Intermediary Option Pricing*

Julian Terstegge
University of Michigan

Abstract

I propose a novel inventory risk for option dealers ("gap risk"), which helps explain salient features of the option risk premium. I define "gaps" as equity price changes over nights or weekends, when reduced equity market liquidity impedes continuous delta-hedging, thus leading to unhedgeable inventory risk. Gaps function like large jumps, even if prices move continuously. Using comprehensive option trade data, I estimate intermediaries' option inventory and find substantial exposure to gap risk. This finding helps explain why the option risk premium is heavily concentrated over nights and weekends. Supporting a causal impact of intermediary gap risk on the option risk premium, I show that the emergence of overnight equity trading around 2006 led to significantly less option risk premium over week-nights relative to weekends, particularly in options that appear in intermediaries' inventory.

Keywords: Asset Pricing, Derivatives, Intermediation, Inventory Risk, Liquidity

JEL Codes: G10, G12, G13, G14

* *This version: December 26, 2025, initial version: November 10, 2024.* I thank Svetlana Bryzgalova; Pierre Collin-Dufresne; Ian Dew-Becker; William Diamond; Itamar Drechsler; Paul Fontanier; Stefano Giglio; Paul Huebner; Kris Jacobs; Theis Ingerslev Jensen; Petri Jylhae; Ron Kaniel; Joren Koeter; Tim Kroencke; Lars Christian Larsen; Anton Lines; Dmitriy Muravyev; Chayawat Ornthanalai; Lasse Heje Pedersen; Markus Pelger; Tobias Sichert; Alessandro Spina; Johannes Stroebel; Stephen Szaura; Anders Bjerre Trolle; Marliese Uhrig-Homburg; and Paul Whelan for helpful comments and suggestions. I thank seminar participants at the 2025 FMA Derivatives and Volatility Conference; the University of Michigan; the University of Wisconsin–Madison; Imperial College Business School; the Stockholm School of Economics; the University of Warwick; the Rotman School of Management at the University of Toronto; the Wharton School of the University of Pennsylvania; Frankfurt School of Finance; the 2024 Macro Finance Society Workshop (poster session); the 8th HEC Paris Finance PhD Workshop; and the 2024 Nordic Finance Network PhD Workshop. I gratefully acknowledge financial support from the Center for Financial Frictions (DNRF102) and the Center for Big Data in Finance (DNRF167). All remaining errors are my own. *E-mail:* juter@umich.edu. *Website:* julianterstegge.com.

Does intermediary inventory risk help explain the option risk premium? Average delta-hedged returns on S&P 500 options are about –0.5% per day, and compensation for intermediaries' market-making activities has been proposed as one potential source of these large returns. Assessing whether inventory risk drives the option risk premium is challenging, however, because theory predicts both negative (Garleanu, Pedersen, and Poteshman, 2009) and positive (Chen, Joslin, and Ni, 2019) correlations between intermediary inventory and option prices. As a result, regressions of option returns on intermediary inventory alone do not provide conclusive evidence. The issue is important because option prices are widely used to infer investors' expectations and risk premia (Andersen, Fusari, and Todorov, 2015; Backus, Chernov, and Martin, 2011). Understanding whose expectations and risk preferences are embedded in option prices—only those of end investors or also those of constrained intermediaries—is therefore of broad interest.

To study the impact of intermediary inventory risk on the option risk premium, I exploit a unique setting in the S&P 500 index options market. This setting combines (*i*) comprehensive trade data on S&P 500 options, which reveal intermediaries' option positions and inventory risk; and (*ii*) a large and rapidly growing market that likely constitutes an important share of intermediaries' overall inventories: between 2018 and 2023, open interest and trading volume in S&P 500 options nearly quadrupled, reaching about $250 billion and $2 trillion per year, respectively. In addition, the market for S&P 500 options offers high-frequency variation in hedging frictions along two dimensions: (*iii*) day–night differences in equity liquidity induced by equity trading hours; and (iv) the sharp increase in overnight equity trading around 2006, which relaxed a first-order hedging constraint on weeknights relative to weekends.

The main channel in this paper are hedging frictions from equity illiquidity, which is most pronounced overnight and on weekends. Option dealers hedge their inventories by dynamically trading the underlying index (Ni, Pearson, Poteshman, and White, 2021), so limited overnight liquidity imposes a first-order constraint on their ability to maintain delta-neutral positions. Even if dealers are perfectly delta-hedged at the equity-market close, a large overnight change in the S&P 500 index requires substantial equity trading in order to remain delta-hedged and low overnight equity liquidity prevents dealers from making these hedge adjustments, leading to inventory losses. As a result, option sellers (dealers) carry unhedgeable equity market risk over nights and weekends. Option prices are high - and returns low - to compensate option sellers for this unhedgeable inventory risk.

To summarize, I present three main findings that support this channel. First, following the rise in weeknight equity trading around 2006, the magnitude of the option risk premium between Monday and Friday declines sharply relative to weekend (Friday–Monday) returns, with the effect concentrated in options held in dealers' inventories. Second, I introduce a new option risk measure, "gap risk," which captures the risk associated with price changes over nights and weekends, when limited equity-market liquidity impedes the adjustment of delta hedges. Dealers are substantially exposed to gap risk, and this exposure is missed by standard (gamma-based) risk measures. Third, exploiting day–night variation in hedging frictions, I show that options' exposure to gap risk predicts option returns only overnight, when continuous delta-hedging is constrained, and that this predictive power is stronger when dealers' aggregate gap-risk exposure is high. Taken together, the evidence supports a causal effect of equity-market liquidity on the option risk premium through dealers' inventory risk.

My results differ conceptually from the existing option-pricing paradigm. Since the early studies on option returns in Coval and Shumway (2001) and Bakshi and Kapadia (2003), the central building blocks for explaining the option risk premium have been stochastic-volatility risk and jump risk. I show that a large share of the option risk premium is instead linked to unhedgeable overnight equity price risk arising from equity illiquidity. This "gap-risk" channel implies that a significant portion of the premium compensates constrained intermediaries for inventory risk, rather than investors for exposure to continuous volatility and rare jumps. The concept of gap risk naturally extends beyond S&P 500 options and provides a general framework for understanding any derivatives market whose underlying is intermittently illiquid.

These results are also timely. Major U.S. equity exchanges have announced plans to extend weekday trading toward a nearly 24-hour schedule. Because I show that increases in overnight trading activity compress the option risk premium, the move toward deeper around-the-clock equity markets implies a further decline in option risk premia if overnight liquidity improves substantially.

The paper starts from the finding of Jones and Shemesh (2018) and Muravyev and Ni (2020), that the risk premium in S&P 500 options is heavily concentrated over nights and weekends. The liquidity of the underlying stocks and associated futures contracts is reduced by an order of magnitude at nights relative to the day. Therefore, hedging frictions from reduced equity market liquidity, constitute a possible explanation for the concentration of the option risk premium. However, nights differ in other important ways from days, for example the information environment. Thus, the first step of this paper addresses the question whether equity liquidity affects the option risk premium and thus weather the concentration of the risk premium over nights and weekends might be explained in this way.

Turning first to the role of equity-market liquidity in shaping the option risk premium, I exploit the sharp increase in overnight equity trading around 2006. Several market-design changes in the mid-2000s—most notably the implementation of Regulation National Market System (Reg NMS) and the acquisitions of major electronic communication networks by NYSE and Nasdaq—substantially increased overnight trading volumes from Monday through Friday, while leaving the weekend period largely non-tradable. If overnight equity illiquidity drives the option risk premium, then intra-week option returns, for which hedging frictions were relaxed, should become less negative relative to weekend returns, for which hedging frictions were not. This is exactly what I find.

Figure 1 shows that option returns within the week (Monday to Friday) became significantly less negative after the emergence of overnight trading, relative to Friday–Monday returns. The effect is strongest for short-dated, out-of-the-money puts, precisely where intermediaries' short exposures and hedging constraints are most pronounced. These findings suggest that improvements in overnight equity-market liquidity have a causal impact on option risk premia, operating through intermediaries' gap-risk exposure. The logic parallels Du, Tepper, and Verdelhan (2018), who show that covered interest rate parity deviations spike in forward contracts that appear on quarter-end bank balance sheets and interpret this pattern as evidence that post-crisis regulation causally affects those deviations. The results also relate to Dew-Becker and Giglio (2023), who document a decline in the option risk premium around the Great Financial Crisis; my decomposition of option returns into night and weekend components points to the rise in overnight equity-market liquidity as a natural explanation.

To explain the impact of equity liquidity on the option risk premium, I introduce a new option risk measure, gap risk, which is the risk of equity price changes, or gaps, over nights or weekends, when reduced equity market liquidity impedes continuous delta-hedging. Gaps function like large jumps, even if equity prices moves continuously. Gap risk is frequently discussed by traders, but so far largely absent from academic finance. For example the financial times mentions legal clauses in

Figure 1: The Effect of Equity Liquidity on the Option Risk Premium

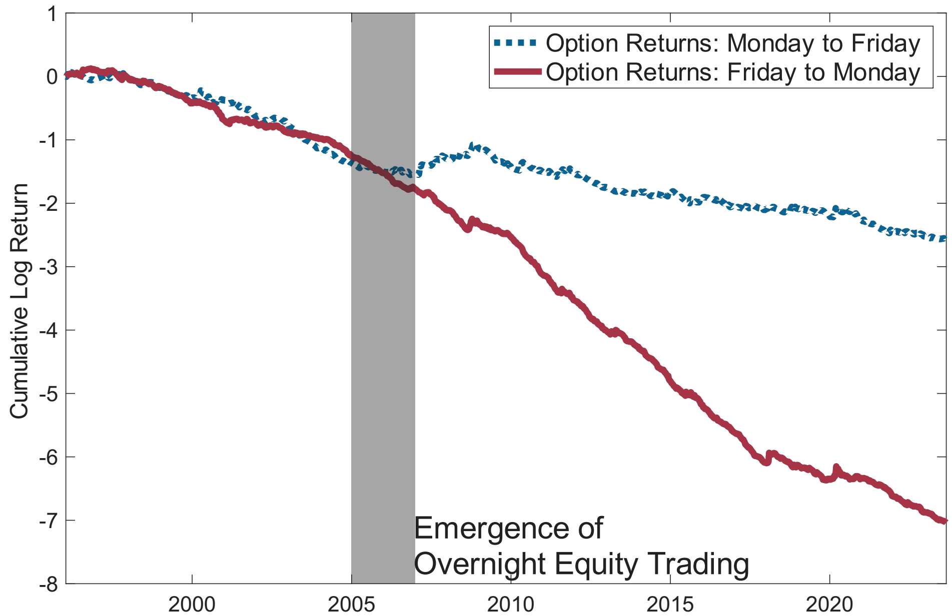

Note: This figure shows cumulative log returns of S&P 500 put options. The dotted line cumulates returns from Monday to Friday, the solid line cumulates returns from Friday to Monday. Returns are measured between trading daily market close and delta-hedged at the beginning of the respective period. Puts are out-of-the-money, with at least two days to expiry. The Monday to Friday return is the average over the four returns between Monday close and Friday close. Returns are in logs and are scaled to the same 10% annualized volatility. The vertical line indicates the emergence of overnight equity trading around 2006.

levered oil products "known as a severe overnight gap event",[1] and Bloomberg reports that "gap risk — sudden moves with little trading in between — is growing".[2] Gap risk is largest at short maturity and deep out of the money, as these options experience returns of often several hundred percent, when the underlying moves by 10% over an illiquid window.

To estimate intermediaries' exposure to gap risk, I use comprehensive trade data from the CBOE. I estimate dealers' aggregate position between 2011 and 2023. Intermediaries hold a persistent net short in puts (about 19 million contracts) while call positions hover near zero. Given this imbalance, a large overnight decline exposes dealers to losses unless hedges are rebalanced in thin markets. Under a hypothetical -10% S&P 500 move, dealers loose a significant fraction of their option inventory value absent re-hedging, and maintaining delta-neutrality requires multi-billion-dollar equity sales. Typical overnight E-mini turnover is well below $1 billion per

[1] https://www.ft.com/content/2fd21b4e-e28c-4e56-8c95-a45e877df9c8?, accessed on November 22, 2025.

[2] https://www.bloomberg.com/news/articles/2019-05-03/stock-volatility-isn-t-dead-it-s-just-got-freakish-and-extreme?, accessed on November 22, 2025.

hour, versus more than \$20 billion per hour intraday (plus roughly \$27 billion per hour in underlying stocks), making continuous re-hedging infeasible overnight and creating overnight equity market risk.

Standard (gamma) risk measures miss this exposure. Aggregate dealer inventory gamma is near zero or slightly positive, indicating that the dealer option portfolio experiences gains from small equity price moves. However, re-calculating dealer inventory gamma for the scenario where the S&P 500 has fallen by, for example, 10%, reveals a large negative aggregate dealer inventory gamma. I refer to such scenario gamma as "shadow gamma". Shadow gamma captures dealers exposure to large equity price moves over (night) periods, where the delta-hedge adjustments are prohibitively expensive. Option dealer shadow gamma predicts option returns, particularly in out-of-the-money puts, where shadow gamma is concentrated.

Although I develop gap risk in the context of S&P 500 index options, the concept applies more broadly. Any derivative that is dynamically hedged in an underlying whose liquidity varies over the calendar—across day and night, across time zones, or around holidays—exposes its intermediaries to gap risk. Examples include equity index options in other markets, single-stock and small-cap options where even intraday order books are thin, volatility derivatives such as VIX futures and options, commodity and credit options hedged in illiquid OTC markets, and structured products replicated with option portfolios. In all of these settings, a shadow-gamma–type measure provides a general way to quantify and price the risk that hedge adjustments cannot be executed when large moves occur, suggesting a wide scope for applying the framework beyond S&P 500 options.

Finally, I estimate the price of gap risk in the options market. I regress option returns over days and nights onto their lagged gap risk and find significant predictive power only over night periods. The predictive power of gap risk for options overnight returns is increasing in intermediaries inventory gap risk exposure. showing that the price of gap risk is related to intermediaries inventory.

