The decline of the S&P 500 variance risk premium

Ian Dew-Becker and Stefano Giglio*

June 2, 2026

Abstract

Equity index options historically displayed sharply negative returns and CAPM alphas. This paper examines whether that is still true. It studies returns on S&P 500 options between 1987 and 2025 and provides novel evidence that there was a structural break in risk premia around 2012. After 2012 – a sample equally long as those in the original studies that found negative premia – standard options strategies no longer have statistically significant alphas or information ratios. The decay in the returns is contemporaneous with a change in dealer option portfolios, and also comes around the same time that many other prominent strategies decayed. The results have implications for the interpretation of the VIX and variance risk premium along with portfolio choice.

1 Introduction

A major empirical fact in financial markets is that equity index options have been overpriced historically relative to simple benchmark models. Investors who purchase options have, on average, earned highly negative returns and CAPM alphas.[1]

*Dew-Becker: Federal Reserve Bank of Chicago. Giglio: Yale University and NBER. The views in this paper are those of the authors and do not represent the views of the Federal Reserve Bank of Chicago. We appreciate helpful comments from Gurdip Bakshi, Mike Chernov, Itamar Drechsler, Bjorn Eraker, Niels Gormsen, and seminar participants at ITAM, Yale, the University of Connecticut, Northwestern, Johns Hopkins, the University of Miami, INSEAD, HEC Paris, Berkeley, the Canadian Derivatives Institute, Boston University, the SFS Cavalcade, Booth, HBS, SAIF, Kansas University, CFIRM, Rice, SITE, FIU, the NBER, MIT, and Georgetown.

[1]For analyses of returns, see Coval and Shumway (2001), Bakshi and Kapadia (2003), Broadie, Chernov, and Johannes (2007), Constantinides, Jackwerth and Savov (2013), Chambers, Foy, Liebner, and Lu (2014), Dew-Becker et al. (2017), and Muravyev and Ni (2020), among many others. For structural models, see Backus, Chernov, and Martin (2011), Drechsler and Yaron (2011), Gabaix (2012), Drechsler (2013), Seo and Wachter (2019), and Schreindorfer (2020). Note, though, that those models are almost exclusively endowment economies.

The basic contribution of this paper is to examine whether S&P 500 options remain overpriced. While there was clear evidence prior to around 2012 that options earned negative returns, that is no longer true. The initial papers finding that S&P 500 options appeared to be overpriced were based on 10-15 years of data.[2] Over the most recent 10-15 years, the same returns are no longer statistically or economically significantly negative and in many cases are actually positive.

More specifically, this paper examines mean returns and CAPM alphas for S&P 500 options, both individually and in portfolios (straddles and the VIX portfolio) with and without delta hedging (i.e. hedging the returns for their local exposure to the S&P 500). It replicates findings from past work that to varying degrees and levels of statistical significance there were negative returns on S&P 500 options strategies between the late 1980's and early 2000's. Consistent with Broadie, Chernov, and Johannes (2007), henceforth BCJ, the results are strongest for straddles and delta-hedged strategies because the distributions of the returns are better behaved.

Subsequent to about 2012, the observed premia largely disappear. While there is a literature that has studied time-variation in various option premia, no other recent work quantitatively analyzes the robustness of the most basic features of returns over time.[3] Statistical tests for a break in the returns process, even accounting for the effects of multiple testing, reject the null of no break in most of the cases. The paper also provides simulation evidence that the statistical tests have correct size.

The paper then develops an equilibrium model in which heterogeneous beliefs among investors drive trade in options.[4] When investors are able to buy options frictionlessly but face barriers to selling them – a situation we argue characterized the period prior to the early 2000's – the model implies that options are overpriced in the sense that their returns relative to the underlying (again, with and without delta hedging) are negative. As the friction eases

[2]E.g. Coval and Shumway (2001), Bakshi and Kapadia (2003), and Broadie, Chernov, and Johannes (2007).

[3]The closest analysis to what is in this paper is in Bates (2022), who also looks at cumulative option returns, in his case up to the end of 2020. The additional data here is important for increasing power. Additionally, that paper uses a weekly instead of daily delta hedge, which also reduces statistical power (see Broadie, Chernov, and Johannes (2009)). Bates (2022) also finds a decline in premia, but dates it somewhat later than us – 2017 instead of 2012.

For other work on time-variation in options premia, see, among many others, Bollerslev and Todorov (2011), Konstantinidi and Skiadopoulos (2016), Andersen, Fusari, and Todorov (2020), and Fournier, Jacobs, and Orlowski (2024).

[4]For related work on intermediaries, see Jackwerth (2000), Bollen and Whaley (2004), Bates (2008), Han (2008), Garleanu, Pedersen, and Poteshman (2008), Jurek and Stafford (2015), Haddad and Muir (2021), Frazzini and Pedersen (2022), among many others. A related literature has also explored the link between option markets, returns in various asset classes, and the role of risk taking by intermediaries: for example, Brunnermeier, Nagel and Pedersen (2008), Bao, Pan and Wang (2011), Longstaff et al. (2011), Nagel (2012), and Chen, Joslin and Ni (2019).

and retail investors can also supply options, the equilibrium price falls and returns move to zero, consistent with the empirical results.

In addition to being able to explain why there would have been a decline in option premia, the model has an additional prediction, which is that the decline should have coincided with a decline in the asymmetry of positions held by unconstrained investors, which we model as dealers. Similar to Garleanu, Pedersen, and Poteshman (2008), dealers in the model absorb excess demand from the retail agents. The model's prediction is that frictions cause asymmetry in demand and hence overpricing. Dealer positions reveal that asymmetry, so dealer positions should predict returns, and they should shift to zero at the same time that returns do. And in fact that is what we observe in the data. The options premium goes away around 2012, which is the same time that the net positions of dealers shift from negative to neutral. Additionally, dealer option exposures have significant predictive power for future option returns.

Beyond the positions, an extension of the model also predicts that declines in the frictions faced by dealers should have reduced premia. At the same time that net dealer positions shifted from negative to neutral, the frictions that they face in hedging options positions, including the bid/ask spread of the underlying and basis risk, also declined.

The model's third implication is that the change in the returns on delta-hedged options comes from a change in the return on the option leg of the trade, not the hedge. The return on the hedge itself – which is meant to approximately replicate an option return via dynamic trading in the underlying, is predicted to have been stable over time. In the data, that is what we observe. We show that a delta hedge does a surprisingly good job of capturing option returns, including providing economically meaningful nonlinearity in returns. Additionally, constructing a delta hedge does not actually require options to exist – it just requires a volatility forecast – and we find that not just in the modern period but all the way back to 1926, there is no evidence for a nonzero CAPM alpha. In other words, there was only a relatively brief period between 1987 and about 2010 where traded options earned a (negative) CAPM alpha. Outside that period, neither traded options nor their dynamic replication strategy had a nonzero alpha in either direction.

This paper adds to the list of prominent asset pricing anomalies whose performance declined over the past two decades. Predictors of both aggregate stock market returns and returns on individual stocks are well known to often have significant trouble out of sample. In addition, average returns across equity anomalies have declined to near zero since the early 2000's. The value and momentum portfolios on Kenneth French's website had approximately zero cumulative returns between 2002 and 2025. In currency markets, the carry trade has shown no outperformance since the Global Financial Crisis. The behavior of options returns thus appears consistent with what has been observed more broadly, which, if not an increase in overall market efficiency, is at the very least a decline in the strength of many of the most famous strategies.

