_2\,(\Delta IV_{j,t}^{ann} \times n_{j,t}^{prior}) + \varepsilon_{i,t} \tag{4}$$

$n_{j,t}^{prior}$ denotes the number of firms in the same industry that have already announced earnings in the same quarter before firm j. β_0 represents the direct influence of previous announcements on peer implied volatility (i.e., through resolving uncertainty). β_2 addresses whether or not the sensitivity of peer implied volatility to the current announcer's signal is different based on how many prior signals were received by peers. Because n^{prior} varies by industry for each given date, it cannot be captured by the date-specific fixed effect and must therefore be entered separately.

4.6. Realized Volatility and the Variance Risk Premium

To determine if the peer firms' implied volatility transmission is consistent with their future realized risk, the peers' changes in realized volatility are regressed against their changes in implied volatility as follows:

$$\Delta RV^{peer}_{i,[+4,+T]} = \alpha + \gamma\, \Delta IV^{peer}_{i,t} + \varepsilon_{i,t} \tag{5}$$

where $\Delta RV^{peer}_{i,[+4,+T]}$ represents the percentage change in peer firm i's realized volatility for the window $[+4, +T]$ relative to the pre-event baseline $[-40, -11]$, and $T \in \{10, 20, 30\}$ trading days. The four-day gap between the end of the implied volatility measurement window (day +3) and the start of the realized volatility window (day +4) eliminates any mechanical association between the two measures. Under rational pricing, $\gamma = 1$: on average, the changes in implied volatility should be equivalent to changes in subsequent realized volatility. If $\gamma < 1$, it implies that implied volatility has an upward bias compared to realized volatility. Similarly, if $\gamma > 1$, then implied volatility appears to have a downward bias relative to subsequent realized volatility. This test does not require time fixed effects, since both the independent and dependent variables are measured for the same firm, and the question is whether the magnitude of the implied volatility adjustment is justified by subsequent realized risk.

5. Results

5.1. Does IV Spill Over Across Industry Peers?

Table 3 provides an overview of the empirical results on the spillover of implied volatility across a range of alternative specifications. Specification 7 represents the preferred specification. In this case, the estimated coefficient is 0.064 with a t-statistic of 3.68 ($p < 0.001$). This specification does so by absorbing all common day-to-day variations in aggregate volatility, such as changes in the VIX or macroeconomic news, and identifying spillovers based only on within-day variations in each announcer.

The remaining columns show how the coefficient responds to less demanding controls. Specification 5 includes week fixed effects. These produce estimates of 0.072 with a t-statistic of 4.43. Specification 3 includes month fixed effects. These produce estimates of 0.165 with a t-statistic of 7.57. Finally, Specification 1 includes no fixed effects. Estimates of 0.431 are generated with a t-statistic of 12.84. As more time variation is absorbed in subsequent specifications, there is a monotonic decline in the size of the estimates. Thus, it appears that much of the raw co-movement is due to peer exposure to overall market shocks, rather than peer-specific spillovers.

Finally, Table 3 also presents similar ladders of specifications for a larger set of peer definitions. Specifically, Table 3 contains a broader set of peer definitions at the two-digit GICS industry level. Using these larger sets of peers generates estimates of spillovers across all specifications; even when date fixed effects are included. For example, with this larger peer definition set (393,452 observations), the estimated coefficient is 0.023 with a t-statistic of 2.66 ($p = 0.008$).

The results in Table 4 further describe the evidence. In Panel A the model is estimated with month fixed effects which allow for greater statistical power in the subsamples. Using only single-announcer days – where the signaling from each individual announcer will be equal to the industry-wide signaling – the estimate is $\beta = 0.105$ ($t = 3.02$). For the multi-announcer subsample the β is slightly smaller yet still statistically significant at 0.050 ($t = 2.62$).

In the industry mean control case, as shown in Panel A, Column 4, this supports an industry level effect; when the date/sector averages for announcer implied volatility are included in the model, the coefficient on the announcer drops to 0.005 ($t = 0.29$) while the coefficient on the industry average is 0.148 ($t = 4.58$). Thus the individual announcer's signal is being absorbed by the industry wide signal.

Finally, Panel B shows the placebo tests based upon the use of alternative peer definitions. As stated above the main test is to compare the spillover from industry peers versus other types of peers. The primary test utilizes date fixed effects because it isolates peer specific variation from common time series variation. Therefore, under date fixed effects, there is a significant spillover from industry peers ($\beta = 0.064$, $t = 3.68$); however, there is no spillover from same sector but different industries peers ($\beta = 0.011$, $t = 1.18$), or from different sectors peers ($\beta = -0.013$, $t = -0.66$), or from random non-peer firms ($\beta = -0.029$, $t = -1.58$).