I conclude that option prices and risk-premia reflect the risk-perceptions and risk-aversion of constrained intermediaries as well as those of households. My result suggest that if equity markets where perfectly liquid around the clock then the option risk premium would be substantially reduced, and possibly indistinguishable from zero. Failing to account for this will likely overestimate the disaster risk aversion of the average household. How best to infer household risk perceptions in a market where intermitted intermediary hedging constraints are crucial for the risk premium, is a potentially promising direction for future research.

Contributions. This paper contributes to the literature on option returns. Black and Scholes (1973) and Merton (1973) establish the risk-free rate as the benchmark for delta-hedged option returns. Coval and Shumway (2001) and Bakshi and Kapadia (2003) show that delta-hedged returns of equity index options are negative, especially for puts. Broadie, Chernov, and Johannes (2009) show that delta-hedge option returns point towards a priced jump risk premium. Jones and Shemesh (2018) and Muravyev (2016) show that option returns are especially negative over weekends and nights, and provide explanations based on market mispricing. Bollen and Whaley (2004), Garleanu, Pedersen, and Poteshman (2009), Fournier and Jacobs (2015), among others, study the relation between intermediary inventory and the option risk premium. Closest to this paper is Gruenthaler (2022), who introduces an option intermediary risk factor based on stress-test scenario P&L of dealers' option inventory and finds predictive power for option returns. In contrast, I quantify dealers inventory risk from overnight equity illiquidity, and find predictive power over (night) periods when hedging frictions are elevated.

This paper contributes to the literature on liquidity premia in options markets. Cao and Han (2013) and Kanne, Korn, and Uhrig-Homburg (2023) show that stock option risk premia decrease in the liquidity of the underlying stocks. Christoffersen, Feunou, Jeon, and Ornthanalai (2021) estimate a model where the crash probability of the S&P 500 depends on its' liquidity. To the best of my knowledge, this is the first paper to provide evidence in support of a causal impact of market liquidity on option risk premia.

Finally, this paper contributes to the literature on intermediary asset pricing. Haddad and Muir (2021) and He, Kelly, and Manela (2017) show that intermediary constraints can explain variation in returns across many asset classes. Du, Tepper, and Verdelhan (2018) show increased cip deviations around quarter-ends, when banks face tighter capital requirements. Bates (2022) notes that the role of intermediary constraints in shaping option returns remains unresolved. I exploit high-frequency variation in hedging frictions to link the option risk premium to hedging frictions and intermediary constraints.

I. Markets and Data

This section outlines the market for S&P 500 options, and briefly comments on stocks, futures, and ETFs as the most common delta-hedging instruments. The section states data sources, while details on variable construction are in the respective sections.

S&P 500 Options. This paper studies S&P 500 equity index options (SPX options), that is, put- and call options written on the S&P 500 equity index of U.S. large-cap stocks. SPX options are exchange-traded exclusively on the Chicago Board Options Exchange (CBOE) and were initially listed in 1983. While SPX option volumes were initially small, volumes have grown to about $\$2tn$ a year in 2023 and open interest has grown to about $\$250bn$. The SPX options market is the worlds' largest and most liquid equity options market. The high option liquidity enables a return decomposition at high frequency and the large market size makes SPX options an economically relevant market to study.

The original SPX options expired once a month on that months' 3^{rd} Friday. Recently, the CBOE has successively added SPXW options with different expiry dates.[3] I include both SPX and SPXW options into the analysis of this paper and refer to both jointly as S&P 500 options. SPX options are liquid across a broad range of strike prices, which occur every $5. Liquidity is particularly high for out-of-the-money options, which are puts (calls) with strike prices below (above) the current value of the underlying index. S&P 500 options are European, which can only be exercises at expiry.

A major advantage of the S&P 500 index options market for the study of intermediary asset pricing is the availability of comprehensive trade data. S&P 500 index options trade exclusively on the Chicago Board Options Exchange (CBOE) and the CBOE makes datasets commercially available that allow for the daily measurement of the options position of the intermediary sector. Intermediaries' options position provide information on their risk exposures and risk management.

S&P 500 Stocks, Futures and ETFs. The most important delta hedging instruments for S&P 500 options are the underlying stocks, as well as associated Futures and ETFs. S&P 500 constituent stocks and SPY (the S&P 500 exchange-traded fund) trade on U.S. exchanges (NYSE, Nasdaq). Regular exchange hours are 09:30–16:00 (E.T.); by the mid-2000s, exchanges expanded

[3] Specifically, the CBOE added weekly Friday expiries in 2011.09, Wednesday expiries in 2016.02, Monday expiries in 2016.08, Tuesday expiries in 2022.04, and Thursday expiries in 2022.05.

extended hours to 04:00–20:00, yet there is little equity trading from 20:00 to 04:00. Pre-market activity is thin (about 0.27% of S&P 500 stock volume) and post-market thinner still (about 0.12%), and intraday equity volumes follow a U-shape, with weekend trading closed entirely. ETFs share the same venues and hours as their underlying stocks. S&P 500 exposure via futures is provided by the CME's E-mini contract on Globex, which trades nearly 24/5, from Sunday 18:00 to Friday 17:00 with a daily 17:00–18:00 maintenance break; E-mini trading supplanted the standard pit-traded contract over time. Volumes differ starkly between day and night: average intraday E-mini dollar volume (2011–2022) is about \$164bn per day versus only about \$7bn overnight (a day to night ratio near 25:1), and at higher frequency, intraday futures turnover runs roughly \$8–\$15bn per 30-minute window compared with about \$250m per 30-minute window overnight. Consistent with this, aggregate equity volumes rose from roughly \$20bn per day in the late 1990s to about \$200bn per day after 2020, but remain small in the pre- and post-market windows. The online appendix contains tables and figures on equity returns and volumes over day and night periods.

Data Sources. From CBOE, I obtain S&P 500 option prices and quotes at 15-minute intervals. Further, I obtain the daily "Open Close Volume" files that allow for the construction of intermediary positions. From OptionMetrics, I obtain S&P 500 option prices and quotes at the daily frequency. I obtain data on S&P 500 E-mini futures from Boyarchenko, Larsen, and Whelan (2023), who sample tick-level data of CME traded futures contracts. I obtain data on risk-free rates from the OptionMetrics IvyDB zero-curve file. Data on daily stock trading volume is from CRSP. High-frequency stock volumes are from Reuters. From Eikon, I obtain tick level data on S&P 500 E-mini futures traded on the CME, comprising best bid offers, trade prices, and volumes.

II. Gap Risk and Options' Shadow Gamma

This section introduces the concept of "gap risk". Gap risk is the risk of price changes over periods where liquidity is low or absent. The risk of equity price changes over nights or weekends are natural examples.

Subsequently, this section discusses the riskiness of options, due to equity price gap risk. Equity price gap risk contributes to option risk, since delta hedges can be adjusted only imperfectly over periods where the underlying is illiquid. Options' gap risk is not well captured by the conventional option risk measure of gamma. Instead the new concept of ”shadow gamma” proves more useful.

This paper studies gap risk in the context of S&P 500 options and intermediaries' unhedgeable inventory risk. However, the potential applications of gap risk are more broad. Gap risk emerges whenever an asset or its underlying experience sporadic drops in liquidity to a value close to zero. Natural examples are exchange close times, holidays, volatility- or price limit trading halts, long time zone gaps between derivative and underlying trading hours (for instance U.S. traded options on Asian or European indices), and derivatives written on thin OTC underlyings such as corporate bonds or CDS, where trading frequently drops to zero for extended periods.

II.A. Gap Risk

Gap risk, the risk of price changes over periods where liquidity is low, is frequently discussed by traders. For example the financial times mentions legal clauses in levered oil products "known as a severe overnight gap event",[4] and Bloomberg reports that "gap risk — sudden moves with little trading in between — is growing".[5] The term "gap" in relation to derivatives hedging has previously been used by Tankov (2008) when analyzing exotic equity derivatives (gap options). I apply the concept of gap risk more generally to plain vanilla equity index options.

First, and most importantly, gaps differ from jumps in their magnitude. Jumps in the S&P 500 tend to be small, while equity returns over nights or weekends can amount to several percent. Table A.6 shows S&P 500 returns over days, nights and weekends. For this paper, especially the left return tail over nights and weekends is important. These exceed -7% and -12% since 1996. Returns of these magnitudes are not well conceptualized as small jumps and, as I argue below, options' exposure to such gaps is not well approximated via gamma.

A second difference between gaps and jumps is that the timing of the former is typically predictable. Except around pre-scheduled economic announcements, jumps tend to occur at random times and tend to be modeled with some probability of arrival (Merton, 1976; Duffie, Pan, and Singleton, 2000) Since gaps tend to be anticipated, especially in case of nights, risk then lies in the magnitude of the gap.

[4] https://www.ft.com/content/2fd21b4e-e28c-4e56-8c95-a45e877df9c8?, accessed on November 22, 2025.

[5] https://www.bloomberg.com/news/articles/2019-05-03/stock-volatility-isn-t-dead-it-s-just-got-freakish-and-extreme?, accessed on November 22, 2025.

II.B. Gamma and Shadow Gamma

Gamma, small jumps, and delta-hedged P&L. Options' exposure to jumps is typically approximated via gamma (Γ). Γ is the second-order derivative of the option price with respect to the price of the underlying. As such, Γ captures the *local* curvature or convexity of the option value function. While the slope of the option value function can be hedged against equity price changes with a static equity position, the curvature cannot.

Consider an option with value $V(t, S)$, delta $\Delta_t = \partial_S V(t, S_t)$, and gamma $\Gamma_t = \partial^2_{SS} V(t, S_t)$. A dealer who delta-hedges holds the option and $-\Delta_t$ units of the underlying, so the value of the hedged portfolio is

$$\Pi_t = V(t, S_t) - \Delta_t S_t. \tag{1}$$

Over a short interval $[t, t + \Delta t]$ the underlying may experience a price change $\Delta S = S_{t+\Delta t} - S_t$. The change in the value of the hedged portfolio due to this move is

$$\Delta\Pi(\Delta S) = V(t + \Delta t, S_t + \Delta S) - V(t, S_t) - \Delta_t \, \Delta S. \tag{2}$$

For a small move ΔS use a second-order Taylor expansion around (t, S_t):

$$V(t + \Delta t, S_t + \Delta S) \approx V(t, S_t) + \partial_t V(t, S_t) \, \Delta t + \Delta_t \, \Delta S + \tfrac{1}{2}\Gamma_t (\Delta S)^2. \tag{3}$$

Plugging this into (2) and canceling the first–order term in ΔS yields

$$\Delta\Pi \approx \partial_t V \, \Delta t + \tfrac{1}{2}\Gamma_t \, (\Delta S)^2. \tag{4}$$

Ignoring the smooth time-decay term $\partial_t V \, \Delta t$, the jump-induced P&L of the delta-hedged position is

$$\Delta\Pi(\Delta S) \approx \tfrac{1}{2}\Gamma_t (\Delta S)^2. \tag{5}$$

Equation (5) shows why gamma is the standard measure of option exposure to small jumps: once delta risk has been hedged out, the remaining exposure to small price moves is proportional to Γ_t and scales with the square of the jump size. An options portfolio with a positive gamma gains in value from ΔS when hedges remain unadjusted.

Shadow gamma and large gaps. For larger "gaps" ΔS the local approximation in (3) becomes inaccurate, because the curvature (Γ) of V changes along the move. The delta-hedged P&L for any finite jump can be written as

$$\Delta\Pi(\Delta S) = V(t, S_t + \Delta S) - V(t, S_t) - \Delta_t\, \Delta S. \tag{6}$$

This can be expressed in "gamma form"

$$\Delta\Pi(\Delta S) = \tfrac{1}{2}\Gamma_t^{\text{sh}}(\Delta S)\,(\Delta S)^2, \tag{7}$$

by defining the *shadow gamma* for price move size ΔS as

$$\Gamma_t^{\text{sh}}(\Delta S) \equiv \frac{2}{(\Delta S)^2}\Big[\,V(t, S_t + \Delta S) - V(t, S_t) - \Delta_t\, \Delta S\,\Big]. \tag{8}$$

By construction, (7) is an identity for any ΔS.

By integrating the option's gamma along the price path from S_t to $S_t + \Delta S$, $\Gamma_t^{\text{sh}}(\Delta S)$ can be written exactly as

$$\Gamma_t^{\text{sh}}(\Delta S) = 2\int_0^1 (1-u)\,\Gamma\big(t, S_t + u\,\Delta S\big)\,\mathrm{d}u. \tag{9}$$

That is, shadow gamma is a weighted average of the option's future curvature along the path from S_t to $S_t + \Delta S$ and shadow gamma captures the exact change in the value of a delta-hedged options portfolio for large changes in the price of the underlying.

Note that, $\lim_{\Delta S\to 0}\Gamma_t^{\text{sh}}(\Delta S) = \Gamma_t$ so that shadow gamma coincides with local gamma for small changes in S, but for large S changes $\Gamma_t^{\text{sh}}(\Delta S)$ loads on the curvature of $V(\cdot)$ in the post-jump state.

For deep out-of-the-money puts, Γ_t can be close to zero, even though a large downward move would take the option close to or into the money, where curvature is large. Local gamma therefore understates the overnight risk of such positions, whereas $\Gamma_t^{\text{sh}}(\Delta S)$ remains large in magnitude because it incorporates the future curvature after a sizeable move.

Equation (9) defines the exact shadow gamma $\Gamma_t^{\text{sh}}(\Delta S)$ as a weighted average of local gammas along the entire price path from S_t to $S_t + \Delta S$. In the empirical analysis I use a tractable proxy,

$$\widetilde{\Gamma}_t^{\text{sh}}(\Delta S) \equiv \Gamma\big(t, S_t + \Delta S\big), \tag{10}$$

and in particular

$$\widetilde{\Gamma}_t^{\text{sh}}(-10\%) = \Gamma\big(t, 0.9 S_t\big), \tag{11}$$

For the short-dated, deep out-of-the-money puts that dominate dealers' inventories, $\Gamma(t, S)$ is negligible at S_t and along most of the path, and spikes only once the index has fallen close to the strike, so the integral in (9) is numerically dominated by the neighborhood of $S_t + \Delta S$. Evaluating gamma at the shocked underlying therefore provides a simple and economically meaningful proxy for the theoretical shadow gamma. In what follows, I refer to $\widetilde{\Gamma}_t^{\text{sh}}$ simply as shadow gamma.