The fact that options no longer earn the premia that they used to also calls into question some of the theories used to explain their previous behavior. For example, past work argued that ambiguity, rare disasters, or long-run risk could explain the pricing of options.[5] But if option prices have declined, then those types of aggregate consumption-based risks should have also declined. As Chen, Dou, and Kogan (2024) discuss, though, testing that is difficult. Consumption-based models often rely on what they call "dark matter" that is not directly observable. Intermediary-based models have the nice feature that they are driven by relatively more easily observed state variables, and those variables line up well with the empirical observed decline in option premia.

That is not to say that the consumption-based models are ruled out by the results – the point of Chen, Dou, and Kogan (2024) is precisely that they are difficult to test and reject. Instead, it is simply that intermediary models have the ability to explain the observed shift. More generally, the paper's results on the performance of both options and the delta hedge represent new benchmarks for models to match.

The results are also relevant for the interpretation of the VIX index. In the period when there was a large variance risk premium, the VIX's average value was about 3 points above average realized volatility. As the variance risk premium has declined, the average of the VIX has nearly converged to that of realized volatility so that it is, at least on average, nearly an unbiased predictor of volatility.

The results also have applications to portfolio choice. They imply that for many investors, volatility and jump risk are not relevant risk factors. While they may represent an uninsurable risk, they are not one that has captured any premium in the recent data, at least based on simple option strategies. That means that for an investor who is choosing a static portfolio (since all our results are unconditional), they do not improve the mean-variance tradeoff compared to just investing in the total equity market.

The remainder of the paper is organized as follows. Section 2 describes the data and contains the core analysis of option returns. Section 3 develops a heterogeneous-agent model that is able to rationalize the empirical results, and it then tests the model's overidentifying implications. Finally section 4 briefly discusses broader implications of the results and section 5 concludes.

[5] E.g. Drechsler and Yaron (2011), Gabaix (2012), Drechsler (2013), and Seo and Wachter (2019).

2 The history of S&P 500 index option returns

This section examines returns on different options strategies over time and tests for whether there has been a break in their means.

2.1 Data

The dataset for traded options splices together CME futures options for the period 1987–1995 with CBOE SPX options from Optionmetrics for 1996–2025. Following Broadie, Chernov, and Johannes (2009), we study a monthly rolling strategy, where options are purchased on the third Friday of every month and then held to their maturity on the following month's third Friday.

In calculating returns, we scale the option payoffs by the initial option price. Formally, if the option price is O_t, then the option return is defined as $\left(X_t - O_{t-1}\right)/O_{t-1}$, where X_t is the option's terminal payoff in month t. This is the standard return that an investor purchasing the option would earn per dollar invested.

We discuss robustness to this choice and present results scaling by the price of the underlying, P_{t-1}, in the denominator, which corresponds to the return per unit of insurance (as in Büchner and Kelly (2022)), in external appendix 1.

We focus on returns on five different strategies. The first four use vanilla options: 5% out-of-the-money puts and at-the-money straddles, with and without daily delta hedging. Last, we use the standard result from Carr and Wu (2009) that, if the level of the S&P 500 is approximately a geometric diffusion, the VIX index is the square root of the price of a variance swap. The variance swap return is then $\left(\mathrm{RV}_t - \mathrm{VIX}_{t-1}^2\right)/\mathrm{VIX}_{t-1}^2$, where RV_t is annualized realized variance over the holding period (based on the daily sum of squared log returns on the S&P 500) and VIX_{t-1}^2 is the squared VIX at the beginning of the month.

The analysis focuses on CAPM alphas. The alphas are relevant not because the CAPM is necessarily the "true" model of risk premia, but rather simply as a benchmark to evaluate whether options carry any independent premium (similar to how equity factors are evaluated) beyond what is embedded in the market return.

Finally, the analysis takes August 2012 as its baseline break date, based on the timing of a shift of the net holdings of dealers, discussed in section 3.4 (and see figure 5).[6] To be clear, though, the paper's primary statistical break tests are based on methods that account for multiple testing across many potential dates (section 2.2).

[6] Specifically, that is the last date where a 12-month moving average of intermediary S&P 500 gamma was negative.

2.2 Returns and alphas over time

Figure 1 presents cumulative returns and CAPM alphas for the five different strategies. The cumulative returns (top panels) and alphas (bottom panels) are calculated by summing the monthly returns over time. All returns are rescaled so that they have the same monthly standard deviation as the straddles. The left-hand side gives results for strategies without delta hedging and the right with delta hedging (where the variance-swap strategy is grouped with the delta-hedged strategies). The dashed vertical line in each panel represents the baseline break date of August 2012.

In all four panels, the returns are clearly negative in the first part of the sample up to about 2012. Subsequently, though, to varying degrees, the returns appear to shift towards zero, so that the cumulative return lines flatten out. The flattening is stronger when we control for market returns, either via delta hedging or by looking at CAPM alphas.[7]

To put numbers on that, figures 2 and 3 use bar charts to report measures of the mean and volatility of returns over the full sample and the pre- and post-August/2012 subsamples. In each panel there are four bars: the first represents the full-sample estimates ("Full"), the second the pre-2012 estimates ("Pre"), the third the post-2012 estimates ("Post") and the last the difference between the Post and the Pre estimates ("Diff"). Each bar includes whiskers representing 90- and 95-percent confidence intervals. The columns correspond to the five strategies.

Figure 2's three rows report alpha, beta, and information ratios. The figure shows that alphas and information ratios declined in absolute value (so the difference between post and pre 2012 is positive), while betas were generally stable. The alphas are statistically significantly negative in the first part of the sample for all five strategies, and they all increase post-2012, becoming positive for the straddle, the delta-hedged straddle, and the variance-swap series, while remaining negative but no longer statistically significant for the 95% put and the delta-hedged 95% put. The change in the alpha is statistically significant at the 5% level for the straddle and the delta-hedged straddle, and at the 10% level for the delta-hedged 95% put.

In a mechanical sense, an increase in the alphas (from negative to zero) could be explained by the betas becoming more negative (holding the mean returns fixed), but that is not what we observe – the betas show little change, with no clear pattern across the strategies.

Similar results hold for raw returns: figure 3, whose rows report means, standard deviations, and Sharpe ratios, shows similar, though somewhat weaker, declines post-2012. Again,

[7] Even after delta hedging, the options can still have an exposure to the market for at least three reasons: the model used to calculate the delta is not perfect, the hedging is done discretely rather than continuously, and the underlying can potentially jump.