Thus, the spillover observed is due to industry-level signals being transmitted across firms within the same industry once common day-to-day variation has been removed. When month fixed effects are used instead of date fixed effects,

Table 3
Implied Volatility Spillover: Specification Ladder

	(1) No FE	(2) VIX	(3) Month	(4) Month+VIX	(5) Week	(6) Week+VIX	(7) Date
Panel A: Industry Peers ($N = 66{,}982$)							
β	0.431***	0.271***	0.165***	0.115***	0.072***	0.072***	0.064***
	(0.034)	(0.028)	(0.022)	(0.019)	(0.016)	(0.016)	(0.018)
FE groups	–	–	174	174	489	489	1,143
Panel B: Sector Peers ($N = 393{,}452$)							
β	0.423***	0.262***	0.102***	0.060***	0.029***	0.028***	0.023***
	(0.034)	(0.028)	(0.013)	(0.010)	(0.008)	(0.008)	(0.009)
FE groups	–	–	174	174	495	495	1,164

Notes: This table reports regressions of peer firm ΔIV on announcer ΔIV across a ladder of increasingly aggressive time controls. Panel A uses 6-digit GICS industry peers; Panel B uses 2-digit GICS sector peers (which additionally include same-sector, different-industry firms). VIX control is the percentage change in VIX around the announcement. Standard errors (in parentheses) are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

then even the same sector, but different industries, and random non-peer firms experience significant coefficients showing that in order to remove common time-series variation and identify peer-specific spillovers, date fixed effects are necessary rather than month fixed effects.

5.2. Is the Spillover Informative? Persistence at Own Announcement

Table 5 tests whether the implied volatility spillover will persist when the peer firm subsequently announces its earnings. If the spillover represents an increase in the uncertainty surrounding the entire industry, then the implied volatility adjustment the peer received should be consistent with the peer firm's own announced changes. If instead the spillover represented nothing but noise, one would expect this relationship to reverse as well.

Column 1 contains the unconditional estimates of β (the implied volatility change induced by the spillover) without fixed effects: the spillover-induced implied volatility change is positively related to the peer firm's subsequent announcement change in implied volatility ($\beta = 0.203$, $t = 6.69$). Including date fixed effects in column 2 does not affect the significance of the coefficient $\beta = 0.112$ ($t = 4.08$). Thus, information incorporated into peer implied volatility at the time of the industry peer's announcement continues to influence the peer firm's own implied volatility at subsequent reporting periods.

In comparison to all other firms, peers announcing shortly after the original announcer have much stronger persistence. Column 3 includes only those firms that announce within thirteen trading days of the original announcer (approximately equal to the sample median) and finds $\beta = 0.504$ ($t = 11.42$) with date fixed effects. In contrast, column 4 limits the analysis to those firms announcing more than thirteen days after the original announcer and finds $\beta = -0.064$ ($t = -2.00$), which is marginally significant and negative. The concentration of the information contained in the spillover in the near-term is therefore very clear: for peers announcing within roughly two weeks of each other, the initial implied volatility adjustment experienced by the original announcer is highly correlated with the peer firm's own announced implied volatility change. For peers announcing later than this, the relationship becomes negative, indicating that intervening events render earlier signals obsolete.

Column 5 conducts a horse-race between the peer firm's processing of spillovers received from other firms and the original implied volatility change experienced by the announcer, with date fixed effects included in estimation. Both variables remain statistically significant; the peer firm's processing of spillovers received from other firms yields $\beta = 0.112$ ($t = 4.14$) while the announcer's initial implied volatility change yields $\beta = 0.106$ ($t = 5.05$). These findings suggest that both the peer firm's processing of spillovers received from other firms and the announcer's original implied volatility change contain predictive information regarding the peer firm's own announced implied volatility.

Table 4
Implied Volatility Spillover: Identification and Placebo Tests

	(1)	(2)	(3)	(4)
Panel A: Identification (Month FE)				
	Full Sample	Single Ann.	Multi Ann.	Ind. Mean
Announcer ΔIV	0.165***	0.105***	0.050***	0.005
	(0.022)	(0.035)	(0.019)	(0.019)
Industry mean ΔIV				0.148***
				(0.032)
N	66,982	20,907	46,075	66,982
Panel B: Placebo Tests (Month FE / Date FE)				
	β (Month)	β (Date)	N	
Industry peers	0.165***	0.064***	66,982	
Same-sector, diff. industry	0.089***	0.011	326,470	
Different sector	0.085***	−0.013	19,991	
Random non-peers	0.068***	−0.029	19,734	

Notes: Panel A reports regressions of peer ΔIV on announcer ΔIV with month fixed effects. Column (1) uses the full sample. Column (2) restricts to dates with a single announcer per industry. Column (3) restricts to dates with multiple announcers. Column (4) adds the sector-date average of announcer ΔIV as a control, estimated on the full sample using GICS sector classifications for all firms. Panel B reports placebo tests using alternative peer definitions, showing coefficients with both month and date fixed effects. The "same-sector, diff. industry" group consists of firms in the same 2-digit GICS sector but a different 6-digit GICS industry, excluding the treated industry peers. Standard errors (in parentheses) are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

5.3. Cumulative Learning Across the Earnings Season

Table 6 shows whether or not the spillover effect will grow as more firms within the same industry report their earnings. Whether the spillover changes with the number of prior announcers is tested using an interaction, as discussed in Section 4.5. The specification used includes the number of firms reporting before the firm being analyzed (n^{prior}) as both a main effect and an interaction with the announcer's implied volatility change.