Interpreting the magnitude of $\widetilde{\Gamma}_t^{\text{sh}}(-10\%)$ is straightforward. $\widetilde{\Gamma}_t^{\text{sh}}(-10\%)$ is the curvature of the option value function (both hedged and un-hedged) in the scenario of a -10% underlying return. An option portfolios shadow gamma is the sum-product of the option positions (in no. of options) and the respective options' individual shadow gamma. Thus, an option with a shadow gamma of 0.05 adds 0.05 units of curvature to a portfolio. The price of shadow gamma, as estimated in subsequent sections, quantifies how much investors are willing to pay in terms of risk premium for an additional unit of curvature in their portfolio value function.

Table I shows the average $\widetilde{\Gamma}_t^{\text{sh}}(-10\%)$ across option portfolios, that I estimate for options over my sample. The properties of gamma naturally translate to the properties of $\widetilde{\Gamma}_t^{\text{sh}}(-10\%)$, such as strict positivity for options. Shadow gamma is heavily concentrated in short-maturity deep out-of-the-money puts , with a a value of 1.2076 against, a value of 0.0023 for short-maturity in-the-money puts. A value of 1.2 indicates that buying one short-maturity deep out-of-the-money put on average added a curvature of 1.2 units at -10% to an investors portfolio value function. Subsequent section find that inventors pay a substantial premium for this curvature, at least partially due to intermediaries' inventory risk.

Table I: Options' Shadow Gamma at −10% (scaled by 1000)

Puts		Days to Expiry 2-70	71-	All
$0.00 \leq \|\Delta\| \leq 0.25$	Deep Out of the Money	1,207.6	176.0	1,464.7
$0.25 < \|\Delta\| \leq 0.50$	Out of the Money	132.5	109.9	245.0
$0.50 < \|\Delta\| \leq 0.75$	In the Money	23.0	21.8	45.4
$0.75 < \|\Delta\| \leq 1.00$	Deep In the Money	2.2	1.6	3.9
All		1,365.4	310.3	1,762.8
Calls				
$0.00 \leq \|\Delta\| \leq 0.25$	Deep Out of the Money	5.4	9.6	15.2
$0.25 < \|\Delta\| \leq 0.50$	Out of the Money	28.4	41.6	70.8
$0.50 < \|\Delta\| \leq 0.75$	In the Money	73.7	52.8	128.0
$0.75 < \|\Delta\| \leq 1.00$	Deep In the Money	130.8	17.5	155.8
All		252.6	121.6	392.5

Note: The table shows options' shadow gamma. Shadow gamma is the expected option gamma, conditional on a hypothetical return in the underlying asset (here: −10%). Shadow gamma is multiplied by 1000 for readability. Days to expiry are calendar days. The sample period is 2011 to 2023.

Numerical example. A simple numerical example highlights the three key points of this paper: (*i*) the option portfolio has a positive gamma, indicating gains from small movements in S, but negative shadow gamma, indicating losses from large movements in S. (*ii*) Simulations confirm large negative returns on the options portfolio from large movements in S. (*iii*) Delta-hedge adjustments would require relatively large trade volume in stocks. The next subsection shows that dealers actual inventory is similar to this stylized example.

Consider the following portfolio of options: A long position in one put option with strike $K = 5200$ and a short positions in five puts with strike 4800. Assume the underlying is the S&P 500 equity index with a current value $S = 5200$, the risk free rate is 0, time to expiry is 7 calendar days and the expected equity return volatility is 30%. Under BSM assumptions the initial portfolio value is \$75, with a *positive* gamma of 4.6bps, while shadow gamma is *negative* at −85bps. Delta-hedging this position would involve a long position in the S&P 500 of \$1888. Assume now that the S&P 500 crashes by 10% over a weekend to $S = 4680$. The delta-hedged option portfolio experiences a return of −772%. Adjusting the initial delta-hedge and thus avoiding these losses would require selling \$14129 worth of S&P 500 stocks.

III. The Effect of Equity Liquidity on the Option Risk Premium

First, this section confirms the findings of Jones and Shemesh (2018) and Muravyev and Ni (2020), that the option risk premium materializes mostly over weekends and nights. Subsequently, I provide evidence that links the option risk premium to equity illiquidity, thus presenting a possible explanation for overnight concentration of the option risk premium.

To that end, I exploit the sharp increase in overnight equity trading around 2006. Substantial weeknight trading activity emerged only after Nasdaq and NYSE acquired major electronic communication networks, marking a structural shift in market design. This institutional change provides a natural experiment for a difference-in-differences estimation: option returns realized within the week (Monday to Friday) serve as the treated group, and weekend returns (Friday to Monday) as the control.

I find a substantial reduction in the treatment group risk premium relative to the control group risk premium, especially in (out of the money put) options where market makers' inventory risk is concentrated. These results suggest a causal link between equity market liquidity and the option risk premium, operating through intermediaries' ability to manage inventory risk.

III.A. The Option Risk Premium is Concentrated at Night

Option Returns and Delta-Hedging. Throughout the paper, I calculate delta-hedged option returns as:

$$R_t^i = \frac{P_t^i - P_{t-1}^i \; - \; \Delta_{t-1}^i \; \times \; (\text{SP500}_t - \text{SP500}_{t-1})}{P_{t-1}^i}, \tag{12}$$

where P_t^i is the option mid-quote at the end of period t, $SP500$ is the S&P 500 level and Δ_{t-1}^i is the lagged option delta. This return removes first-order exposure to the equity risk premium and follows the standard approach in option-pricing studies (for example, Bakshi and Kapadia (2003); Jones and Shemesh (2018)). The implicit assumption that trading the underlying requires no capital is reasonable given the depth of S&P 500 futures during regular hours. Delta-hedging neutralizes small moves in the underlying but not larger ones: as prices change, deltas change, forcing continuous hedge rebalancing.

Data: High-Frequency Option Returns. I use high-frequency data on S&P 500 index options from the Chicago Board Options Exchange (CBOE). The dataset aggregates observations at 15-minute intervals, beginning at 09:45 E.T. (fifteen minutes after the regular market open) and ending at 16:15, the official market close. For each interval, it reports bid and ask quotes, as well as first, last, high, and low trade prices. The data further include option volume, open interest, and pre-computed risk measures such as delta and gamma.[6]

To mitigate concerns about liquidity and data quality, I apply a series of standard filters. I exclude options with zero trading volume on any of the previous three days or with zero volume at the start of the return period. Accordingly, to be included in the night (day) portfolio, an option must have traded for three consecutive days and have a valid trade between 16:00–16:15 (09:30–09:45) prior to the return interval. I remove observations with negative lagged bid–ask spreads, zero bids, or lagged mid-quotes below $0.05. I also exclude extreme reversal returns, exceeding 1,000% in one period and –90% in the next. Finally, I discard all violations of no-arbitrage bounds. Merging the resulting dataset with intermediary position data yields the main sample period of 2011–2023. All filters are based solely on lagged information and thus avoid look-ahead bias. These procedures closely follow Jones and Shemesh (2018) and Muravyev and Ni (2020).

I measure night returns from 16:15 to 09:45 (E.T.) and day returns from 09:45 to 16:15. Open prices are recorded at 09:45 because the dataset aggregates quotes into 15-minute intervals; this timing also mitigates concerns about illiquid opening quotes. Throughout the analysis, I use mid-quotes as prices. Option returns are delta-hedged with S&P 500 E-Mini futures at the start of each period: the night-time hedge is established at 16:15 and held constant until the next morning. Option deltas are estimated from the Black–Scholes–Merton model, setting the volatility of the underlying equal to the option's lagged implied volatility. Implied volatility is lagged to avoid bias from the negative correlation between volatility and market returns. Option returns are reported in excess of the risk-free rate, which is negligible over intraday or overnight horizons and thus does not affect the results.

Figure 2 shows summary statistics for delta-hedged option returns. The two left bars contain out-of-the-money put options, the two right bars show out-of-the-money calls. Over 2011 t0 2023,

[6] A detailed description of U.S. high-frequency option data and its construction is provided in Andersen, Archakov, Grund, Hautsch, Li, Nasekin, Nolte, Pham, Taylor, and Todorov (2021).

Figure 2: Option Risk Premia During the Day and Overnight

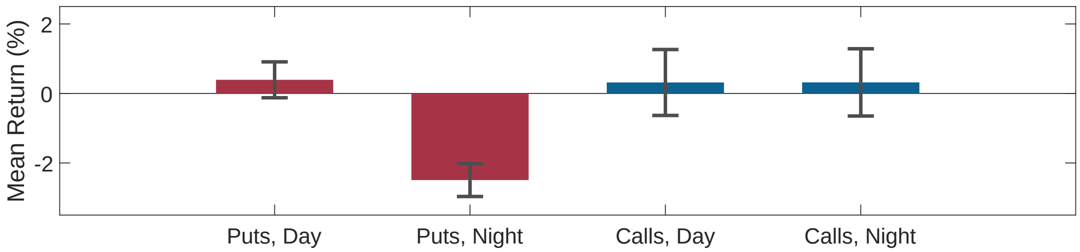

Note: This figure plots average delta-hedged returns of out-of-the-money S&P 500 put and call options. Bars show mean returns with 95% confidence intervals. Day returns are measured between 09:45 and 16:15 (E.T.), and night returns between 16:15 and 09:45. Returns are delta-hedged at the start of each period and expressed in percent. The sample period is 2011 to 2023.

S&P 500 put option experienced an average night return of -2.49%.[7] The associated Newey-West t-statistic exceeds 10. Put option intraday returns average only 0.39%. The difference between night and day returns is highly significant. In contrast, S&P 500 call options experienced an average night return of 0.32% and an equal average day return. Neither the night return, the day return or the difference between the two is significantly different from zero for call options. Thus, negative option risk premia are concentrated in puts and arise almost entirely overnight.

The strong concentration of option risk premia overnight suggests that returns are linked to periods when intermediaries face the greatest difficulty adjusting their hedges. To test this mechanism directly, I next relate overnight option returns to measures of option level risk and intermediary inventory risk.

III.B. The Emergence of Overnight Equity Trading

The institutional foundations of overnight equity trading developed gradually, but the key structural changes can be traced to the mid-2000s. In the early 1990s, the New York Stock Exchange (NYSE) began limited off-hours trading sessions, and by the late 1990s, regulatory and technological innovations accelerated the transition toward electronic markets. Regulation ATS (Alternative Trading Systems) of 1998 formally distinguished ATSs from registered exchanges and heightened transparency requirements, prompting several electronic communication networks (ECNs) to consolidate and register as full exchanges.

A sequence of mergers and technological upgrades followed. In 2001, Archipelago ECN merged

[7] Average returns are not annualized. Option risk premia are very large relative to most other traded assets.

Figure 3: The Emergence of Overnight Equity Trading

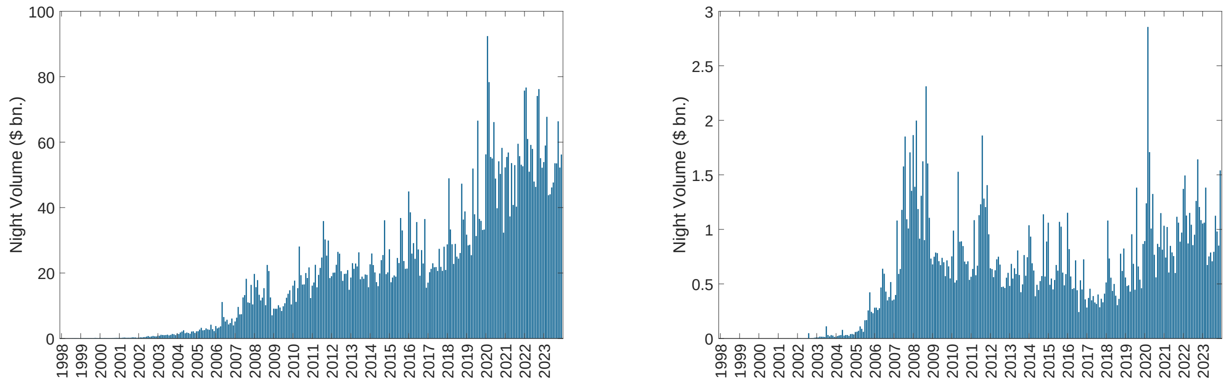

(a) S&P 500 Futures

(b) S&P 500 ETF

Note: This figure shows the rise of overnight equity trading activity in S&P 500 instruments. Panel (a) reports the monthly average dollar volume transacted overnight in the most actively traded S&P 500 E-mini futures contract, while Panel (b) shows the corresponding series for the SPY ETF. Overnight trading is defined as transactions between 16:00 and 09:30 (E.T.). Volumes are expressed in billions of dollars.

with the Pacific Exchange to form ArcaEx, the first fully electronic stock exchange. In 2005, the NYSE launched its Hybrid Market, combining electronic order matching with traditional floor trading, and one year later, after becoming a public company, it acquired Archipelago and phased out the open-outcry system. Around the same time, Nasdaq sought to regain competitiveness by acquiring Instinet and, following its 2005 IPO, expanded trading hours from 08:00–18:30 to 04:00–20:00. These reforms were completed by 2006 and established the extended stock trading hours that define modern U.S. equity markets.

Consistent with these market changes, meaningful overnight equity trading activity emerged only around 2006. Although S&P 500 E-mini futures had traded nearly around the clock between Mondays and Fridays since their 1996 launch, overnight volumes remained negligible for almost a decade. Figure 3 documents the subsequent take-off in trading activity. Panel (a) plots the monthly average overnight dollar volume in the most actively traded S&P 500 E-mini futures contract, while Panel (b) shows the corresponding series for the SPY ETF. Both series rise sharply beginning around 2006. While overnight ETF volumes are smaller in absolute terms, their emergence indicates that meaningful stock trading had become possible outside regular exchange hours.