Figure 1: Cumulative returns and alphas

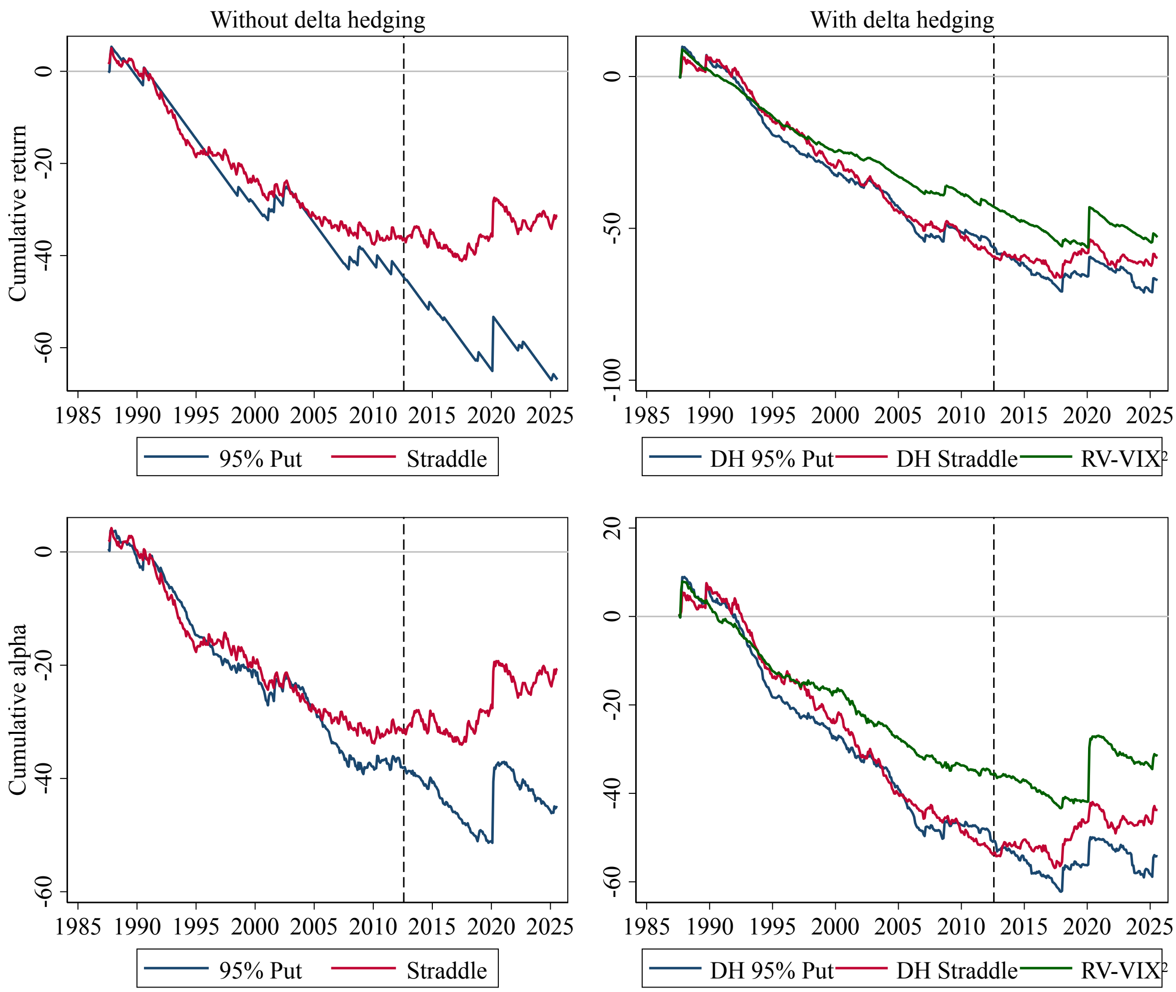

Description: The top panels show cumulative returns, and the bottom panels show cumulative CAPM alphas, for the five option strategies, with returns scaled by the initial option price ($\left(\mathrm{RV}_t - \mathrm{VIX}^2_{t-1}\right)/\mathrm{VIX}^2_{t-1}$ for the variance-swap series). Cumulation is done by simple summation rather than log-compounding, since option-price-scaled returns can fall below -1. The left-hand panels show strategies without delta hedging (95% put and ATM straddle) and those on the right show delta-hedged strategies (delta-hedged 95% put, delta-hedged ATM straddle, and $\left(\mathrm{RV} - \mathrm{VIX}^2\right)/\mathrm{VIX}^2$). All series are standardized to have the same monthly standard deviation as the straddle before cumulating. The dashed vertical line marks August 2012. Sample period: monthly returns, 1987:08–2025:08, $N = 457$ monthly observations for each of the five strategies. **Interpretation:** Cumulative returns and alphas trend strongly negative through about 2012 and then flatten (more clearly so for alphas), indicating that the negative option premia documented in the earlier literature have sharply declined in the post-2012 period.

Figure 2: Alphas, betas, and information ratios

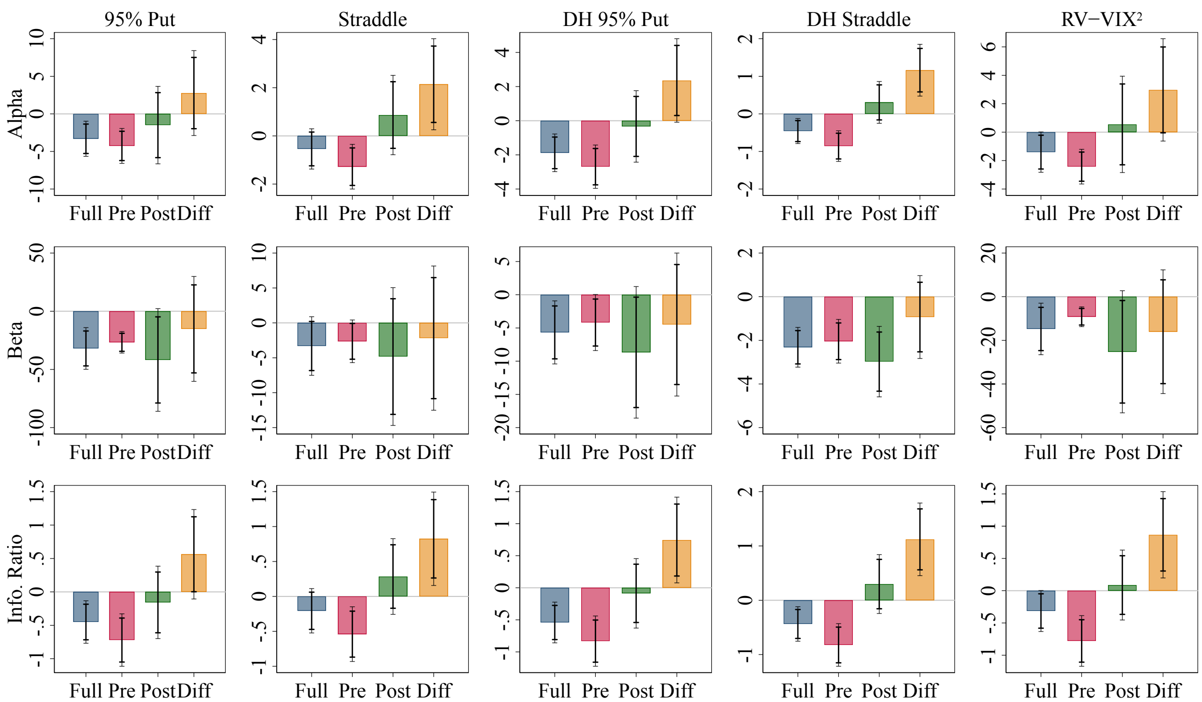

Description: Each column corresponds to one of the five strategies, with returns scaled by the initial option price ($\left(\mathrm{RV}_t - \mathrm{VIX}_{t-1}^2\right)/\mathrm{VIX}_{t-1}^2$ for the variance-swap series). The top row shows annualized CAPM alphas (computed using the full-sample beta), the middle row shows CAPM betas, and the bottom row shows annualized information ratios. The four bars in each panel represent the full sample, the pre-August 2012 subsample, the post-August 2012 subsample, and the difference (post minus pre). Thick and thin whiskers represent 90- and 95-percent confidence intervals based on robust standard errors. Sample period: monthly returns, 1987:08–2025:08, $N = 457$ monthly observations (pre-August 2012 subsample $N = 301$; post-August 2012 subsample $N = 156$). **Interpretation:** Alphas and information ratios were significantly negative in the pre-2012 sample but rise toward (and in some cases above) zero post-2012, with the change statistically significant for the straddle, the delta-hedged straddle, and the delta-hedged 95% put; CAPM betas show no comparable shift.