A negative and significant relationship exists between n^{prior} and the peer's implied volatility adjustment, indicating that with every additional firm reporting their earnings, the peer's implied volatility adjustments decrease by approximately 0.2% ($\beta = -0.002$, $t = -3.61$). This indicates that uncertainty resolves itself throughout the earnings season as more firms within the industry report. However, the interaction term, which determines if the sensitivity of peer implied volatility to the announcer's signal changes based on the number of previous announcements, is positive, however it was statistically insignificant in all specifications (baseline specification $p = 0.135$, column 2, $t = 1.02$, column 4, $t = 1.60$). A statistically marginally larger interaction existed when controlling for the number of firms within an industry (column 3, $t = 1.74$).

Overall these results indicate that the spillover effect does not increase in magnitude as more firms within an industry report their earnings. Instead, the sensitivity of peer implied volatility to the first reporter is statistically equal to subsequent reporters. Evidence of this can be found in that the coefficient of the first reporter ($\beta = 0.097$, $t = 2.60$) is statistically equivalent to the coefficients of subsequent reporters (the fourth-and-later reporters' coefficient is $\beta = 0.068$, $t = 1.75$; difference $p = 0.62$).

5.4. Implied Volatility, Realized Volatility, and the Variance Risk Premium

The relationship between the implied volatility spillover and subsequent realized volatility has been examined within Table 7. In the area of options markets there exists an empirically-identified regularity; namely that the variance risk premium (implied volatility > realized volatility) is well documented at the index level (Bakshi and Kapadia, 2003;

Table 5
Persistence at Own Announcement

	(1) No FE	(2) Date FE	(3) Short Gap	(4) Long Gap	(5) Horse Race
Peer ΔIV (spillover)	0.203***	0.112***	0.504***	−0.064**	0.112***
	(0.030)	(0.027)	(0.044)	(0.032)	(0.027)
Announcer ΔIV					0.106***
					(0.022)
FE	None	Date	Date	Date	Date
Sample	All	All	≤13d	>13d	All
N	16,205	16,205	8,140	8,065	16,205

Notes: This table regresses the peer firm's ΔIV at its own subsequent earnings announcement on the ΔIV it received as a spillover from the earlier announcer. Column (3) restricts to peers announcing within 13 trading days of the announcer (approximately the sample median gap); column (4) restricts to peers announcing more than 13 days later. Column (5) includes both the peer spillover and the announcer's original ΔIV as regressors, with date fixed effects. Standard errors (in parentheses) are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Table 6
Cumulative Learning Across the Earnings Season

	(1) Baseline	(2) Density	(3) Ind. Size	(4) Kitchen Sink	(5) Capped	(6) Excl. 1st
n^{prior}	−0.002***	−0.002***	−0.002***	−0.002***	−0.002***	−0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$\Delta IV^{ann} \times n^{prior}$	0.009	0.009	0.015*	0.014	0.002	0.015*
	(0.008)	(0.008)	(0.009)	(0.009)	(0.011)	(0.009)
t-statistic (interaction)	1.15	1.02	1.74	1.60	0.19	1.67
FE type	Date	Date	Date	Date	Date	Date
Controls	No	Yes	Yes	Yes	Yes	No
Capped at 5	No	No	No	No	Yes	No
Excl. 1st ann.	No	No	No	No	No	Yes

Notes: This table reports the interaction between announcer ΔIV and the number of prior announcers (n^{prior}) in the same industry during the same quarter, including n^{prior} as a main effect. All columns use date fixed effects. Column (2) controls for earnings announcement density. Column (3) controls for industry size. Column (4) includes all controls. Column (5) caps n^{prior} at 5. Column (6) excludes first announcers ($n^{prior} = 0$). Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Bollerslev, Tauchen and Zhou, 2009; Drechsler and Yaron, 2011) and extends to individual firms (Carr and Wu, 2009), and therefore this test requests if the cross-firm spillover inherits these properties. That is, it requests if the cross-firm spillover will inherit the characteristics of the variance risk premium.

The peer firm's implied-minus-realized volatility gap is estimated as a function of announcer's implied volatility change using non-overlapping realized volatility windows that start four trading days post-announcement.