III.C. Risk Premia Around the Emergence of Overnight Equity Trading

The sharp increase in week-night trading, contrasted with persistently closed weekend markets, provides a natural setting to study how equity liquidity affects the option risk premium.[8] Intra-week option returns, which include nights when equity markets are tradable after 2006, serve as the treated group, while weekend returns, which span periods of more than forty-eight hours without equity or futures trading, serve as the control. If higher equity liquidity relaxes dealers' hedging frictions, option returns within the week should become less negative relative to those over weekends following the onset of overnight trading. Because reliable high-frequency option quotes are not vailable before the late 2000s, I implement the test using daily close-to-close data, which allows for consistent measurement of option returns over the full 1996–2023 sample, while avoiding biases from sparse intraday observations.

Data: Daily Option Returns. OptionMetrics aggregates option trades at the daily frequency, such that all available observations are at 16:00. I obtain option's bid quote, ask quote, and delta. While OptionMetrics applies a proprietary method for calculating options' deltas, their deltas are typically close to Black-Scholes deltas where sigma is set equal to the options implied volatility. To alleviate concerns of liquidity and data errors, I apply several filters to the data. I exclude options which fulfill any of the following conditions: zero trade volume on any of the previous three days, less than two days to expiry, the mid-quote violates no-arbitrage bounds, lagged mid quotes below \$0.05, lagged bids of 0, non-positive lagged bid-ask spreads, lagged spreads above \$10, lagged spreads above the mid-quote, lagged spreads above twice the bid, a large hedged or un-hedged returns above 1000% immediately followed by −90%, or vice versa.

To examine how the option risk premium changes with the emergence of overnight equity trading, I estimate the following difference-in-differences specification:

$$R_t^i \;=\; \beta_1 \,\text{Intra-week}_t \;+\; \beta_2 \,\text{Post}_t \;+\; \beta_3 \,\text{Intra-week}_t \;\times\; \text{Post}_t \;+\; \epsilon_t^i, \tag{13}$$

where R_t^i denotes the delta-hedged return of portfolio i of out-of-the-money S&P 500 put options. $IntraWeek_t$ is an indicator for returns that fall within the week (i.e., any close-to-close return

[8] A potential concern is that the start of overnight equity trading might have changed the equity return distribution in significant ways and thus affected the option risk premium. In unreported robustness tests, I do not find a significant change in equity return volatility or equity return skewness over nights relative to weekends from before 2006 to after 2006.

Table II: The Effect of Equity Liquidity on Put Option Returns: DiD

	(1) Baseline	(2) Winsorized	(3) Excl. Crashes
Intra-week	86.7	95.0**	86.7
	(1.42)	(2.08)	(1.42)
Post	-264.5***	-292.5***	-358.0***
	(-3.16)	(-5.75)	(-5.11)
Intra-week × Post	355.8***	304.2***	429.4***
	(4.05)	(5.33)	(5.41)
Constant	-231.4***	-278.8***	-231.4***
	(-4.05)	(-6.89)	(-4.05)
Observations	6,958	6,958	6,812
R2-adjusted	0.01	0.02	0.01

Note: The table shows regression estimates of Equation (13), where option returns are regressed on indicators for treated periods (intra-week returns), and the post event period. Option returns are for the portfolio of out-of-the-money S&P 500 puts. Intra-week returns comprise of all daily close-to-close returns, except Friday to Monday. Column (1) presents the baseline regression, column (2) contains option returns winsorized at percentiles 5 and 95, column (3) excludes the crisis months of 2018.02 and 2020.02 to 2020.04. Returns are in basis points and are delta-hedged. Standard errors are clustered within every month. The sample period is 1996 to 2023.

that is not Friday-to-Monday), and $Post_t$ indicates observations after January 2006. Standard errors are clustered by month.

Table II reports estimates from Equation 13. Column (1) presents the baseline results. The constant (-231.4 bps) shows that before 2006, average weekend option returns were substantially negative. The coefficient on Post (-264.5 bps) implies that weekend returns became even more negative after the emergence of overnight equity trading, consistent with a general increase in option risk premia during the post-period. In contrast, intra-week returns were moderately less negative before 2006 (-144.7 bps $= -231.4 + 86.7$) and turned nearly flat afterward ($-53.4bps = -231.4 + 86.7 - 264.5 + 355.8$). The interaction term Intra-week × Post (355.8 bps, t = 4.05) provides the key difference-in-differences estimate: option returns within the week became significantly less negative relative to weekend returns once overnight equity trading emerged.

Columns (2) and (3) show that this result is robust to winsorizing returns at the 5th and 95th percentiles and to excluding major crisis months (2018.02 and 2020.02–2020.04). Across specifications, intra-week returns remain consistently higher than weekend returns after 2006. These findings indicate that greater overnight equity liquidity mitigated the negative option risk premium, consistent with the interpretation that improved liquidity relaxes dealers' hedging frictions. When equity trading is limited, intermediaries face higher inventory risk because they cannot continuously adjust their delta-hedges; the increase in overnight liquidity alleviates this constraint.

The increase in the weekend risk premium over the post-period, visible in the second row of Table II, possibly reflects the timing of the 2008 Global Financial Crisis, which followed shortly after the onset of overnight equity trading. Risk premia across many asset classes were unusually compressed in the years preceding the crisis and increased sharply thereafter. The more negative weekend returns observed after 2008 therefore may capture a broader repricing of macroeconomic risk and heightened investor demand for crash protection.

III.D. Enhancing Identification: A Triple-Difference Test

To further isolate the liquidity channel and assess where the effect is most pronounced, I extend the analysis to a triple-difference specification that interacts the treatment with cross-sectional variation in option type. Specifically, I compare out-of-the-money puts, where dealer inventory risk is concentrated, to other option portfolios. This approach sharpens identification and reveals whether the liquidity effect operates primarily through the options that contribute most to intermediaries inventory risk.

Table III extends the baseline difference-in-differences framework of Table (II) by adding a third dimension of variation across option types. To that end, the table contains option returns for eight out-of-the-money portfolios (four puts and four calls) with breakpoints defined in Table A.1. While the regression in Table II compares Intra-week and weekend option returns before and after the emergence of overnight equity trading, the triple-difference specification exploits cross-sectional heterogeneity between out-of-the-money puts and other option portfolios. This additional layer of identification isolates the liquidity effect from broader time-series shifts in option risk premia by comparing treated and untreated options within the same periods. In doing so, the specification tests whether the impact of overnight equity liquidity is strongest precisely where hedging frictions are expected to bind most tightly and helps rule out confounding explanations driven by aggregate shocks or changes in overall market volatility.

Column (1) shows that the key triple-interaction term, $Intra-weekPostOTMputs$, is large and statistically significant (470.9 bps, $t = 2.93$). This coefficient indicates that the post-2006

Table III: Equity Liquidity and Option Returns: Triple-Difference Estimates

	(1) Baseline	(2) Winsorized	(3) Excl. Crashes
Intra-week × Post × OTM Puts	470.9***	421.3***	561.0***
	(2.93)	(4.38)	(3.97)
Intra-week × Post	121.9***	92.9***	174.2***
	(2.60)	(2.74)	(3.87)
Intra-week × OTM Puts	61.7	55.9	61.7
	(0.56)	(0.73)	(0.56)
Post × OTM Puts	-305.6**	-348.1***	-409.7***
	(-1.97)	(-4.13)	(-3.24)
Lower Order Terms	Y	Y	Y
Observations	54,842	54,842	53,675
R2-adjusted	0.00	0.01	0.01

Note: This table reports estimates from a triple-difference specification based on Equation (13), in which option returns are regressed on indicators for treated options (OTM puts), treated periods (Intra-week), and the post-event period. Option returns are computed for eight out-of-the-money portfolios (four puts and four calls) with breakpoints defined in Table A.1. OTM Puts indicates the portfolio of deep out of the money puts, with less than 70 days to expiry. Intra-week returns include all close-to-close observations, except Friday to Monday. Column (1) presents the baseline regression; Column (2) uses returns winsorized at the 5th and 95th percentiles; Column (3) excludes the crisis months 2018.02 and 2020.02–2020.04. Returns are expressed in basis points and are delta-hedged. Standard errors are clustered by $month \times portfolio\ ID$. The sample period is 1996 to 2023.

improvement in intra-week option returns documented in Table II is concentrated in out-of-the-money puts—the contracts that dominate dealers' short inventory positions and thus carry the highest hedging frictions. The positive estimate implies that overnight equity liquidity reduced the magnitude of negative returns specifically for those options most exposed to dealers' inventory risk. In contrast, the lower-order interactions ($Intra-week \times Post, Post \times OTM\ puts$) are smaller in magnitude, confirming that the observed effect is not a broad repricing across all options but rather a targeted easing of the liquidity constraint faced by intermediaries holding downside risk. Appendix Table A.2 displays lower order terms.

Columns (2) and (3) show that the triple-difference results are robust to outliers. Column (2) winsorizes option returns at the 5th and 95th percentiles, while Column (3) excludes the crisis months of 2018.02 and 2020.02–2020.04. Across both specifications, the coefficient on $Intra-week \times Post, Post \times OTM\ puts$ remains positive, large, and statistically significant, confirming that the attenuation of the negative option risk premium is not driven by outliers or by periods of exceptional market stress. The persistence of the effect across samples reinforces the interpretation that improved overnight equity liquidity eases dealers' hedging constraints and thereby reduces the compensation investors pay for bearing downside option risk.

Tables II and III report results based on portfolios of options rather than individual contracts. Aggregating options into portfolios mitigates measurement error arising from idiosyncratic pricing noise, thin trading, and bid–ask bounce that affect individual option quotes. Portfolio-level returns smooth out microstructure frictions and ensure that each observation represents a well-diversified position within a given moneyness–maturity bin, which improves comparability across periods and reduces the influence of outliers. This approach follows standard practice in the options literature (e.g., Coval and Shumway (2001); Bakshi, Charles, and Chen (1997)), where portfolios serve as more reliable estimators of systematic return patterns than individual option-level regressions.

A potential concern for identification is that neither "weekend" nor "overnight" periods are completely non-tradable. Pre-market equity trading begins at 04:00 on Monday mornings, and E-mini S&P 500 futures reopen as early as Sunday 18:00. Hence, the weekend control group is only partially untreated. I address this concern in several ways. First, the expansion of equity trading hours around 2006 was far more pronounced for weeknights than for weekends. A typical weeknight spans 17.5 hours (16:00 – 09:30). The introduction of post-market (until 18:00) and pre-market (from 04:00) trading reduced its non-tradable share by roughly 30 percent. By contrast, a weekend period spans about 65.5 hours (17.5 plus 48), for which the same extension lowers non-tradable hours by less than 10 percent. Consistent with this difference, Figure A.4 shows that weeknight trading volumes in futures and ETFs increased much more sharply than weekend volumes. Finally, Table A.4 presents estimates from a continuous-treatment specification that uses measured overnight liquidity instead of the 2006 breakpoint. The results confirm that greater overnight equity liquidity significantly compresses the option risk premium, particularly for deep out-of-the-money puts.

The appendix provides several additional robustness analyses. Figure A.1 shows no evidence of pre-event trends in the weekend or intra-week returns of out-of-the-money puts, supporting the parallel-trend assumption underlying the difference-in-differences design. Table A.3 reports a decline in option bid–ask spreads for intra-week relative to Friday-close observations around 2006, consistent with the notion that market makers' required compensation for providing liquidity fell once overnight equity trading improved. Table A.5 documents a similar pattern in VIX futures: intra-week returns become less negative relative to weekend returns after 2006. The evidence from VIX futures reinforces the main result in a liquid, transparent, and easily replicable return series, underscoring that the effects of improved overnight liquidity extend beyond the option market itself.

IV. Intermediaries' Exposure to Gap Risk

This section documents intermediaries' exposure to gap risk. I estimate option dealers option inventory and estimate that, given their inventory, large negative equity market returns would lead to significant inventory losses, if these negative equity market returns occur overnight, when reduced equity liquidity impedes the adjustment of delta-hedges.

Supporting this interpretation, I quantify dealers' liquidity demand for delta-hedge adjustments, and find that this measure exceeds typical overnight stock or futures trade volumes. In addition, I show that dealers inventory gamma is on average positive over my sample, but their shadow gamma is large and negative, highlighting that conventional measures do not capture intermediaries exposure to gap risk. Finally, I provide evidence that intermediaries actively manage their option positions in response to changing equity market volatility, consistent with behavior aimed at mitigating this equity market risk.

IV.A. Intermediaries' Option Inventory and Short Put Exposure

In addition to the papers already mentioned, Stoikov and Sağlam (2009) show in a theoretical inventory-risk model that option market makers must be compensated for residual gamma and vega risk when they cannot hedge continuously. I complement their analysis by empirically quantifying gap risk in dealers' aggregate S&P 500 option positions and linking this gap risk to the empirical option risk premium. Fournier and Jacobs (2015) develop a structural dynamic model of an index option market maker with limited capital and estimates the model on index options plus inventory data. Hitzemann, Hofmann, Uhrig-Homburg, and Wagner (2023) provide evidence that the option risk premium is related to the capital constraints of option dealers. This paper, proposes a new type of inventory risk, which helps explain key moments of the option risk premium.

Data: Option Trade Volume by Trader Type. I use the CBOE Open–Close Volume files, which report daily option volumes by contract (put or call, strike, and maturity), trader group (market maker, broker-dealer, firm, customer, and professional customer), and trade direction (buy versus sell). Throughout the paper, I use the terms market makers and intermediaries interchangeably. The analysis focuses on the period 2011–2023, during which the data explicitly identify market-maker trades. Before 2011, market makers are not separately classified and must be inferred as the counterparty to firms and customers; restricting attention to the post-2011 sample also benefits from markedly higher option liquidity, yielding more reliable estimates of intermediaries' trading patterns at high frequency.

S&P 500 index options have a contract multiplier of 100, that is, one contract represents 100 units of the underlying index. To facilitate interpretation, I therefore express all option volumes and positions in underlying-equivalent units by multiplying by 100.

Examining market-maker trading activity, I find that intermediaries buy an average of 29.1 million S&P 500 put options per day and sell roughly 30.3 million, implying net purchases of −1.2 million puts daily. Net buys are defined as:

$$\text{Net Buys}_t^i = \text{Buys}_t^i - \text{Sells}_t^i, \tag{14}$$

where i indexes the option contract and t the trading day. In contrast, buy and sell volumes in call options are nearly identical at about 17.5 million contracts per day, resulting in net purchases close to zero. These patterns reveal that the intermediary sector persistently supplies put options to end-users while maintaining balanced positions in calls, foreshadowing the short-put exposure documented below.