all the mean returns are statistically significantly negative at the 5% level in the first half of the sample, but that no longer holds in the second half. The change itself is statistically significant at the 5% level for the delta-hedged straddle and at the 10% level for the straddle.

Figure 3: Returns, volatility, and Sharpe ratios

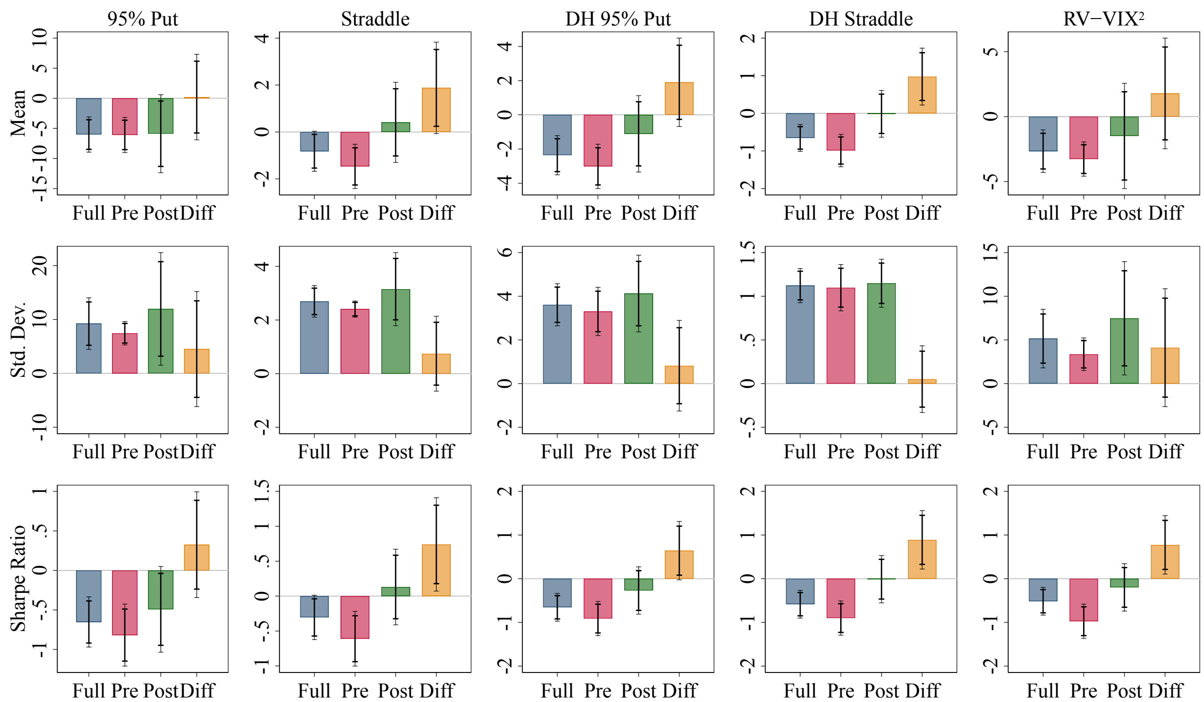

Description: Each column corresponds to one of the five strategies, with returns scaled by the initial option price ($\left(\text{RV}_t - \text{VIX}_{t-1}^2\right)/\text{VIX}_{t-1}^2$ for the variance-swap series). The top row shows annualized mean returns, the middle row shows annualized standard deviations, and the bottom row shows annualized Sharpe ratios. The four bars in each panel represent the full sample, the pre-August 2012 subsample, the post-August 2012 subsample, and the difference (post minus pre). Thick and thin whiskers represent 90- and 95-percent confidence intervals based on robust standard errors. Sample period: monthly returns, 1987:08–2025:08, $N = 457$ monthly observations (pre-August 2012 subsample $N = 301$; post-August 2012 subsample $N = 156$). **Interpretation:** Mean returns and Sharpe ratios are significantly negative in the pre-2012 subsample but shift toward zero in the post-2012 subsample; standard deviations are essentially unchanged, so the decline in Sharpe ratios is driven by the change in means rather than in volatility.

One common feature across the figures is that the results are weakest for the raw returns on the unhedged 95% put. The reason for that is that market returns accelerated post-2012 – from 5.9 to 13.2%. The 95% put inherits that due to its negative beta. The results that remove market returns, either through delta or beta hedging, avoid that issue.

Overall, then, figures 2 and 3 provide the simplest evidence that the CAPM alphas and mean returns earned by standard option strategies have changed significantly in the last 15 years compared to the values observed prior to 2012. The results are clearly strongest for the cases where returns are adjusted for the market, or have weaker exposure to the market to begin with (straddles and variance swaps). That result is exactly what to expect based on the results of BCJ who noted that the extreme behavior of unadjusted vanilla option returns makes them unattractive both for testing models and trying to measure average returns.

2.3 Tests for a break in the returns process

2.3.1 Methods

The results so far are meant to give a basic description of the data and show that it appears that realized returns have shifted such that they are no longer distinguishable from zero on average. There are some concerns that one should have with those results, though. First, the confidence bands are based on usual Gaussian asymptotics that might be a bad representation of behavior in the samples we have. Second, while we provide evidence motivating the use of August 2012 as a break date, one might naturally worry that that date was selected after some snooping.

The usual econometric method to address the second concern is to use asymptotics based on the functional central limit theorem to characterize the behavior of test statistics calculated on many dates in the sample (Andrews (1993) and Andrews and Ploberger (1994)). Those are the baseline starting point.

Those methods, however, impose even *stronger* versions of the central limit theorem than in the previous section because they effectively assume that in every nontrivial subsample of the data (Andrews and Ploberger (1994) propose at least 15%), the mean is well approximated as being normally distributed. While that is reasonable in many economic settings, it is highly questionable in the case of options returns with only about 40 years of data.

We therefore also apply the median-based estimator of Hodges and Lehmann (1963), which imposes weaker distributional assumptions. At the same time, it tests only for shift in location, holding the rest of the shape of the distribution fixed, which can help increase statistical power.