Panel A estimates a direct calibration regression model where the dependent variable is peer ΔRV and the independent variable is peer ΔIV. Under rational pricing with no risk premium, the coefficient should be exactly one. The observed values for the coefficients are 0.946, 1.062, and 1.033 at the 10, 20, and 30 day horizons respectively. The hypothesis that $\gamma = 1$ can not be rejected at any of the three time horizons. Peer firms' implied volatility changes appear to be very close to being perfectly calibrated to subsequent realized volatility.

Table 7
Realized Volatility Calibration and Variance Risk Premium

	[+4, +10]	[+4, +20]	[+4, +30]
Panel A: Peer ΔIV → Peer ΔRV (No FE)			
γ	0.946***	1.062***	1.033***
	(0.060)	(0.071)	(0.084)
H_0: $\gamma = 1$	$t = -0.91$	$t = 0.87$	$t = 0.39$
Panel B: IV-RV Gap = α_t + β Announcer ΔIV (Date FE)			
β	0.100	0.097*	0.098**
	(0.060)	(0.051)	(0.046)
t-statistic	1.68	1.93	2.13
Decomposition (Date FE, [+4, +20])			
IV side (β_{IV})	0.065*** ($t = 3.36$)		
RV side (β_{RV})	−0.027 ($t = -0.50$)		
N	66,847		

Notes: Panel A regresses the peer firm's change in realized volatility (ΔRV) on the peer firm's own change in implied volatility (ΔIV). Realized volatility is computed over non-overlapping windows beginning on day +4. Under rational pricing, $\gamma = 1$. Panel B regresses the peer's implied-minus-realized volatility gap on the announcer's ΔIV with date fixed effects. A positive coefficient indicates that the spillover widens the variance risk premium. The decomposition shows whether the effect operates through the IV side (announcer ΔIV → peer ΔIV) or the RV side (announcer ΔIV → peer ΔRV). Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Panel B tests whether the cross-firm channel increases the variance risk premium. The peer firm's IV-RV gap is estimated as a function of announcer's ΔIV with date fixed effects. At the 30 day horizon the coefficient is positive and statistically significant ($\beta = 0.098$, $t = 2.13$, $p = 0.033$), while at shorter horizons it is marginally statistically significant. The decomposition indicates that most of the effect occurs through the IV side: Announcer's signal predicts peer firm's ΔIV ($\beta = 0.065$, $t = 3.36$), however it does not predict peer firm's ΔRV ($\beta = -0.027$, $t = -0.50$).

5.5. Supporting Evidence and Null Results

The idea that the spillover was at an industry level, informationally relevant and approximately efficiently priced is supported by several other conclusions.

Firstly, there is no evidence of predictive ability using equity return data when considering time periods greater than just a few days. Regressing peer cumulative abnormal returns over different time periods (one, three, five, ten and twenty trading days) against the announcer's implied volatility change, without controlling for date fixed effects there is a marginally negatively related coefficient; however once date fixed effects are included, all of the predictive power after five days is eliminated. On the other hand there are some marginal negative coefficients at very short horizons (one and three days), suggesting that peers with larger implied volatility spillovers experience slightly lower short-term returns, consistent with the options market reflecting cross-firm uncertainty before equity prices fully adjust.

Secondly, there is little systematic variation in the size of the spillover across firms. The sample is divided into terciles based upon four variables: peer market capitalization, peer baseline realized volatility, options market depth and the magnitude of the earnings surprise. In most cases the hypothesis that the high-minus-low differences in spillovers are zero cannot be rejected. This indicates that there is no systematic relation between the size of the spillover and the specific characteristics of individual firms such as market cap, volatility, etc.

Lastly, the spillover appears to be symmetric with regard to whether the announcers' implied volatility changed due to good (decrease) or bad (increase) uncertainty news. Two separate coefficients representing implied volatility changes (both positive and negative) are estimated. The t-statistic for the interaction term used to test for asymmetric transmission of bad vs. good news is equal to $t = 1.62$ ($p = 0.105$). Although the point estimates indicate somewhat larger transmission for bad news (i.e., increased implied volatility), the hypothesis of symmetry cannot be rejected at conventional significance levels, suggesting that the spillover is part of a rational update process.

5.6. Robustness

The principal results are stable under numerous alternative specifications discussed fully in the Appendix. Omitting the COVID Period (March 2020–June 2021) produces an estimate for β of 0.061 ($t = 2.90$) when using date-specific fixed effects; thus, these results appear to be unaffected by the high levels of volatility during the COVID period. In addition, the corrections for degrees of freedom increase standard errors by less than 1% and have little or no impact on the ultimate conclusions. Fitting peer firm fixed effects to account for time-invariant differences in implied volatility sensitivity does not significantly alter the estimated values compared to the initial estimates as shown in Table A3.