I construct intermediaries' cumulative option position by aggregating their daily net buys over the life of each contract:

$$\text{Net Position}_t^i = \begin{cases} \sum_{k=1}^{t} \text{Net Buys}_k^i & \text{if } t \leq \text{Expiry} \\ 0 & \text{if } t > \text{Expiry} \end{cases} \tag{15}$$

Here, k indexes trading days from the start of the sample through day t. Thus, NetPosition_t^i measures the number of contracts of option i that intermediaries are long minus the number they are short at the end of day t. Because options are continuously listed and expire on a rolling basis, this cumulation produces a time series of intermediaries' outstanding inventory after a burn in period. The frequent expiration of derivative contracts is a major advantage when studying intermediary inventory risk. In contrast to equities, where precise measurement of the initial

Table IV: Intermediaries' Net Position in S&P 500 Options

Puts		Days to Expiry 2-70	71-	All
$0.00 \leq \lvert\Delta\rvert \leq 0.25$	Deep Out of the Money	-13.54	-3.32	-16.86
$0.25 < \lvert\Delta\rvert \leq 0.50$	Out of the Money	-0.47	-1.47	-1.94
$0.50 < \lvert\Delta\rvert \leq 0.75$	In the Money	0.38	-0.07	0.32
$0.75 < \lvert\Delta\rvert \leq 1.00$	Deep In the Money	0.47	0.17	0.63
All		-13.15	-4.70	-17.85
Calls				
$0.00 \leq \lvert\Delta\rvert \leq 0.25$	Deep Out of the Money	-0.84	-0.50	-1.33
$0.25 < \lvert\Delta\rvert \leq 0.50$	Out of the Money	1.55	0.87	2.42
$0.50 < \lvert\Delta\rvert \leq 0.75$	In the Money	1.40	0.47	1.87
$0.75 < \lvert\Delta\rvert \leq 1.00$	Deep In the Money	0.61	-0.01	0.59
All		2.72	0.83	3.55

Note: This table reports the average net positions of S&P 500 option market makers, sorted by moneyness and days to expiry. The net position is the number of contracts that intermediaries are long minus the number they are short. Negative values indicate that intermediaries are, on average, net short in the respective option category. Positions are expressed in millions of contracts. Days to expiry are calendar days. The sample period spans 2011 to 2023.

position is crucial and often unobservable, options naturally reset at expiry, eliminating concerns about initial conditions or position drift. I adopt a six-month burn-in period, resulting in a main sample that spans July 2011 to August 2023. Positions are reset to zero at contract expiry.

Table IV reports intermediaries' average net positions in S&P 500 options across portfolios sorted by moneyness and days to expiry. Options are assigned to portfolios at the daily market close and reclassified at the next close. The intermediary sector holds an aggregate short position of approximately −17.9 million put contracts, concentrated in short-maturity, deep out-of-the-money puts, where the net position reaches −13.5 million. Negative values indicate that intermediaries are, on average, net short in the respective option category. In contrast, market makers' net positions in call options are slightly positive and display no consistent pattern across portfolios. This asymmetry is consistent with persistent end-user demand for deep out-of-the-money puts as crash protection, leaving intermediaries as the residual sellers of such options. The next subsection examines how this short put exposure translates into equity market risk for intermediaries.

The literature attributes intermediaries' persistent short positions in S&P 500 put options largely to the hedging demands of sophisticated investors. Lemmon and Ni (2014) link index option trading to institutional hedging motives, noting that trading in single-stock options is more often driven by retail investors. Bollen and Whaley (2004) show that institutional investors hold long positions in index puts as portfolio insurance, and Chen, Joslin, and Ni (2019) interpret customer demand for puts as evidence of investors' aversion to economic crash risk. Similarly, Goyenko and Zhang (2019) document net buying pressure in S&P 500 put options but net selling pressure in calls. Taken together, this evidence indicates that institutional investors are the primary buyers of index puts, motivated by hedging or insurance considerations, and that there is no natural counterparty to supply these contracts. Intermediaries therefore fill this role, resulting in a persistent short position in out-of-the-money puts.

IV.B. Intermediaries' Inventory Exposure to Gap Risk

The short-put exposure documented above implies that intermediaries are exposed to large declines in the underlying equity index. This subsection quantifies that exposure and illustrates the mechanics through which limited overnight equity liquidity transforms intermediaries' option positions into equity market risk. When the underlying price falls sharply, delta-hedged positions must be rebalanced through substantial equity sales; when trading is illiquid, such adjustments are infeasible, leaving dealers exposed to losses. In the next step, I quantify this risk using intermediaries' actual option positions to estimate their aggregate profit-and-loss under different hypothetical equity return scenarios.

I quantify intermediaries' exposure to equity market risk by estimating the profit and loss (P&L) of their aggregate option portfolio under hypothetical equity return scenarios. I estimate dealers' exposure to gap risk via their hypothetical P&L for different potential returns of the S&P 500.

Specifically, I compute the Dealer P&L as:

$$\text{Dealer P\&L}_{t+1}(\Delta S) = \sum_{i=1}^{I} \text{Net Position}_t^i \times \left[P_{t+1}^i(\Delta S) - P_t^i \; - \; \Delta_t^i \; \times \; (\text{SP500}_{t+1}(\Delta S) - \text{SP500}_t)\right], \tag{16}$$

where $\text{P\&L}_{t+1}(\Delta S)$ is the estimated P&L over period $t+1$; Net Position_t^i denotes intermediaries' net position in option i at the end of day t; $\text{P}_{t+1}^i(\Delta S)$ is the estimated option price under the simulated market scenario, Δ_t^i is the option's delta, and $\text{SP500}_{t+1}(\Delta S)$ is the assumed value of the S&P 500.

This framework estimates the losses the intermediary sector would incur under different joint realizations of option prices and underlying index levels. The method is flexible and can accommodate various dynamics for both P and $SP500$. To map future option prices to underlying values, I adopt the Black–Scholes–Merton model for simplicity. The choice of pricing model is immaterial for the analysis, as long as option values remain convex in the underlying price, such that large underlying moves induce nonlinear changes in delta. I assume that intermediaries are initially fully delta-hedged but subsequently hold their positions without adjustment, capturing their risk exposure over nights and weekends when equity market liquidity is low.

To evaluate Scenario $\text{P\&L}_{t+1}$ in Equation (16) for a hypothetical -10% return in the underlying S&P 500, I compute the option value implied by the scenario $(P^i_{t+1}(\Delta S))$ as the Black-Scholes-Merton price of option i at the end of period $t+1$, assuming that $\sigma^i_{t+1} = \sigma^i_t$ and the assumed S&P 500 level under the scenario $\text{SP500}_{t+1}(\Delta S) = \text{SP500}_t \times 0.9$, where σ describes options' implied volatility relative to the Black-Scholes-Merton model. That is, I change the equity price, under the Black-Scholes-Merton pricing framework, while holding implied volatilities constant. In practice, equity returns and volatility are negatively correlated, and increases in volatility raise option prices. Consequently, the losses reported here likely understate intermediaries' true downside exposure, as incorporating volatility spikes would make losses during market crashes even more pronounced.

The use of option dealer scenario P&L was pioneered by Gruenthaler (2022), who constructs an option intermediary risk factor (OIR) from the slope of the scenario P&L around a -10% crash, and then shows that OIR predicts the option risk premium. My approach differs in several ways: For one, I keep option implied volatilities constant across scenarios, to focus specifically on dealer risk from inability to adjust delta-hedges. Second, I focus on the *curvature* of the scenario P&L, that is, I focus on the new concept of shadow gamma. Third, I show that shadow gamma is priced at night, and especially so when dealers are constrained.

Figure 4 shows intermediaries' exposure to gap risk. The figure shows the intermediary scenario P&L for different hypothetical S&P 500 returns, in % of the upper bound on dealers' option portfolio value. The figure shows that a -10% return in the S&P 500 index would lead to a portfolio return of about -2%. This number is highly significantly different from zero, but represents a lower bound on dealers' gap risk exposure, since the denominator is an upper abound on dealers' option portfolio value. I calculate dealers' daily option portfolio value as the sum-

Figure 4: Dealer P&L Across Hypothetical S&P 500 Returns

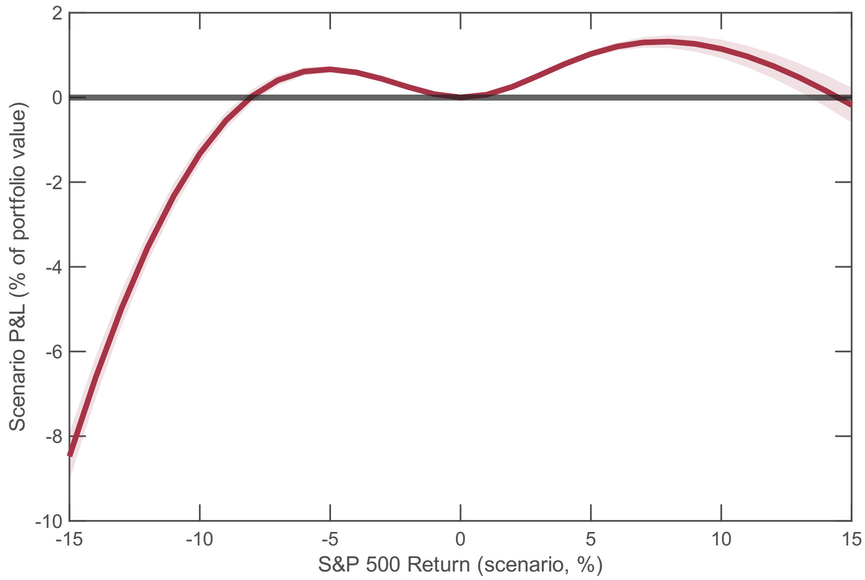

Note: The figure shows estimated profits of market makers' option inventory for different hypothetical returns of the underlying S&P 500 equity index. Profits are in percent of the upper bound of the dollar value of dealers' option inventory. Option returns are initially delta-hedged, but the hedges are not adjusted as the underlying price changes. The shaded area shows the 95% confidence interval around the mean estimate. The sample period is 2011 to 2023.

product of the *absolute value* of dealers' net-position across options and the respective options' mid-quote. Thus, I assume (i) netting of longs and shorts within the same option, (ii) no netting of longs and shorts for very similar options, (iii) dealers' margin requirement is the option value for both their longs and shorts. These assumptions are very conservative, since dealers' margin requirements are significantly lower especially for high-price long-maturity in-the-money options (Hitzemann, Hofmann, Uhrig-Homburg, and Wagner, 2023). As a result, the estimated percent values are a lower bound on dealers' gap risk exposure. The estimation does not account for the empirically large option risk premium, and thus does not display the intermediary profits that materialize when absolute equity returns are small.

IV.C. Intermediaries' Liquidity Demand for Hedge Adjustments

Periods of limited equity liquidity expose intermediaries to equity market risk because they constrain their ability to maintain delta-neutral portfolios. This subsection quantifies the volume of equity trading required for intermediaries to remain delta-hedged when the underlying index experiences large price movements. By estimating the liquidity demand implied by intermediaries'

Figure 5: Dealers' Liquidity Demand for Hedge Adjustments

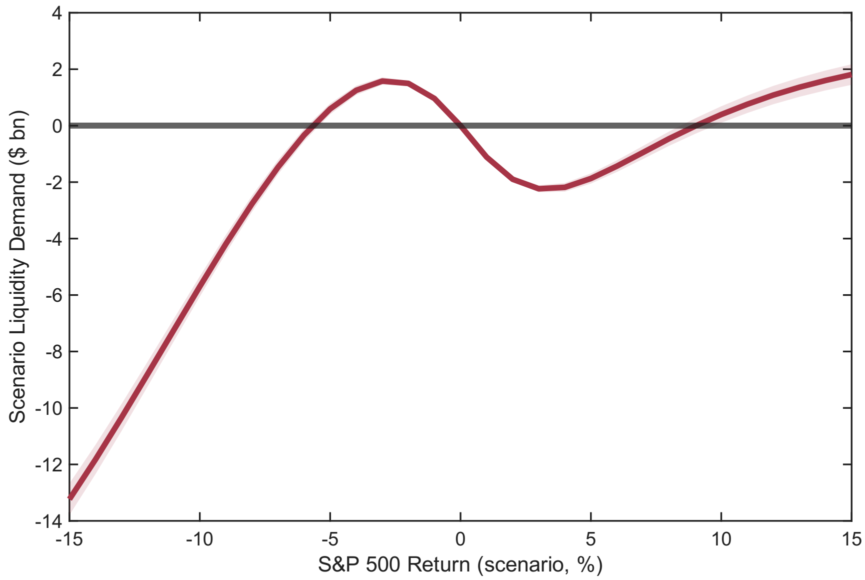

Note: This figure plots intermediaries' estimated liquidity demand required to maintain delta-hedged positions under hypothetical S&P 500 return scenarios, as defined in Equation (17). Liquidity demand is measured in billion dollars and represents the trading volume needed to offset changes in portfolio delta. The shaded area shows the 95% confidence interval around the mean estimate. The sample period is 2011 to 2023.

option inventory under different equity return scenarios, I show that the scale of these required hedge adjustments far exceeds actual overnight trading volumes for most hours of the night, highlighting the practical limits of continuous hedging.

I estimate intermediaries' liquidity demand as:

$$\text{Liquidity Demand}_{t+1}(\Delta S) = \sum_{i=1}^{I} NetPosition_t^i \times \left[\Delta_{t+1}^i - \Delta_t^i(\Delta S)\right] \times \text{SP500}_t, \tag{17}$$

where Liquidity $\text{Demand}_{t+1}(\Delta S)$ denotes the estimated dollar trading volume required for intermediaries to maintain delta-neutral positions following a hypothetical change in the underlying. $NetPosition_t^i$ is intermediaries' net position in option i at the end of day t as defined earlier. $\Delta_t^i(\Delta S)$ is options' Black-Scholes-Merton delta, calculated using the option's implied volatility σ_t^i. Δ_{t+1}^i is the estimated delta under the simulated market scenario, where $\sigma_{t+1}^i = \sigma_t^i$ and $\text{SP500}_{t+1} = \text{SP500}_t \times \mu$. The parameter μ takes values of $[0.85, 0.9, ..., 1.15]$ to simulate S&P 500 returns of $[-15\%, -10\%, ..., 15\%]$.