Dehling, Fried, and Wendler (2020) show how to generate p-values for the Hodges–Lehmann estimator when the break date is unknown, as is the case here. Formally, the assumption is that the returns process is

$$r_t = \mu_t + \varepsilon_t \tag{1}$$

where μ_t is a variable that is constant outside of a shift at some unknown $\bar{t}$ (i.e. $\mu_t = a$ for

$t < \bar{t}$ and $\mu_t = b$ for $t \geq \bar{t}$). ε_t is a residual with an unconditional mean of zero, but its moments may be serially correlated (e.g. it may be an ARMA or GARCH process), and in fact the moments of ε_t (beyond the mean) need not even exist.[8]

Controls can be added to (1). The most natural control is the market return, in which case we have a standard market model,

$$r_t = \alpha_t + \beta r_{m,t} + \varepsilon_t \tag{2}$$

and the test is for a shift in α_t. Note that even delta-hedged returns can still have nonzero beta since they are still exposed to jumps in the underlying, which have a market component when the jumps are asymmetrical.

The usual test for a shift in the intercept is to estimate the change in the conditional mean of r (or $r - \beta r_m$) before and after each t. However, in the presence of heavy-tailed and skewed distributions, such a test may have both low power and poor sampling properties since the central limit theorem will be a poor approximation for the sampling distribution of the mean. The Hodges–Lehmann (HL) estimator is instead based on the *medians* before and after each t. Under the assumption that the distribution of ε is stationary, an estimate of a shift in the median also is an estimate of a shift in the mean.[9]

It is important to note that the assumption that the distribution of ε_t is the same before and after the break is stronger than what is used by estimators such as Andrews (1993). The HL estimator therefore uses weaker assumptions than the typical mean-based estimators about the existence of moments and convergence to the normal distribution, making it more robust to heavy tails, but at the same time it uses a *stronger* assumption about stationarity of the shape of the distribution, which means that it will be less robust in the presence of higher-order shifts in the distribution. For example, if μ (and/or α) are constant over time but the skewness of ε changes, that could cause the HL estimator to reject the hypothesis of stationarity. It is therefore necessary to first test for evidence of a change in the shape of the distribution of ε. If the shape of the distribution is in fact well modeled as stationary, then the HL estimator will gain power from imposing that assumption.

[8]That follows from the use of the median, which exists for all distributions. Confidence bands here come from a central limit theorem applied to relative ranks, which are bounded between 0 and 1, instead of scaled sums which can be unbounded.

[9]To be clear, the idea here is not to test whether mean or median CAPM alphas are zero, but rather whether they shifted. The HL estimator looks for a shift in the location of the distribution of r or $r - \beta r_m$. When the shape of the distribution is stationary, a shift in its location corresponds exactly to a shift in its mean.

2.3.2 Stationarity of the residuals

To test for the stationarity of the ε's, we demean the five returns processes before and after August 2012 and then apply the Kolmogorov–Smirnov and Cramér–von Mises tests of a change in distribution (e.g. Darling (1957)) to the pre- and post-August/2012 samples.

Table 1 reports those results. For beta-adjusted returns (i.e. CAPM alphas) the tests do not reject the hypothesis that the distribution of the ε_t is stable over time for the put, straddle, and delta-hedged versions of each, and are generally not close to doing so. The variance-swap series $\left(\left(\text{RV} - \text{VIX}^2\right)/\text{VIX}^2\right)$ rejects at the 5% level; however, the tests do not reject the null for the variance-swap series that is not scaled by the VIX. That implies that for variance swaps, the unscaled version is better suited to the HL test. In what follows, we report both results.

For raw returns – unadjusted by the market – the 95% put and the variance-swap series both reject significantly, while the straddle and the delta-hedged strategies do not. This is consistent with the warning of BCJ that in general the returns are better behaved when using delta hedging and straddles. Again, as with beta-adjusted returns, the unscaled version of the variance swap is better behaved than the scaled version.

Overall, then, the assumption that the residuals have stable distributions over time, required for use of the HL estimator, appears generally reasonable for the put and straddle strategies, especially when applied to delta-hedged or market-adjusted returns. For the scaled variance-swap series the distribution does appear to shift, which means HL results for that strategy should be interpreted with some caution; the unscaled variance-swap series has a more stable distribution, especially when beta-adjusted.

2.3.3 Simulation results

To validate both the Hodges–Lehmann and Andrews–Ploberger (AP; 1994) estimators here, appendix figures A.1 and A.2 report results from two sets of simulations. Figure A.1 simulates option returns (for the two delta-hedged series plus $(\text{RV} - \text{VIX}^2)$, all market-adjusted; note that we use the unscaled version of the variance swap series as it satisfies the stability assumption required by the HL estimator) by drawing randomly from the full sample of returns, and adding a mean shift in the second part (post August 2012) of the sample. The x-axis varies the size of that shift as a fraction of the estimated shift from the empirical data. It therefore examines the ability of the two estimators to detect a change in mean option returns, holding the distribution otherwise equal.

For all three strategies, when the shift is zero (so there actually is no shift and the null is true), figure A.1 shows that both methods reject the null at the 5% level almost exactly

Table 1: Tests for changes in the distribution of returns

	KS test p-value	CvM test p-value
Panel A: Beta-adjusted returns		
95% Put	0.50	0.60
Straddle	0.17	0.18
DH 95% Put	0.54	0.52
DH Straddle	0.72	0.86
$\left(\text{RV}-\text{VIX}^2\right)/\text{VIX}^2$	0.03	0.01
$\text{RV}-\text{VIX}^2$ (unsc.)	0.41	0.26
Panel B: Raw returns		
95% Put	<0.01	<0.01
Straddle	0.23	0.19
DH 95% Put	0.52	0.50
DH Straddle	0.46	0.45
$\left(\text{RV}-\text{VIX}^2\right)/\text{VIX}^2$	<0.01	<0.01
$\text{RV}-\text{VIX}^2$ (unsc.)	0.10	0.12

Description: Kolmogorov–Smirnov and Cramér–von Mises two-sample tests comparing the pre- and post-August 2012 distributions of monthly returns, after demeaning each subsample separately. Panel A uses beta-adjusted returns (subtracting the full-sample CAPM beta times the market return); Panel B uses raw returns. Cramér–von Mises p-values are computed by permutation. Sample period: monthly returns, 1987:08–2025:08, $N = 457$ monthly observations (pre-August 2012 $N = 301$; post-August 2012 $N = 156$). **Interpretation:** For beta-adjusted series, the within-subsample distribution of returns (after demeaning) is statistically indistinguishable across the pre- and post-2012 subsamples, supporting the use of break tests that assume distributional stability; the scaled variance-swap series is the one exception.

5% of the time. So the tests have the correct size in the sense that if there is truly no break in the data, the tests reject the null at the correct rate.

As the size of the shift grows, so do the rejection probabilities. For an ideal test, the rejection probabilities jump to 1 when there truly is a shift, regardless how small. The rejection probabilities for the HL test do in fact rise quickly, passing 50% for a shift that is about 2/3 as large as what is observed empirically. For a shift of what we estimate empirically, the HL test is expected to reject nearly 100% of the time for all three cases.

For the AP mean-based method, on the other hand, power is much lower. For a shift of the size observed in the data, its probability of rejecting the null is markedly below that of the HL test in all three cases.