Figure 1 displays the daily spillover coefficients (calculated without fixed effects to display raw spillovers) from day -10 until day $+10$ relative to the announcement. The coefficients were near zero before the announcement, jumped up on the announcement date ($\beta = 0.445$, $t = 10.28$), and stabilized around 0.54 after the announcement. These results do not indicate a reversal of effect until at least day $+10$. The flatness of the coefficients prior to the announcement provides some assurance that there was not a pre-trend present in the data. The stabilization of the coefficients following the announcement also supports the persistence estimates provided in Table 5.

6. Conclusion

This study analyzes how the options market uses cross-firm uncertainty information when there are earnings announcements. The results show significant spillovers in terms of implied volatility from announcing firms to other firms in the same industry, and these spillovers are observed across a variety of different types of time controls, including date fixed effects, week fixed effects, and month fixed effects. Also, the spillover is based upon the average implied volatility of all announcers within an industry, and thus, the spillover operates at the industry level. There are two separate analyses that demonstrate that the spillover contains useful information. First, the spillover persists at the next earnings announcement made by a firm's peers. Moreover, such persistence is particularly pronounced for those firms that have very short intervals between earnings announcements. Second, the magnitude of the spillover generated by an earnings announcement does not decline as the number of firms which announce after an initial set of firms increases. This suggests that each new earnings announcement has a similar amount of useful information regarding industry-wide uncertainty. Finally, while the previous paragraphs describe how uncertainty is transmitted among firms, the empirical analysis shows that there are no corresponding patterns of return-based predictability between firms.

Thus, these findings contribute to the literature on cross-firm information transmission by showing that the process described above can occur via the options market. In contrast to prior studies, such as Hann et al. (2019) which focused solely on the first firm(s) to make an earnings announcement, this study shows that cross-firm spillovers continue across all firms making subsequent announcements and persist until the subsequent earnings announcement made by a firm's peers. This addresses one of the key questions posed by Hao (2025): whether such spillovers contain useful information versus noise. Additionally, the findings suggest that cross-firm channels may increase the variance risk premium; thereby connecting this study to a broader body of literature on option pricing and risk compensation. As such, the finding that uncertainty is transmitted to peers through changes in implied volatility represents a unique channel for transmitting such uncertainty versus the equity return contagion found in prior research.

The sample consists primarily of large, liquid U.S. firms with actively traded options. Thus, date fixed effects and placebo tests are used to address potential endogeneity issues as opposed to using exogenous variations in announcer implied volatility. Future research may include extending this analysis internationally, examining the term structure of implied volatility spillovers, and/or utilizing intraday data to identify when exactly cross-firm repricings occur.

Data availability

The option chain and stock price data that support the findings of this study are available from FirstRate Data. Restrictions apply to the availability of these data, which were used under licence for this study. Data are available

from the author with the permission of FirstRate Data. GICS classifications are proprietary to MSCI and S&P Global and were obtained from LSEG (Refinitiv Eikon) under licence. Restrictions apply to the availability of these data, which are available from the author with the permission of LSEG. Earnings announcement dates are public information and were collected via LSEG.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used Claude Opus 4.6 (Anthropic) in order to proofread the manuscript, improve readability, and review code. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

An, B.J., Ang, A., Bali, T.G., Cakici, N., 2014. The joint cross section of stocks and options. The Journal of Finance 69, 2279–2337.

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. The Review of Financial Studies 16, 527–566.

Barth, M.E., So, E.C., 2014. Non-diversifiable volatility risk and risk premiums at earnings announcements. The Accounting Review 89, 1579–1607.

Bhojraj, S., Lee, C.M., Oler, D.K., 2003. What's my line? a comparison of industry classification schemes for capital market research. Journal of Accounting Research 41, 745–774.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. The Review of Financial Studies 22, 4463–4492.

Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. Journal of Business & Economic Statistics 29, 238–249.

Carr, P., Wu, L., 2009. Variance risk premiums. The Review of Financial Studies 22, 1311–1341.

Chakravarty, S., Gulen, H., Mayhew, S., 2004. Informed trading in stock and option markets. The Journal of Finance 59, 1235–1257.

Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. Journal of Financial Economics 50, 125–150.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. The Journal of Finance 63, 1977–2011.

Cremers, M., Weinbaum, D., 2010. Deviations from put-call parity and stock return predictability. Journal of Financial and Quantitative Analysis 45, 335–367.

Diebold, F.X., Yilmaz, K., 2009. Measuring financial asset return and volatility spillovers, with application to global equity markets. The Economic Journal 119, 158–171.

Drechsler, I., Yaron, A., 2011. What's vol got to do with it. The Review of Financial Studies 24, 1–45.

Driessen, J., Maenhout, P.J., Vilkov, G., 2009. The price of correlation risk: Evidence from equity options. The Journal of Finance 64, 1377–1406.

Dubinsky, A., Johannes, M., Kaeck, A., Seeger, N.J., 2019. Option pricing of earnings announcement risks. The Review of Financial Studies 32, 646–687.