Intermediaries' liquidity demand for maintaining delta-hedges amounts to several billions of dollars. Figure 5 plots the estimated liquidity demand from Equation (17) across hypothetical S&P 500 return scenarios. If the underlying index declines by 10%, intermediaries would need to sell approximately \$8 billion worth of equities to remain delta-hedged. This required trading volume far exceeds typical overnight activity in equities and futures between 20:00 and 04:00, when markets are thinly traded. Consequently, intermediaries face severe constraints on hedge adjustments during these periods, exposing them to substantial inventory risk.

These estimates likely represent a lower bound. Investors' strong demand for crash protection implies that intermediaries hold short positions in options beyond the S&P 500 contracts examined here. Moreover, option deltas rise with expected volatility, which typically spikes during large negative equity market moves, further amplifying intermediaries' liquidity needs precisely when market depth is at its lowest.

IV.D. Intermediaries' Shadow Gamma

Intermediaries' inventory risk is driven by options' shadow gamma—the state-contingent curvature that emerges after large underlying moves—rather than by contemporaneous gamma. Gamma captures the local curvature of option values and thus the extent to which delta hedges become imperfect for small price changes. Over 2011–2023, however, intermediaries' aggregate gamma hovers around zero and, if anything, is slightly positive, implying gains from underlying price changes, not losses. This measure misses the core exposure because dealers' dominant position is short deep out-of-the-money puts: these contracts suffer large losses in a crash yet contribute little to current gamma, since deep-OTM gamma is near zero. Evaluated under large price moves, by contrast, intermediaries' inventories load on substantial negative shadow gamma. Positions that are well balanced for small fluctuations become acutely exposed when the underlying moves by a large amount, generating equity-market risk and inventory losses, especially overnight when hedge adjustments are constrained.

I estimate intermediaries' inventory gamma for different scenarios of underlying S&P 500 returns as:

$$\text{Dealer Shadow Gamma}_{t+1}(\Delta S) = \sum_{i=1}^{I} \text{Net Position}_{i,t} \times \widetilde{\Gamma}^{\text{sh}}_{i,t}(\Delta S), \tag{18}$$

where $NetPosition_t^i$ denotes intermediaries' net holdings of option i at time t, and $\widetilde{\Gamma}^{\text{sh}}_{i,t}(\Delta S)$ is the option's shadow gamma evaluated under a hypothetical equity-market return (ΔS). I compute shadow gamma across a grid of counterfactual S&P 500 returns from -15% to $+15\%$, which captures the curvature of intermediaries' portfolio value under large price moves rather than infinitesimal changes. This construction links measured portfolio risk directly to the nonlinear exposure dealers face when equity liquidity constraints impede continuous re-hedging.

Figure 6 plots intermediaries' inventory gamma across hypothetical S&P 500 returns. The gamma at a return of zero corresponds to intermediaries' current exposure and is slightly positive, implying that small equity moves generate gains rather than losses. This measure therefore fails to capture intermediaries' inventory risk over nights and weekends, when large price changes are most relevant. For both positive and negative equity moves, the estimated shadow gamma turns negative, reflecting the expected losses on intermediaries' short positions in deep out-of-the-money puts following a market decline. Figure A.3 contains a binscatter of intermediary market risk against intermediary shadow gamma and shows a strong positive relation between the two.

At-the-money gamma remains positive, indicating that dealers profit from small price fluctuations while bearing risk from large ones. This pattern is consistent with Hu, Kirilova, and Muravyev (2023), who find limited delta-hedging by option intermediaries in Korean markets. My results suggest that intermediaries' portfolios are locally well balanced against small underlying moves but become increasingly exposed as price changes grow in magnitude.

Option dealers inventory gamma still matters. For example Soebhag (2023) and Baltussen, Da, Lammers, and Martens (2021) conclude that option dealers inventory gamma captures their need to trade for delta-hedges and create price pressure in the underlying equities.

My approach to measuring intermediaries' exposure to equity market risk differs fundamentally from existing work. Since Garleanu, Pedersen, and Poteshman (2009), the literature has typically estimated intermediary risk using net gamma and net vega, which proxy for exposure to jump and volatility risk, respectively. Net gamma, the sum product of intermediaries' net option positions and each contract's gamma, captures how much dealers' delta-hedges must adjust in response to small price movements of the underlying. Gamma, the second derivative of the option value with respect to the underlying price, thus measures the local curvature of intermediaries' portfolio value around current market conditions. This local measure is informative when jumps are small and trading is continuous, since in such settings the underlying rarely departs far from its current level.

Figure 6: Dealer Gamma Across Hypothetical S&P 500 Moves

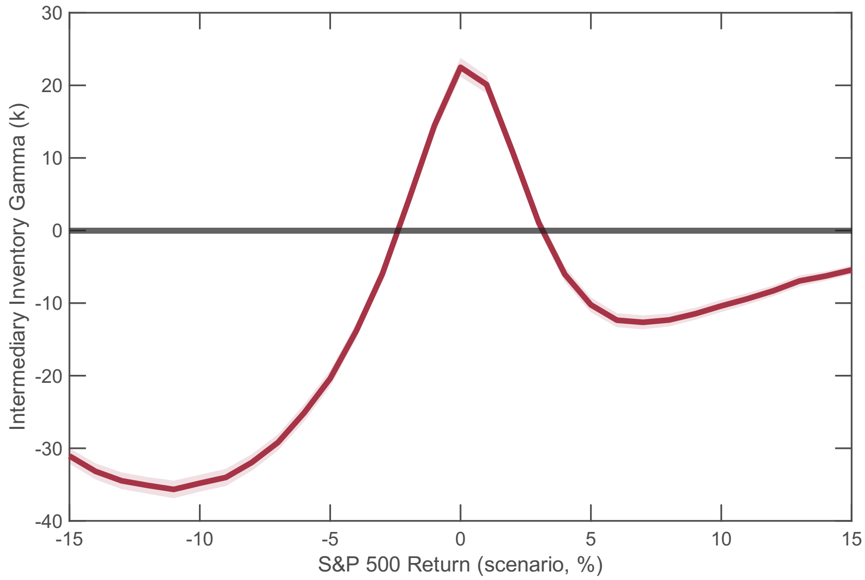

Note: This figure plots intermediaries' aggregate inventory gamma for the actual and several hypothetical values of the S&P 500 index. Inventory gamma is computed as the sum product of intermediaries' net positions across option contracts and the respective options' gamma, defined as the second derivative of the option value with respect to the underlying price. Negative gamma indicates expected inventory losses under large equity moves. The shaded area shows the 95% confidence interval around the mean estimate. The sample period is 2011 to 2023.

However, this framework cannot account for the pronounced premia in out-of-the-money puts relative to calls. Because puts and calls with identical contract specifications share the same gamma, net-gamma–based measures would imply symmetric risk premia, contrary to the data. Moreover, gamma is highest for at-the-money options, whereas observed risk premia are largest for short-maturity, deep out-of-the-money puts. If conventional gamma captured intermediaries' primary exposure, the most expensive contracts should be at-the-money, which the evidence rejects.

In contrast, I show that intermediaries' risk arises from exposure to large, infrequent price movements, rather than small price fluctuations. This risk emerges from the daily market close of equity exchanges and the resulting overnight illiquidity, which prevents continuous delta-hedging. Consequently, out-of-the-money puts are especially risky for intermediaries, and their elevated risk premia reflect compensation for this shadow-gamma exposure.

Figure 7: Intermediary Put Trading By Lagged Equity Return Volatility

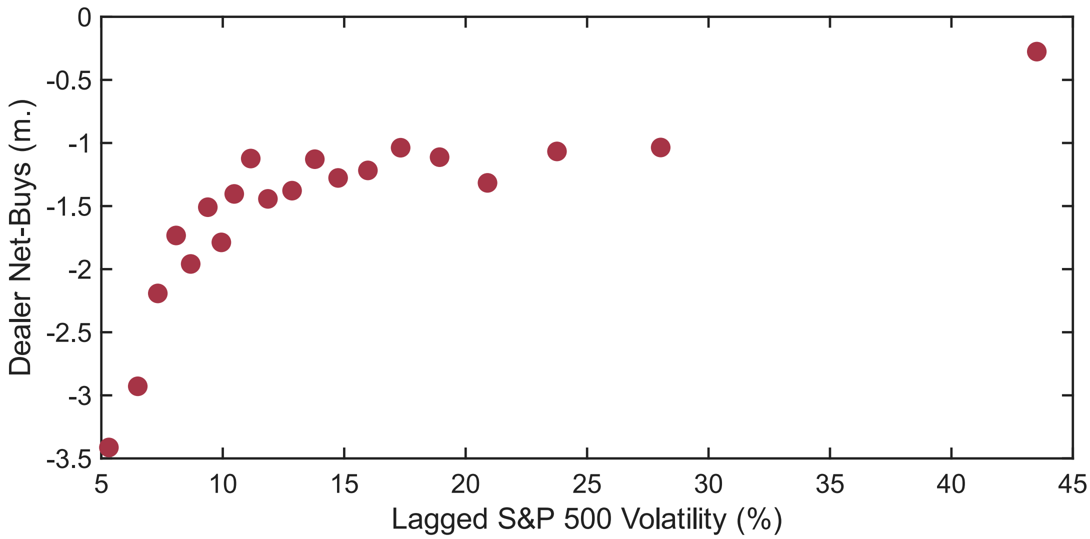

Note: This figure plots intermediaries' net purchases of S&P 500 put options against lagged equity market volatility. Net purchases are measured as daily buys minus sells, and are expressed in millions of contracts. Volatility is computed as the standard deviation of close-to-close S&P 500 returns over a rolling 10-day window, lagged by one day. The downward slope indicates that intermediaries sell more puts when recent volatility is low, consistent with active inventory risk management. The sample period is 2011 to 2023.

IV.E. Intermediaries' Inventory Adjustment and Risk Management

Intermediaries adjust their option inventories systematically in response to changing market conditions. They sell more put options when equity market volatility is low, indicating that (*i*) they actively manage exposures rather than passively absorb investor demand, and (*ii*) equity market risk is a central component of their overall inventory risk. Gruenthaler (2022) documents that option intermediaries manage their net-vega positions in anticipation of spikes in implied volatility. I provide complementary evidence that intermediaries manage their exposure to equity market risk. This behavior reinforces the interpretation that option risk premia arise primarily overnight, when limited equity liquidity constrains delta-hedging and leaves intermediaries exposed to unhedgeable inventory risk.

Intermediaries' trading behavior provides evidence of active risk management in response to changing equity market conditions. Figure 7 plots intermediaries' daily net purchases of S&P 500 puts against lagged realized equity volatility. The relationship is strongly negative: when volatility is low at around 5% intermediaries sell roughly 3.4 million puts per day; when volatility is high at around 17% daily sales decline to about 1 million contracts. The pattern is nearly monotonic and appears concave, suggesting diminishing sensitivity at higher volatility levels. Although the figure does not separately identify shifts in customer demand, it is more plausible that dealers scale back put sales when volatility rises than that customers reduce demand when markets are calm. This trading pattern supports the interpretation that intermediaries manage their short-put exposure countercyclically, increasing risk-taking when equity market risk is low and retrenching when risk is high. This behavior is consistent with inventory management motives rather than passive order flow accommodation.

V. The Price of Gap Risk

This section examines whether gap risk is priced at night, especially when dealers are constrained. I regress day and night option returns on lagged shadow gamma and find strong predictive power only for overnight returns. This pattern is consistent with shadow gamma capturing the risk from the inability to adjust delta hedges in illiquid night markets. I further show that the explanatory power of shadow gamma is stronger when intermediaries are more constrained, in the sense that they hold larger shadow-gamma inventories.

To test whether shadow gamma is priced in options, I estimate the following specification:

$$R_t^i = \beta_0 \, \mathbf{1} + \beta_1 \, \widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%) + \beta_2 \times \text{Night}_t + \beta_3 \, \widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%) \times \text{Night}_t + \epsilon_t^i, \tag{19}$$

where $R_{i,t}$ denotes the delta-hedged return of option i, in basis points. $\widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%)$ is the lagged shadow gamma of option i evaluated for hypothetical S&P 500 returns of -10%, and $Night$ indicates overnight returns, as opposed to intraday returns. Shadow Gamma is standardized to zero mean and unit variance. Standard errors are double-clustered by trading day and option identifier.

Table V shows regression estimates of Equation (19). Column (1) displays the baseline specification. Most importantly, options' shadow gamma predicts overnight option returns. A one standard deviation increase in shadow gamma is associated with a relative reduction of overnight option returns of 177.4bps. More shadow gamma indicates a higher curvature of the option value function in the scenario of large negative equity returns, thus making an option more risky for sellers, which tend to be intermediaries. Overnight option returns are large and negative (-266bps on average). In contrast, intraday option returns are indistinguishable from zero, and are not predicted by options' shadow gamma. Bali, Goyal, Moerke, and Weigert (2025) find momentum and

Table V: The Price of Shadow Gamma

	(1) Baseline	(2) Winsorized	(3) Excl. Crashes
Shadow Gamma	29.1	26.5	19.9
	(1.26)	(1.19)	(1.06)
Night	-266.1***	-265.8***	-233.2***
	(-4.71)	(-4.93)	(-4.89)
Shadow Gamma × Night	-177.4***	-180.4***	-169.7***
	(-5.46)	(-6.16)	(-6.70)
Constant	43.7	34.9	-6.4
	(0.92)	(0.75)	(-0.17)
Observations	3,259,955	3,259,955	3,125,142
R2-adjusted	0.00	0.00	0.00

Note: This table presents regression estimates of Equation (19) which regresses day and night delta-hedged option returns on lagged option-level shadow gamma. Option shadow gamma is for the scenario of a -10% S&P 500 return. Option returns are measured in basis points. Shadow gamma is standardized have zero mean and unit variance. Column (1) presents the baseline regression; Column (2) uses returns winsorized at percentiles 5 and 95; Column (3) excludes the crisis months 2018.02 and 2020.02–2020.04. $t-$stats are in round brackets, and standard errors are double-clustered by trading day and option identifier. The sample period is 2011 to 2023.

reversal in day-night option returns. Nonetheless, I find predictability only for night returns, without a subsequent reversal prediction. Columns (2) and (3) provide show robustness of the findings against outliers in option returns. These findings are consistent with an interpretation where the option risk premium predominantly derives from unhedgeable risk due to equity illiquidity.