Figure A.2 repeats the simulations but draws separately from the pre- and post-2012 subsamples, thus allowing for different distributions before and after 2012. The results are similar to the previous picture, with the HL test showing correct size but much greater power compared to AP.

Figure A.3 reports results that, instead of changing the size of the shift, change the length of the sample, while keeping the size of the shift as large as what we estimate empirically. The samples of varying length are obtained by drawing randomly from the full sample, and adding the estimated shift to the second half of the simulated sample. The simulation therefore illustrates how much data the two methods need in order to be able to consistently detect changes of the size observed empirically. For HL, with a sample the size of the actual data, rejection probabilities are 80-95% and they rise from there. For AP, the sample needs to be 2-5 times larger to get similar power. Given that the sample is nearly 40 years long, we would need up to 200 years of data for the mean-based estimator to perform as well in this setting as the median-based estimator.

2.3.4 Empirical break test results

Table 2 reports results from three tests for a break:

1. A standard test for a shift in the mean (using heteroskedasticity-robust standard errors) assuming a known break of August 2012

2. The HL-based test (using the sampling distribution of Dehling, Fried, and Wendler (2020))

3. The AP exponential-Wald statistic.

For each test, it reports a p-value along with the estimated change in mean returns and the estimated break date.[10] Panel A reports results for beta-adjusted returns and panel B

[10]Since the Andrews–Ploberger (1994) test is based on a sum of exponential Wald statistics, the reported break date and value are based on the date that maximizes the Wald statistic.

raw returns. The HL estimator rejects the null of no change in all specifications, at varying levels of significance; for beta-adjusted returns the HL estimator dates the change to the period 2006–2010. Additionally, comparing the results here to those in figures 2 and 3, the estimated magnitude of the change in each case is similar.

The fixed-date Wald test gives broadly consistent results. For beta-adjusted returns it rejects the null at the 10% level or stronger for the straddle, the delta-hedged 95% put, the delta-hedged straddle, and the (preferred) unscaled variance swap. It does not reject for the 95% put, and is marginal for the scaled variance-swap. As before, the test is strongest for delta-hedged and market-adjusted strategies, consistent with BCJ.

Finally, the AP exp-Wald statistic is the weakest of the three statistically, consistent with the simulations. Even though the estimated mean shifts have very similar magnitudes to the estimates from the other two methods, the AP p-values are larger, especially compared to the HL statistic. After adjusting for market beta, the AP test rejects at the 5% level for the straddle and the delta-hedged straddle and at the 10% level for the delta-hedged 95% put. The 95% put and variance-swap series are not statistically significant under AP.

The lower significance of the AP results is consistent with the simulation results finding that that test has low power in samples of our size for changes of the magnitude observed empirically. The more powerful HL test, on the other hand, finds clear evidence of a change, consistent with what is visible from the cumulative returns in figure 1. Beyond the specific findings on option returns, the results here suggest that the HL test may be an attractive alternative to AP for break tests in future work.

External appendix 1 shows that the results are robust to scaling returns by the level of the underlying instead of by the initial option price. External appendix 2 shows that the results extend across a range of strikes from 10% below the level of the underlying index to 2% above, for both unhedged and delta-hedged returns.[11]

2.4 Summary

This section contains the paper's main empirical results. It replicates the well known findings that for the first few decades of their existence, S&P 500 index options appear to have been overpriced in the sense that they had large negative CAPM alphas and mean returns. However, since 2012 there is no longer evidence for such behavior. The change in returns itself is statistically significant according to most (though certainly not all) of the estimators we study.

[11] That range is similar to what is chosen by BCJ and we use it both for consistency and because that is where there is most consistently liquidity through the full sample.

Table 2: Break tests for option strategy returns

		Change	p-value	Date
Panel A: Beta-adjusted returns				
	Fixed date	0.2307	0.337	8/2012
95% Put	HL	0.3378	0.076	4/2009
	Andrews	0.3376	0.352	1/2008
	Fixed date	0.1787	0.026	8/2012
Straddle	HL	0.1948	0.044	1/2010
	Andrews	0.1958	0.041	3/2010
	Fixed date	0.1964	0.059	8/2012
DH 95% Put	HL	0.2515	<0.001	3/2007
	Andrews	0.2515	0.071	3/2007
	Fixed date	0.0969	<0.001	8/2012
DH Straddle	HL	0.0813	0.001	2/2007
	Andrews	0.0981	0.007	11/2012
	Fixed date	0.2477	0.105	8/2012
RV − VIX^2	HL	0.2324	<0.001	8/2006
	Andrews	0.2569	0.161	3/2009
	Fixed date	0.0093	0.074	8/2012
RV − VIX^2 (unsc.)	HL	0.0080	0.007	9/2010
	Andrews	0.0096	0.347	12/2012
Panel B: Raw returns				
	Fixed date	0.0166	0.956	8/2012
95% Put	HL	-4.0137	<0.001	1/1988
	Andrews	0.2802	0.822	8/1998
	Fixed date	0.1565	0.059	8/2012
Straddle	HL	0.1716	0.098	1/2010
	Andrews	0.1631	0.072	10/2005
	Fixed date	0.1585	0.149	8/2012
DH 95% Put	HL	0.2334	<0.001	3/2007
	Andrews	0.2334	0.156	3/2007
	Fixed date	0.0814	0.011	8/2012
DH Straddle	HL	0.0632	0.014	10/2004
	Andrews	0.0832	0.057	11/2012
	Fixed date	0.1488	0.411	8/2012
RV − VIX^2	HL	0.1442	0.003	8/1998
	Andrews	0.1922	0.504	3/2007
	Fixed date	0.0044	0.500	8/2012
RV − VIX^2 (unsc.)	HL	-0.3386	0.019	12/1987
	Andrews	0.0046	1.000	2/2013

Description: Each panel reports results from three break tests for each of the five strategies, with returns scaled by the initial option price ($\left(\mathrm{RV}_t - \mathrm{VIX}^2_{t-1}\right)/\mathrm{VIX}^2_{t-1}$ for the variance-swap series). *Fixed date* is a Wald test (HC-robust) for a shift in the mean at the known break date of August 2012. *HL* is the Hodges–Lehmann test of Dehling, Fried, and Wendler (2020), which estimates the break date endogenously. *AP* is the Andrews–Ploberger (1994) exponential-Wald statistic. "Change" is the estimated shift in the monthly mean return. "Date" is the estimated or assumed break date. Panel A uses beta-adjusted returns (subtracting the full-sample CAPM beta times the market return); Panel B uses raw returns. Sample period: monthly returns, 1987:08–2025:08, $N = 456$ monthly observations. **Interpretation:** The HL test, the most powerful in our setting, rejects the null of no break for every strategy at conventional levels and dates the change to 2006–2010 for beta-adjusted returns; the fixed-date and AP tests deliver broadly consistent point estimates of the change but with weaker statistical significance, in line with the simulations.

That said, both the initial findings on option returns from the 1990's and ours for the post-2012 period are based on samples that are economically relatively short. Things change, and nothing here is about forecasting the future. The point is simply that the evidence that the literature has relied on in driving the conventional wisdom that options are overpriced and that there is a clear and significant variance risk premium (in the sense of an alpha on $(\mathrm{RV} - \mathrm{VIX}^2)/\mathrm{VIX}^2$) has dissipated.