Foster, G., 1981. Intra-industry information transfers associated with earnings releases. Journal of Accounting and Economics 3, 201–232.

Gleason, C.A., Jenkins, N.T., Johnson, W.B., 2008. The contagion effects of accounting restatements. The Accounting Review 83, 83–110.

Han, J.C., Wild, J.J., 1990. Unexpected earnings and intraindustry information transfers: Further evidence. Journal of Accounting Research , 211–219.

Hann, R.N., Kim, H., Zheng, Y., 2019. Intra-industry information transfers: Evidence from changes in implied volatility around earnings announcements. Review of Accounting Studies 24, 927–971.

Hao, Z., 2025. Overreaction in intra-industry information transfers: Evidence from implied volatility changes around earnings announcements. Available at SSRN .

Herskovic, B., Kelly, B., Lustig, H., Van Nieuwerburgh, S., 2016. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. Journal of Financial Economics 119, 249–283.

Hoberg, G., Phillips, G., 2016. Text-based network industries and endogenous product differentiation. Journal of Political Economy 124, 1423–1465.

Hou, K., 2007. Industry information diffusion and the lead-lag effect in stock returns. The Review of Financial Studies 20, 1113–1138.

Kim, Y., Lacina, M., Park, M.S., 2008. Positive and negative information transfers from management forecasts. Journal of Accounting Research 46, 885–908.

Moskowitz, T.J., Grinblatt, M., 1999. Do industries explain momentum? The Journal of Finance 54, 1249–1290.

Pan, J., Poteshman, A.M., 2006. The information in option volume for future stock prices. The Review of Financial Studies 19, 871–908.

Patell, J.M., Wolfson, M.A., 1979. Anticipated information releases reflected in call option prices. Journal of Accounting and Economics 1, 117–140.

Patell, J.M., Wolfson, M.A., 1981. The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices. Journal of Accounting Research , 434–458.

Petersen, M.A., 2009. Estimating standard errors in finance panel data sets: Comparing approaches. The Review of Financial Studies 22, 435–480.

Ramnath, S., 2002. Investor and analyst reactions to earnings announcements of related firms: An empirical analysis. Journal of Accounting Research 40, 1351–1376.

Rogers, J.L., Skinner, D.J., Van Buskirk, A., 2009. Earnings guidance and market uncertainty. Journal of Accounting and Economics 48, 90–109.

Thomas, J., Zhang, F., 2008. Overreaction to intra-industry information transfers? Journal of Accounting Research 46, 909–940.

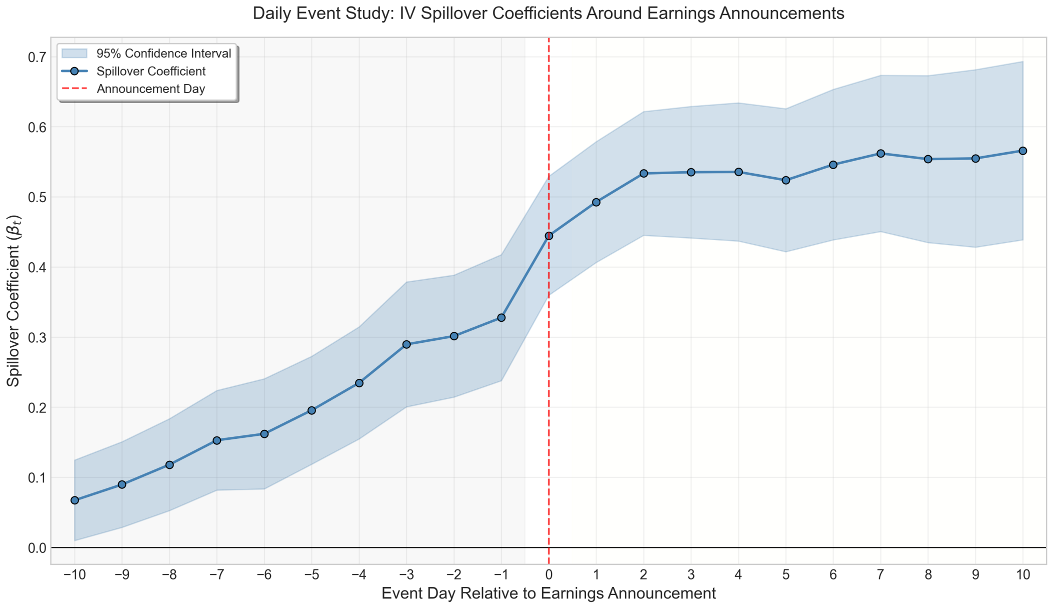

Figure 1: Daily Spillover Coefficients Around Earnings Announcements.
This figure plots the coefficient from daily regressions of peer ΔIV on announcer ΔIV (without fixed effects), estimated separately for each day from $t-10$ to $t+10$ relative to the earnings announcement date. Shaded areas represent 95% confidence intervals based on two-way clustered standard errors.