Finally, I turn to the role of intermediaries' option inventory. Table V examines whether an option's own shadow gamma predicts its returns and allows the price of shadow gamma to differ between day and night, but does not consider dealers' aggregate position. Inventory-risk models instead imply that the premium per unit of shadow gamma should be higher when intermediaries are more exposed to gap risk. To test this prediction, I now focus on overnight returns and relate them to the interaction between option-level shadow gamma and dealers' aggregate shadow gamma.

To test whether intermediary constraints impact the option risk premium, I estimate the following specification:

$$R_t^i = \beta_0 \, \mathbf{1} + \beta_1 \, \widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%) \; + \beta_2 \times \text{Dealer Shadow Gamma}_t(-10\%) + \\ \beta_3 \, \widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%) \times \text{Dealer Shadow Gamma}_t(-10\%) \; + \epsilon_t^i, \quad (20)$$

Table VI: Dealer Inventory and the Price of Shadow Gamma

	(1) Baseline	(2) Winsorized	(3) Excl. Crashes
Shadow Gamma	-155.4***	-155.5***	-155.4***
	(-7.98)	(-8.09)	(-7.98)
Dealer Shadow Gamma	30.5	29.9	30.5
	(0.86)	(0.85)	(0.86)
Shadow Gamma	35.8**	35.7**	35.8**
× Dealer Shadow Gamma	(2.44)	(2.46)	(2.44)
Constant	-216.3***	-217.1***	-216.3***
	(-6.04)	(-6.15)	(-6.04)
Observations	958,387	958,387	958,387
R2-adjusted	0.00	0.00	0.00

Note: This table presents regression estimates of Equation (20), which regresses overnight, delta-hedged option returns on lagged option-level shadow gamma and dealers' aggregate shadow gamma. Option shadow gamma and dealer shadow gamma are for the scenario of a -10% S&P 500 return. Option returns are measured in basis points and standardized explanatory variables have zero mean and unit variance. Column (1) presents the baseline regression; Column (2) uses returns winsorized at the 5th and 95th percentiles; Column (3) excludes the crisis months 2018.02 and 2020.02–2020.04. $t-$stats are in round brackets, and standard errors are double-clustered by trading day and option identifier. The sample period is 2011 to 2023.

where $R_{i,t}$ denotes the overnight delta-hedged return of option i, $\widetilde{\Gamma}^{\text{sh}}_{t-1}(-10\%)$ is the lagged shadow gamma of option i evaluated for hypothetical S&P 500 returns of -10%. Dealer Shadow Gamma measures intermediaries' inventory shadow gamma at -10%. Shadow Gamma and Dealer Shadow GAmma are standardized to zero mean and unit variance. Standard errors are double-clustered by trading day and option identifier.

Table VI shows regression estimates of Equation (20). Column (1) contains the baseline estimates. Most importantly, shadow gamma has more predictive power for overnight option returns when intermediaries are more constrained. Specifically, an increase in dealer shadow gamma by one standard deviation is associated with a shadow gamma predictive power that changes from -155 to -129 bps per standard deviation. This is reasonable, since dealer shadow gamma is large and negative, indicating their inventory risk. Higher dealer shadow gamma indicates less risk and should reduce the predictive power of shadow gamma on its own.

Dealer shadow gamma by itself does not significantly predict overnight option returns. This is not evidence against an inventory-risk channel. Dealer inventories are equilibrium objects that respond both to shocks to end-user demand for crash insurance and to shocks to dealers' funding and risk-bearing capacity. A positive demand shock raises option prices and expands dealers' short positions, whereas a tightening of funding conditions induces dealers to cut positions even as required risk premia rise. As a result, the unconditional time-series correlation between the level of dealer shadow gamma and option returns is theoretically ambiguous, and the insignificant coefficient on dealer shadow gamma in Equation (20) is entirely consistent with the model.

Instead of interpreting the level of dealer shadow gamma, I focus on how the price of gap risk varies with dealers' aggregate exposure. The interaction between option-level and dealer shadow gamma in Equation (20) asks whether options that add more shadow gamma to dealers' books earn especially negative overnight returns precisely when dealers' aggregate shadow-gamma exposure is high. Because this test exploits cross-sectional differences in option shadow gamma at a given point in time, it is much less sensitive to the endogeneity of inventory levels driven by aggregate shocks that move all option prices together. Consistent with the inventory-risk interpretation, the evidence in subsection IV shows that intermediaries actively reduce their positions when market risk increases, so that high and low inventory states can be associated with both high and low risk premia. Taken together, these results indicate that it is the state-dependent price of gap risk, rather than the unconditional level of dealer inventory, that links intermediary constraints to the option risk premium.

VI. Conclusion

This paper shows that intermediary inventory risk affects the option risk premium. I exploit the emergence of overnight equity trading around 2006 to show that the overnight concentration of the option risk premium is potentially explained by reduced equity liquidity. Equity il-liquidity raises the magnitude of the option risk premium because option sellers require liquid equity markets to adjust their delta-hedges in case or large underlying returns. A new derivatives risk measure "shadow gamma" captures this risk. Shadow gamma predicts overnight option returns, especially when dealers are constrained.

The results have two broader implications. First, from a market-design perspective, the evidence suggests that increasing around-the-clock equity trading can lower hedging costs and compress derivative risk premia by improving intermediaries' ability to manage inventory risk. Second, from an asset-pricing perspective, the findings highlight that liquidity frictions and intermediary constraints jointly determine the price of risk. Derivative markets therefore provide a clean setting in which to observe how institutional frictions give rise to persistent and economically large risk premia.

A natural direction for future research is to develop methods that explicitly control for intermediary inventory and market frictions when extracting information from option prices, for example for tests of macro models as in Backus, Chernov, and Martin (2011).

References

Andersen, Torben, Ilya Archakov, Leon Grund, Nikolaus Hautsch, Yifan Li, Sergey Nasekin, Ingmar Nolte, Manh Cuong Pham, Stephen Taylor, and Viktor Todorov, 2021, A descriptive study of high-frequency trade and quote option data, *Journal of Financial Econometrics* 19, 128–177.

Andersen, Torben G, Nicola Fusari, and Viktor Todorov, 2015, The risk premia embedded in index options, *Journal of Financial Economics* 117, 558–584.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *The Journal of Finance* 66, 1969–2012.

Bakshi, Gurdip, Cao Charles, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003–2049.

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

Bali, Turan G, Amit Goyal, Mathis Moerke, and Florian Weigert, 2025, In search of seasonality in intraday and overnight option returns, *Available at SSRN 5386128*.

Baltussen, Guido, Zhi Da, Sten Lammers, and Martin Martens, 2021, Hedging demand and market intraday momentum, *Journal of Financial Economics* 142, 377–403.

Bates, David S., 2022, Empirical option pricing models, *Annual Review of Financial Economics* 14, 369–389.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Bollen, Nicolas P.B., and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance* 59, 711–753.

Boyarchenko, Nina, Lars C Larsen, and Paul Whelan, 2023, The overnight drift, *The Review of Financial Studies*.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding index option returns, *Review of Financial Studies* 22, 4493–4529.

Cao, Jie, and Bing Han, 2013, Cross section of option returns and idiosyncratic stock volatility, *Journal of Financial Economics* 108, 231–249.

Chen, Hui, Scott Joslin, and Sophie Xiaoya Ni, 2019, Demand for crash insurance, intermediary constraints, and risk premia in financial markets, *The Review of Financial Studies* 32, 228–265.

Christoffersen, Peter, Bruno Feunou, Yoontae Jeon, and Chayawat Ornthanalai, 2021, Time-varying crash risk embedded in index options: The role of stock market liquidity, *Review of Finance* 25, 1261–1298.

Coval, Joshua D., and Tyler Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1009.

Dew-Becker, Ian, and Stefano Giglio, 2023, Risk preferences implied by synthetic options, *NBER Working Paper*.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2018, Deviations from covered interest rate parity, *The Journal of Finance* 73, 915–957.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Fournier, Mathieu, and Kris Jacobs, 2015, Inventory risk, market-maker wealth, and the variance risk premium: Theory and evidence, *Rotman School of Management Working Paper* 2334842.

Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman, 2009, Demand-based option pricing, *The Review of Financial Studies* 22, 4259–4299.

Goyenko, Ruslan, and Chengyu Zhang, 2019, Demand pressures and option returns, *Working Paper*.

Gruenthaler, Thomas, 2022, Risk premia and option intermediation, *Working Paper*.

Haddad, Valentin, and Tyler Muir, 2021, Do intermediaries matter for aggregate asset prices?, *The Journal of Finance* 76, 2719–2761.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Hitzemann, Steffen, Michael Hofmann, Marliese Uhrig-Homburg, and Christian Wagner, 2023, Margin requirements and equity option returns, *Working Paper*.

Hu, Jianfeng, Antonia Kirilova, and Dimitry Muravyev, 2023, Option market makers, *Working Paper*.

Jones, Christopher S., and Joshua Shemesh, 2018, Option mispricing around nontrading periods, *Journal of Finance* 73, 861–900.

Kanne, Stefan, Olaf Korn, and Marliese Uhrig-Homburg, 2023, Stock illiquidity and option returns, *Journal of Financial Markets* 63, 100765.

Lemmon, Michael, and Sophie Xiaoyan Ni, 2014, Differences in trading and pricing between stock and index options, *Management Science* 60, 1985–2001.

Merton, Robert C., 1973, Theory of rational option pricing, *Bell J Econ Manage Sci* 4, 141–183.

Merton, Robert C, 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Muravyev, Dmitriy, 2016, Order flow and expected option returns, *Journal of Finance* 71, 673–708.

——— , and Xuechuan (Charles) Ni, 2020, Why do option returns change sign from day to night?, *Journal of Financial Economics* 136, 219–238.

Ni, Sophie X, Neil D Pearson, Allen M Poteshman, and Joshua White, 2021, Does option trading have a pervasive impact on underlying stock prices?, *The Review of Financial Studies* 34, 1952–1986.

Soebhag, Amar, 2023, Option gamma and stock returns, *Journal of Empirical Finance* 74, 101442.

Stoikov, Sasha, and Mehmet Sağlam, 2009, Option market making under inventory risk, *Review of Derivatives Research* 12, 55–79.

Tankov, Peter, 2008, Pricing and hedging gap risk, .

A.1. Appendix Figures

Figure A.1: Mean Option Returns by 2-year Period

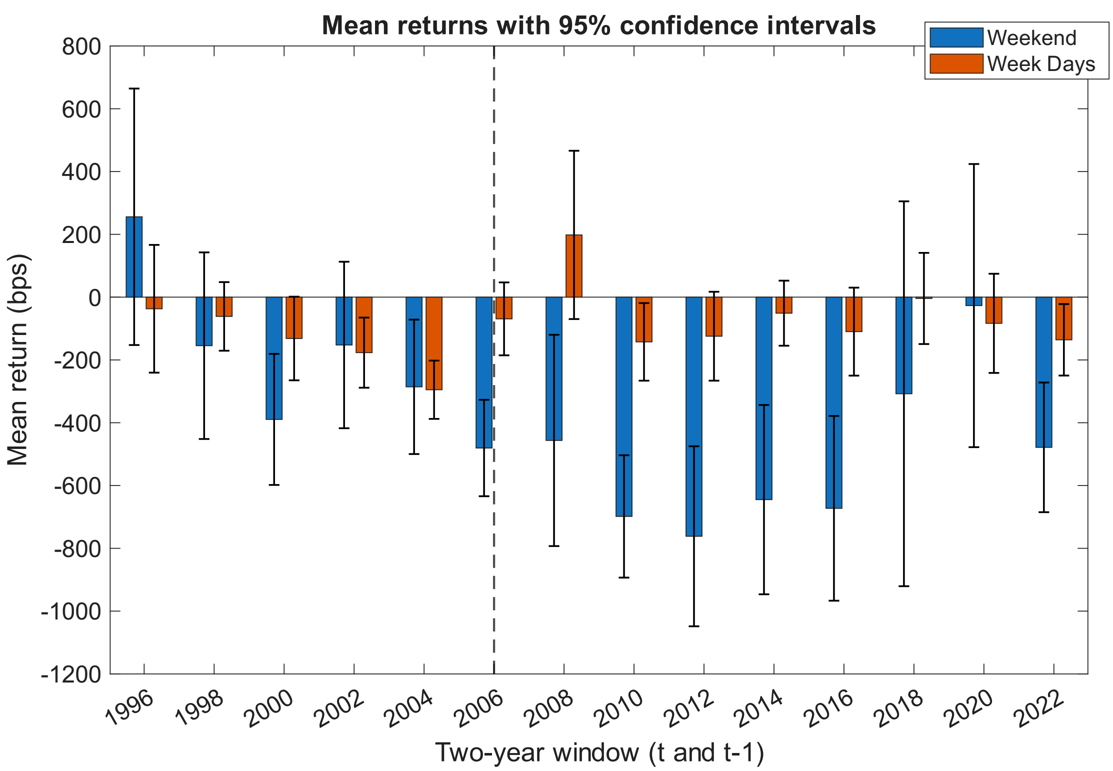

Note: This figure shows average week-night and weekend option returns over 2 year periods. The left bars show mean returns over weekends. The right bars show mean returns over week days. The error bars show the 95% confidence interval. The vertical line indicates the emergence of overnight equity trading around 2006.01. Returns are for the portfolio of out-of-the-money puts, are delta-hedged and are in basis points.