3 The decline of option overpricing and the role of intermediaries

This section presents a heterogeneous agent model that can help rationalize the empirical findings from the previous section – a negative CAPM alpha prior to about 2012 that has now shifted close to zero.

3.1 Model setup

The model has three types of agents: two types of retail investors plus a small set of intermediaries.

In order to model frictions in trade, we consider an overlapping-generations model in which retail investors can only trade equity and options at birth. The lives of the investors will be calibrated to be equal to the maturity of options. In that sense, these agents can be thought of as making myopic portfolio decisions and rebalancing once per month.

There are two types of retail investors, indexed by i. They both live for $J+1$ periods, consume in each period, and have a bequest motive. They differ in how their effective risk aversion varies over time. The retail agents represent those who face investment frictions, which could, more generally, include pensions, insurance companies, or other managed investments that face constraints on trading options.

There is also a set of "dealers" or market makers – more generally, financial intermediaries – that have infinitesimal wealth relative to retail investors, but are unrestricted in their trading. Their wealth is sufficiently small that their supply curve for options is effectively vertical – take it to have slope $1/\varepsilon$ for some very small ε.

3.1.1 Budget constraints

The budget constraint in the first period of an agent's life is

$$C_{i,J,t} + B_{i,J,t} + P_t^X X_{i,t} + P_t^O O_{i,t} = P_t^X \tag{3}$$

where $C_{i,j,t}$ is consumption on date t of an agent of type i with j periods remaining to live, $B_{i,j,t}$ measures their holdings of riskless bonds, and $P_t^X X_{i,t}$ and $P_t^O O_{i,t}$ are their purchases or sales of equity and J-period options, respectively (these do not take j subscripts because equity and options can only be traded at birth). The right-hand side represents the agent's endowment – agents of all types are born with a unit allocation of equity.

In periods after their birth, when agents have $0 < j < J$ periods left to live, they receive dividends, trade bonds, and consume subject to the constraint

$$C_{i,j,t} + B_{i,j,t} = D_t X_{i,t-(J-j)} + R_t^B B_{i,j+1,t-1} \tag{4}$$

where R_t^B is the risk-free rate from date $t-1$ to t and D_t is the dividend paid by equity on date t.

Terminal wealth is then

$$W_{i,t} = X_{i,t-J}\left(P_t^X + D_t\right) + O_{i,t-J} X_t^O + R_t^B B_{i,1,t-1} \tag{5}$$

where X_t^O is the payoff of a J-period option on date t.

3.1.2 Retail agents' objective

The retail agents have log utility over consumption. On the day they are born their objective is

$$\max E_t^i \left[\log C_{i,J,t} + \sum_{k=1}^{J-1} \beta^k \log C_{i,J-k,t+k} + \beta^J \log\left(C_{i,0,t+J}\right) \right] \tag{6}$$

where E_t^i is the expectation operator on date t for agents of type i.

Since the agent has log utility, the consumption-wealth ratio is constant. We therefore model terminal consumption as

$$C_{i,0,t} = \frac{C_{i,J,t}}{P_t^X} W_{i,t} \tag{7}$$

where P_t^X is the wealth of an agent born on date t, so that $C_{i,J,t}/P_t^X$ is the consumption/wealth ratio of new agents of type i. Intuitively, this is just a way to capture a bequest motive, or, more realistically, the marginal utility of wealth at date $t+J$ when the agent is able to reoptimize.

In general log utility does not generate realistically large risk premia. We therefore modify the expectation operator to make agents pessimistic over the distribution of shocks. That pessimism can be thought of as a reduced-form for risk or ambiguity aversion, or simply as behavioral. To generate demand for options, we additionally assume that the pessimism (which determines effective risk aversion) varies over time.[12] Specifically, in the model there will be a single fundamental shock $\varepsilon_t \sim N(0,1)$, and on date t, type-i agents believe that

$$\varepsilon_{t+1} \sim N(-\mu_{i,t}, 1) \tag{8}$$

$$\mu_{i,t} = \phi\mu_{i,t-1} + (1-\phi)\bar{\mu} + \kappa_i\varepsilon_t \tag{9}$$

$\bar{\mu}$ determines average pessimism (and hence the equity premium). Type i's pessimism is procyclical when $\kappa_i > 0$ and countercyclical for $\kappa_i < 0$.

3.1.3 Option specification

For both tractability and simplicity, we model options as quadratic contracts on equity, with payoff

$$X_t^O = \left(\prod_{j=0}^{J-1} R_{t-j} - 1\right)^2 \tag{10}$$

$$\text{where } R_t = \frac{P_t^X + D_t}{P_{t-1}^X} \tag{11}$$

A quadratic contract is equivalent to a particular portfolio of options (Bakshi and Madan (2000)). On any given date, its exposure to the underlying is proportional to the cumulative return since inception. That is why it plays a role for investors in the model. Following positive equity returns, agents with countercyclical pessimism are relatively more optimistic about future returns, increasing their desired allocation to equities. Since they cannot change their exposure after their first period of life, the option is valuable to them for inducing that dynamic reallocation automatically. Agents with countercyclical pessimism will thus tend to demand options, while those with procyclical pessimism (intuitively, agents who think returns show mean reversion rather than momentum) will tend to supply options. Options are valuable because they allow agents to trade on differences of opinion about the persistence of returns.

3.2 Calibration and solution

The equilibrium concept is standard:

Definition 1 *An equilibrium is a set of processes for prices,* $\left\{P_t^X, P_t^O, R_t^B\right\}$*, and the agents'*

[12] Similar to the habit formation of Campbell and Cochrane (1999), or the endogenous time-varying pessimism in Bidder and Dew-Becker (2016) and Maenhout, Vedolin, and Xing (2025), among many others.

demands, $\{C_{i,j,t}, B_{i,j,t}, X_{i,t}, O_{i,t}\}$, *such that markets clear,*

$$\sum_i X_{i,t} = \sum_i 1, \quad \sum_i O_{i,t} = 0, \quad \sum_i \sum_{j=1}^{J} B_{i,j,t} = 0$$

and agents maximize (6).

The assumption here is that the intermediaries are so small that they effectively cannot bear any meaningful amount of options exposure relative to households.[13]

The model is calibrated to the weekly frequency. Since it is meant to match behavior of monthly options, that implies $J = 3$. $\log D_t$ is set to be a Gaussian random walk with innovations that have a standard deviation of $15\%/\sqrt{52}$ to match the volatility of aggregate stock returns. We set $\bar{\mu} = \frac{1}{2}\sqrt{1/52}$ to generate an annualized Sharpe ratio for equities of about 1/2 (the additional risk aversion from log utility makes it slightly higher). We set $\phi = 0.79$, so that pessimism has a half-life of one month. κ for the procyclical and countercyclical agent are set to $\pm\bar{\mu}\sqrt{1-\phi^2}$, which implies that the unconditional standard deviation of pessimism is equal to its mean, and the two types of agents are assumed to have equal mass (and hence equal initial wealth).[14] Finally, the rate of time preference is calibrated to 5% per year in order to generate a plausible risk-free interest rate.