Table A1
Robustness and Specification Checks

	β	SE	t-stat	N	Note
Panel A: Core Spillover Robustness					
Full sample (Date FE)	0.065	0.019	3.34	66,982	Baseline
Exclude COVID (Date FE)	0.061	0.021	2.90	56,237	Mar 2020–Jun 2021
DF correction (Date FE)	0.065	0.019	3.31	66,982	SE × 1.009
Single-announcer (Date FE)	0.123	0.037	3.36	20,907	
Multi-announcer (Date FE)	0.046	0.021	2.15	46,075	
Panel B: Asymmetry (Date FE)					
IV increase (bad news)	0.090	0.045	2.03	14,554	
IV decrease (good news)	0.032	0.031	1.02	52,428	
Interaction (increase − decrease)	0.064		1.62	66,982	$p = 0.105$

Notes: This table reports robustness checks for the core implied volatility spillover. Panel A compares the baseline month FE specification with alternative samples and fixed effects. Panel B tests for asymmetric spillover effects by the sign of the announcer's IV change. Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

A.1 Robustness and Specification Checks

Table A1 lists additional robustness tests on the key spillover finding. In Panel A the robustness of the findings to different specifications is tested. If the COVID-19 period from March 2020 until June 2021 is removed (dropping 10,745 observations) then the estimate is 0.061 ($t = 2.90$) when using date fixed effects as controls; this confirms that the results cannot be explained solely because of the very high level of volatility experienced during the pandemic. There was a <1% inflation in standard errors due to the use of degrees of freedom corrections for the 1,143 date fixed effects but it did not affect conclusions. The single-announcer subsample shows a coefficient of 0.123 ($t = 3.36$), which again shows there is sufficient variation to support estimation of effects without within-day variation.

Panel B provides evidence on asymmetry in the spillover between implied volatility increases and decreases. Using date fixed effects as controls, the t-statistic for the interaction term testing for whether there was a different spillover given whether the change in announcer implied volatility was an increase or decrease was 1.62 ($p = 0.105$).

A.2 Calibration Test: Overlapping vs. Non-Overlapping Windows

Table A2 reports the tests of volatility calibration (using non-overlapping and overlapping realized volatility windows). Non-overlapping windows start from day +4 to exclude overlap with days 0 through +3 when the volatility is calculated by means of the implied volatility. Overlapping windows start on day +1. This allows the use of as much data as possible (i.e., maximize the statistical power), but it also introduces potential issues since both the realized volatility and the implied volatility are computed at the same time during these days.

There are some differences in the estimated coefficients among different ways to define the window, with a smaller estimate being provided by an overlapping window of ten days (0.830) and a larger estimate by a non-overlapping window of twenty days (1.062). For all non-overlapping window horizons, the null hypothesis $\gamma = 1$ cannot be rejected; this implies that the findings reported in the main body of this paper were unaffected by the way the window was defined.

A.3 Peer Firm Fixed Effects

The peer firm fixed effects are added to the core spillover model using Table A3. The peer firm fixed effects control for time-invariant differences in implied volatility sensitivity among the 1,365 unique peer firms. The spillover coefficient using peer firm fixed effects and month fixed effects was 0.153 ($t = 6.98$). This is near the base case estimate of 0.159 using just month time fixed effects. When peer fixed effects and date fixed effects were used together, the coefficient was 0.061 ($t = 4.81$) with strong statistical significance. These results support the trend observed as the number of time controls increases; that is, more time controls will decrease the size of the point estimates but keep the estimates statistically significant.

Table A2
Calibration Test: Overlapping vs. Non-Overlapping Windows

	10-day	20-day	30-day
Non-overlapping $[+4, +T]$			
γ	0.946***	1.062***	1.033***
	(0.060)	(0.071)	(0.084)
H_0: $\gamma = 1$	$t = -0.91$	$t = 0.87$	$t = 0.39$
Overlapping $[+1, +T]$			
γ	0.830***	0.988***	0.982***
	(0.064)	(0.070)	(0.082)
H_0: $\gamma = 1$	$t = -2.66$	$t = -0.18$	$t = -0.22$
N		66,847	

Notes: This table reports the realized volatility calibration test (peer $\Delta IV \rightarrow$ peer ΔRV) using non-overlapping windows beginning on day +4 and overlapping windows beginning on day +1. Under rational pricing, $\gamma = 1$. Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Table A3
Peer Firm Fixed Effects

	(1) Month FE	(2) Peer FE	(3) Peer + Month	(4) Peer + Date
Announcer ΔIV	0.159***	0.444***	0.153***	0.061***
	(0.024)	(0.036)	(0.022)	(0.013)
t-statistic	6.67	12.44	6.98	4.81
N	66,982	66,982	66,982	66,982
Peer FE	No	Yes	Yes	Yes
Time FE	Month	None	Month	Date