Figure A.2: Equity Trade Volume Around the Clock

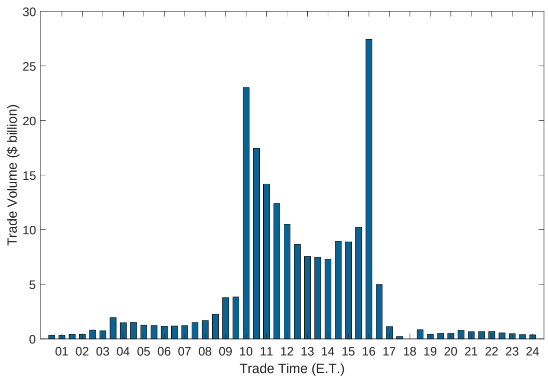

(a) Futures

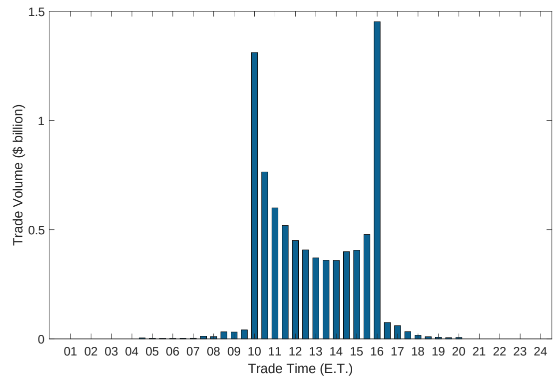

(b) Apple Inc.

Note: The figure shows average trade volumes for each 30-minute interval of the day. Panel (a) contains the most liquid S&P 500 E-mini futures contract, panel (b) contains Apple Inc. stocks, to provide an idea of overnight trade volumes in stocks themselves. The sample period is 2011 to 2023.

Figure A.3: Market Makers' Equity Market Risk and Inventory Shadow Gamma

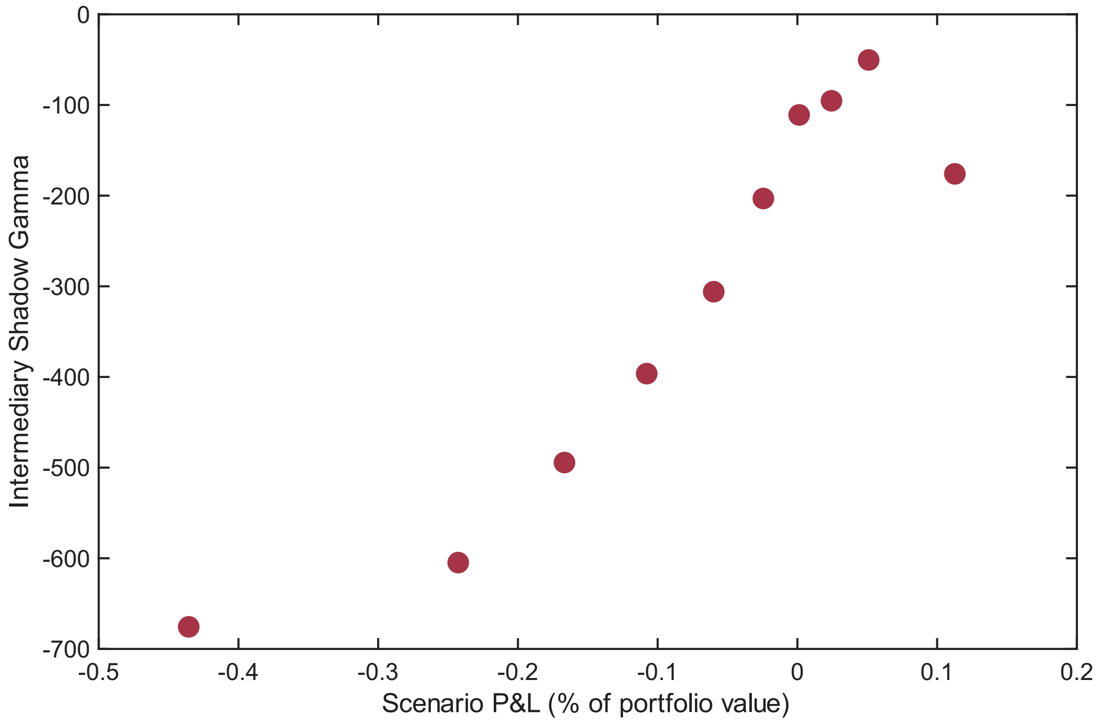

Note: This figure shows dealers' inventory shadow gamma by level of market maker scenario P&L. Shadow gamma and scenario P&L are measured at -15%. The sample period is 2011 to 2023.

Figure A.4: The Emergence of Overnight Equity Trading Relative To Weekends

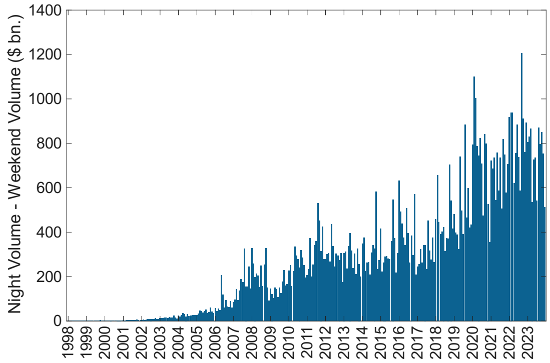

(a) S&P 500 Futures

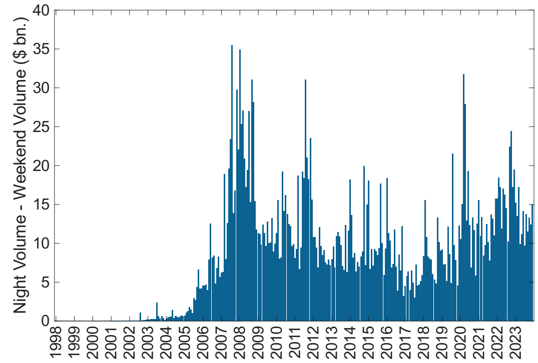

(b) S&P 500 ETF

Note: This figure shows the rise of overnight equity trading activity in S&P 500 instruments over week-nights, relative to weekends. Panel (a) reports the monthly sum dollar volume transacted overnight in the most actively traded S&P 500 E-mini futures contract over weeknights minus weekends. Weekends are measured from Friday 18:00 to Monday 09:30. Week-nights are all other nights. Panel (b) shows the corresponding series for the SPY ETF. Weekends trading is defined as transactions between Friday 16:00 and Monday 09:30 (E.T.). Week-nights are all other nights. Volumes are expressed in billions of dollars.

A.2. Appendix Tables

Table A.1: The Cross-Section of Night Returns

Puts		Days to Expiry 2-70	71-	All
$0.00 \leq \lvert\Delta\rvert \leq 0.25$	Deep Out of the Money	-390.1	-17.6	-307.1
$0.25 < \lvert\Delta\rvert \leq 0.50$	Out of the Money	-87.8	-24.3	-71.7
$0.50 < \lvert\Delta\rvert \leq 0.75$	In the Money	-59.7	-21.6	-47.6
$0.75 < \lvert\Delta\rvert \leq 1.00$	Deep In the Money	-69.0	-63.8	-63.4
All		-297.8	-22.2	-233.3
Calls				
$0.00 \leq \lvert\Delta\rvert \leq 0.25$	Deep Out of the Money	71.3	34.8	65.5
$0.25 < \lvert\Delta\rvert \leq 0.50$	Out of the Money	-8.0	-19.8	-7.1
$0.50 < \lvert\Delta\rvert \leq 0.75$	In the Money	-24.0	-12.9	-21.9
$0.75 < \lvert\Delta\rvert \leq 1.00$	Deep In the Money	-19.8	-7.7	-16.7
All		21.0	-1.2	14.9

Note: The table shows average S&P 500 option returns for eight portfolios, sorted by days to expiry and moneyness. Returns are measured from option market close at 16:15 to the subsequent market open at 09:45. Returns are in basis points and are delta-hedged at the beginning of the respective period. Days to expiry are calendar days. Newey-West t-statistics are in brackets. The sample period is 2011 to 2023.

Table A.2: Equity Liquidity and Option Returns: Triple-Differences

	(1) Baseline	(2) Winsorized	(3) Excl. Crashes
IntraWeek	85.7**	115.1***	85.7**
	(2.40)	(4.41)	(2.40)
Post	-84.4**	-98.6***	-154.7***
	(-1.98)	(-3.27)	(-3.89)
IntraWeek x Post	121.9***	92.9***	174.2***
	(2.60)	(2.74)	(3.87)
OtmPuts	-307.3***	-350.7***	-307.3***
	(-2.94)	(-5.27)	(-2.94)
IntraWeek x OtmPuts	61.7	55.9	61.7
	(0.56)	(0.73)	(0.56)
Post x OtmPuts	-305.6**	-348.1***	-409.7***
	(-1.97)	(-4.13)	(-3.24)
IntraWeek x Post x OtmPuts	470.9***	421.3***	561.0***
	(2.93)	(4.38)	(3.97)
Constant	-134.7***	-184.5***	-134.7***
	(-4.21)	(-7.92)	(-4.21)
Observations	54,842	54,842	53,675
R2-adjusted	0.00	0.01	0.01

Note: This table reports regression estimates from a triple-difference specification based on Equation (13). I regress option returns on indicators for treated options (OTM puts), treated periods (Intra-week), and the post-event period. Option returns are computed for eight out-of-the-money portfolios (four puts and four calls) with breakpoints defined in Table A.1. Intra-week returns include all close-to-close observations except Friday-to-Monday intervals. Column (1) presents the baseline regression; Column (2) uses returns winsorized at the 5th and 95th percentiles; Column (3) excludes the crisis months 2018.02 and 2020.02–2020.04. Returns are expressed in basis points and are delta-hedged. Standard errors are clustered by $month \times portfolio\ ID$. The sample period is 1996 to 2023.

Table A.3: The Impact of Stock Trading on Option Spreads: Diff-in-Diff

	Absolute Spread (cents)			Relative Spread (bps)		
	(1) Baseline	(2) Winsorized	(3) Excl. Crashes	(4) Baseline	(5) Winsorized	(6) Excl. Crashes
Intra-week	4.3***	3.7***	4.3***	1.6***	1.6***	1.6***
	(3.75)	(3.67)	(3.75)	(3.12)	(4.47)	(3.12)
Post	14.9***	8.8**	10.6**	-26.2***	-23.8***	-26.4***
	(2.88)	(2.28)	(2.37)	(-14.69)	(-20.54)	(-14.91)
Intra-week x Post	-7.1***	-5.8***	-6.5***	-1.4**	-1.2***	-1.2**
	(-3.87)	(-4.32)	(-3.93)	(-2.39)	(-3.12)	(-2.25)
Constant	97.6***	97.4***	97.6***	38.7***	36.0***	38.7***
	(34.43)	(36.95)	(34.43)	(23.42)	(37.11)	(23.42)
Observations	6,958	6,958	6,812	6,955	6,955	6,809
R2-adjusted	0.01	0.00	0.00	0.44	0.52	0.44

Note: This table reports estimates from a difference in differences specification that includes indicators for treated periods (IntraWeek), and the post-event period. Intra-week spreads are daily option spreads at close, except Friday. *Post* indicates the period after 2006.01 when overnight stock trading emerged. Spreads are measured for the portfolio of out-of-the-money S&P 500 puts. Columns 1 to 3 show absolute spreads in cents, columns 4 to 6 show relative spreads in basis points. Columns (2) and (5) winsorize option spreads at percentiles 5 and 95. Columns (3) and (6) exclude the crash months of 2018.02 and 2020.02 to 2020.04. Standard errors are clustered within every month. The sample period is 1996 to 2023.

Table A.4: Overnight Equity Liquidity and Option Returns: Continuous-Treatment Estimates

	(1)	(2)	(3)
Liquidity x OTM puts	293.3***	153.5***	347.4***
	(4.77)	(5.18)	(5.72)
Day	Yes	Yes	Yes
ID x Month	Yes	Yes	Yes
Observations	51,696	51,696	50,529
R2-adjusted	0.36	0.29	0.34

Notes: This table reports panel regressions of portfolio-level, delta-hedged option returns on an equity-liquidity measure and its interaction with an indicator for deep out-of-the-money (OTM) puts. Returns are daily close-to-close returns, computed for eight OTM portfolios (four puts and four calls) with breakpoints defined in Table A.1. Liquidity is the trading volume of the front-month E-mini S&P 500 futures contract measured over the night period of the corresponding return window, scaled by the number of hours in that night window and standardized to zero mean and unit variance. Thus, I adjust for the longer periods over for example weekends. The specification includes fixed effects for the return-window date and for option-portfolio $\times$ month; main effects of liquidity and the OTM-put indicator are therefore absorbed. Column (1) reports the baseline estimates; column (2) winsorizes returns at the 5th and 95th percentiles; column (3) excludes February 2018 and February–April 2020. Returns are in basis points. Standard errors are two-way clustered by option portfolio identifier and month. The sample period is 1996 to 2023.

Table A.5: The Impact of Stock Trading on VIX Futures Returns: Diff-in-Diff

	Intra-week	Post	Intra-week × Post
Coefficients	65.4	-63.4	51.2
t-stats	3.4	-2.9	1.5

Note: This table reports estimates from regressing daily VIX futures returns on an indicator for intra-week returns and an indicator for the period post 2006.01, when overnight equity trading emerges. Returns are daily close-to-close, so the intra-week indicator marks returns between Monday close and Friday close. Returns are measured from VIX futures mid quotes and are in basis points. $t-$ statistics are Newey-West with 10 lags. The sample period is 2004.04 to 2007.12.

Table A.6: S&P 500 Equity Returns

	Mean	Std	Min	P1	P50	P99	Max
Day	1.4	98.9	-926.2	-274.1	4.3	277.8	817.8
Night	2.1	67.7	-749.4	-210.3	4.0	182.2	585.8
Weekend	-1.2	84.4	-1,235.7	-227.0	3.3	223.6	494.0

Note: This table displays summary statistics for S&P 500 equity returns. Returns are measured from S&P 500 stock trade prices. Day returns are from 09:30 to 16:00, Night returns are from 16:00 to 09:30 and Weekend returns are from Friday 16:00 to Monday 09:30. P1 indicates the first percentile. The sample period is 1996 to 2023.