Notes: This table adds peer firm fixed effects to the core spillover specification. Peer firm fixed effects absorb time-invariant differences in average implied volatility adjustments across the 1,365 unique peer firms. Columns (1) and (2) report specifications without peer FE for comparison. Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

A.4 Return Predictability

Table A4 reports regressions of peer cumulative abnormal returns on the announcer's implied volatility change. Without controlling for date fixed effects, there is a negative correlation across all horizons. With date fixed effects, all correlations for horizons greater than 5 days became statistically insignificant. For very short horizons (1 day and 3 days), there is evidence of a small negative relationship ($t = -1.80$ and $t = -1.79$, respectively), which suggests that the options market leads the equity markets. The equity markets show little persistence in their ability to forecast future returns based on peer firm implied volatility spillovers.

A.5 Cross-Sectional Heterogeneity

The spillover coefficient estimated separately by terciles of peer firm characteristics is provided in Table A5. All regression models include date fixed effects. For each dimension of peer firm characteristics tested (i.e., market capitalization, baseline realized volatility, options market depth and earnings surprise magnitude), the difference between the highest and lowest tercile is generally insignificantly different. More specifically, none of these high/low

Table A4
Return Predictability

	CAR(1d)	CAR(3d)	CAR(5d)	CAR(10d)	CAR(20d)
No Fixed Effects					
Announcer ΔIV	-0.018^{***}	-0.046^{***}	-0.050^{***}	-0.061^{***}	-0.068^{***}
	(0.004)	(0.007)	(0.010)	(0.014)	(0.017)
Date Fixed Effects					
Announcer ΔIV	-0.004^{*}	-0.007^{*}	−0.007	0.000	0.003
	(0.002)	(0.004)	(0.005)	(0.009)	(0.012)
t-statistic	−1.80	−1.79	−1.26	0.03	0.26
N			66,896		

Notes: This table regresses peer cumulative abnormal returns (CARs) on the announcer's ΔIV. CARs are buy-and-hold returns adjusted for the market return. Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Table A5
Cross-Sectional Heterogeneity (Date FE)

	Low	Mid	High	H−L	t(H−L)
Peer market cap	0.106^{***}	0.074	−0.005	−0.111	−1.81
Peer baseline RV	0.104^{***}	−0.021	0.047^{**}	−0.057	−1.57
Option market depth	0.083^{***}	0.089^{*}	−0.034	−0.116	−1.80
EPS surprise magnitude	0.021	0.076^{***}	0.088^{***}	0.066	1.62

Notes: This table reports the spillover coefficient estimated separately for terciles of peer firm and event characteristics, all with date fixed effects. Market capitalization, realized volatility, and options depth are peer firm characteristics; earnings surprise magnitude is an announcer (event) characteristic. H−L is the difference between the high and low tercile coefficients. Standard errors are two-way clustered by firm and date. ***, **, * denote significance at the 1%, 5%, and 10% levels.

differences reach statistical significance at the 5% level. In general terms, the spillover appears to be uniformly distributed across observable firm characteristics, supporting an industry-wide explanation for the spillover.

A.6 Option Trading Strategy

Quintile portfolio sorts by the announcers' IV change are reported in Table A6. As expected, the Q5−Q1 spread in peer IV changes is 0.183, confirming a monotonically increasing relationship. The Q5−Q1 spread in proxy straddles is −0.015 ($t = -9.73$), showing that selling volatility bets on peers after large positive announcer IV shocks result in losses. While the unconditional IV-RV gap is smaller in Q5 (0.017) compared to Q1 (0.062), this reflects cross-sectional variation over time (i.e., high-spillover days are typically also high-VIX days where the variance risk premium is depressed). Thus, it does not reflect within-day cross-firm variation. As shown in Table 7 Panel B, the spillover channel actually expands the variance risk premium. There exists no exploitable trading strategy: the structural variance risk premium dominates any spillover-based signal.

Table A6
Option Trading Strategy: Quintile Sorts by Announcer ΔIV

	Q1	Q2	Q3	Q4	Q5	Q5–Q1
Announcer ΔIV	−0.280	−0.172	−0.106	−0.034	0.126	0.406
Peer ΔIV	−0.082	−0.037	−0.009	0.021	0.101	0.183***
Straddle P&L	0.003***	0.000	−0.001	−0.003***	−0.012***	−0.015***
IV-RV gap	0.062***	0.044***	0.031***	0.033***	0.017***	−0.045***

Notes: This table sorts peer-event observations into quintiles by the announcer's ΔIV. Straddle P&L is a proxy for the payoff of a delta-neutral straddle position on the peer firm. IV-RV gap is the difference between peer implied volatility and subsequent realized volatility. ***, **, * indicate the Q5–Q1 spread is significant at the 1%, 5%, and 10% levels based on two-way clustered standard errors.