We numerically approximate the model's solution with a fourth-order perturbation using Dynare.[15]

3.3 Results

3.3.1 Equity and option returns

We analyze two regimes for the model: one in which retail investors may buy but not sell options, and a second in which they are free to both buy and sell. When agents can buy but not sell, $O_{i,t} \geq 0$ for all i, meaning that $O_{i,t}$ in fact must equal zero in equilibrium. The shift across the two regimes is meant to capture the decline over time in the frictions that retail investors face in trading options.

In the early regime, when retail can buy but not sell options, the market clearing price is the one such that the maximum option demand across all retail investors is equal to zero

[13] The aggregate financial wealth of US households is currently about \$143 trillion, while the financial assets of broker/dealers (which are almost entirely *not* derivatives) are only \$6.3 trillion (from the *Financial Accounts of the United States*).

[14] Because the agents can only trade equity at birth and because the wealth of the two types is equal at birth, wealth reallocations in subsequent periods of life have no impact on equity prices, even though they do affect consumption.

[15] In order to regularize the problem so the numerical solution is well behaved, we add small quadratic investment costs for bonds and options which are then refunded lump-sum to the households.

(some agents may have negative demand – they would like to sell but cannot). When they can both buy and sell, we look for the standard market-clearing price.

Table 3 reports simulation results for Sharpe and information ratios. The first pair of columns report results for the heterogeneous-agent version of the model where the pessimism of one agent is countercyclical and the other procyclical. Column 1 corresponds to the frictional (early) and column 2 the frictionless (late) regime. Across the two columns, the Sharpe ratio of equities is the same, since options do not change the total equity risk that must be borne. In the frictional regime, traded options are "overpriced", causing them to earn significant negative Sharpe and information ratios, with and without delta hedging.

In the frictionless regime, on the other hand, the options premium goes away. When retail investors can both buy and sell options, then the agents with procyclical pessimism – who trade as though they believe returns are mean-reverting – sell options to the agents with countercyclical pessimism – who trade as though they believe returns are persistent.

3.3.2 Returns on the delta hedge

The fact that options earn negative returns with and without hedging implies that the delta hedge – which is designed to mimic an option return as well as possible – must *not* earn nonzero returns (and alphas). Table 3 reports average returns for the hedge itself, and that is exactly what we see.

To understand the difference in the returns on options and the delta hedge, the first thing to note is that the delta hedge is a dynamic strategy on the underlying, and in particular it is a momentum trade. Specifically, the delta hedge invests an amount in the underlying equal to the option's local exposure, so it is a first-order and hence local hedge. It will have errors when there are jumps in the underlying, when the hedge is not adjusted continuously (which it is not in practice or in the model, which is in discrete time) and when option prices are driven by factors unspanned by the level of the underlying (such as unspanned shocks to volatility).

For both puts and calls (and a quadratic contract), the option delta rises when the value of the underlying rises and falls when it falls. The delta hedge for long options positions is therefore a type of momentum trade. Since effective risk aversion is constant in the model, the equity premium is also constant and a dynamic strategy like that cannot outperform the market. That is what is observed in table 3.

The key driver in the frictional regime is that options, unlike equity, are *not* priced by the average agent. Instead, their price is driven by the demand of the investors who most want to purchase options. The price that sets the demand of those investors to zero is higher than the delta hedge price. That is precisely because these investors cannot perform the dynamic replication on their own.

Overall, the first two columns qualitatively match the empirical results reported so far: there is an early frictional regime in which traded options (both alone and delta-hedged) earn significant negative risk premia and a later period in which that premium disappears.

As a check on the mechanism, the second pair of columns reports results for a version of the model with homogeneous agents with countercyclical pessimism. Since the agents are identical, it is irrelevant whether they can trade or not, so the two columns are identical. As in the heterogeneous case with restricted trade, options again earn negative returns. Intuitively, this is because the agents who effectively drive the price of options in the first column – the countercyclical agents with high option demand – are identical to the agents in the homogeneous case. The difference here, though, is that now delta-hedged options *do not* earn a negative premium, precisely because now option prices reflect the preferences of the average investor, who in this case demand options. The returns on the delta-hedged options therefore reject the single-agent version of the model.

Table 3: Model results

	Heterogeneous		Homogeneous	
	Early	Late	Early	Late
Equity SR	0.65	0.65	0.67	0.67
Traded option SR	-0.22	0.00	-0.22	-0.22
Hedge leg SR	0.00	0.00	-0.28	-0.28
Delta-hedged option SR	-0.37	0.00	0.09	0.09
Traded option IR	-0.29	-0.06	-0.29	-0.29
Hedge leg IR	-0.07	-0.07	-0.35	-0.35
Delta-hedged option IR	-0.37	0.01	0.09	0.09
Corr(option,hedge)	0.81	0.81	0.77	0.77

Description: Table reports various statistics from the calibrated model under two regimes: "Heterogeneous" (intermediaries and investors have different marginal utilities; pre/post-frictions early and late columns) and "Homogeneous" (single representative agent benchmark). SR is the Sharpe ratio and IR is the information ratio. **Interpretation:** Under the heterogeneous model, the negative Sharpe and information ratios on traded and delta-hedged options vanish in the late period, exactly as in the data; under the homogeneous benchmark, the implied premia are unchanged across regimes.

3.4 Empirical implications and tests of the model

3.4.1 Returns on the delta hedge

A key prediction of the model is that the delta hedge should have a CAPM alpha of zero. The fact that the alphas for options and delta-hedged options discussed in section 2.2 are similar immediately implies that the delta hedge must have an alpha of zero.

Typically a delta hedge is constructed based on a model-implied delta that takes as an input the option price (along with the level of the underlying). A key input in that computation is the option-implied volatility (which is just a transformation of the option price). While that is not available before options were traded, we can simply use instead a forecast of volatility, which turns out to produce economically similar results.

Figure 4 plots the information ratio of the delta hedge on a rolling 10-year basis. Before traded options are available, the series is constructed from synthetic straddles going back to 1926; once implied volatilities from traded options are available, the series uses the actual delta-hedge replication. It also plots the information ratio for traded straddles themselves, over the period since 1987 (so that the first 10-year observation is in 1997). From this perspective, there has been an extended period where traded straddles earned strongly negative returns, much more negative than the delta hedge, but the traded straddles have now converged to the delta hedge.

3.4.2 Intermediary gamma exposure and option returns

As noted above, intermediaries here are assumed to be small, with a supply curve with some (very large) slope $1/\varepsilon$, which implies that in the frictional regime when there is net demand from retail investors, the dealers will be short options, while in the frictionless regime their net position will shift to zero. The model therefore implies that dealers' net positions should have moved to zero when option returns did.

We measure intermediaries' net option exposure by their total SPX options gamma. In the model, that is exactly the correct measure, since the quadratic contract in the model has constant (positive) gamma. More generally, gamma measures the local convexity of their position, which locally measures optionality. Using gamma allows us to put options with different strikes and maturity into consistent units and is common in the literature (e.g. Garleanu, Pedersen, and Poteshman (2008), Barbon and Buraschi (2021), and Ni et al. (2021)).

Figure 5 plots net dealer S&P 500 gamma over the period 1996-2025, constructed using the CBOE open-close data.[16]

[16]This is the same dataset used in the previous literature on option frictions, e.g. Garleanu, Pedersen, and

Figure 4: Rolling 10-year information ratios: traded straddle and delta hedge

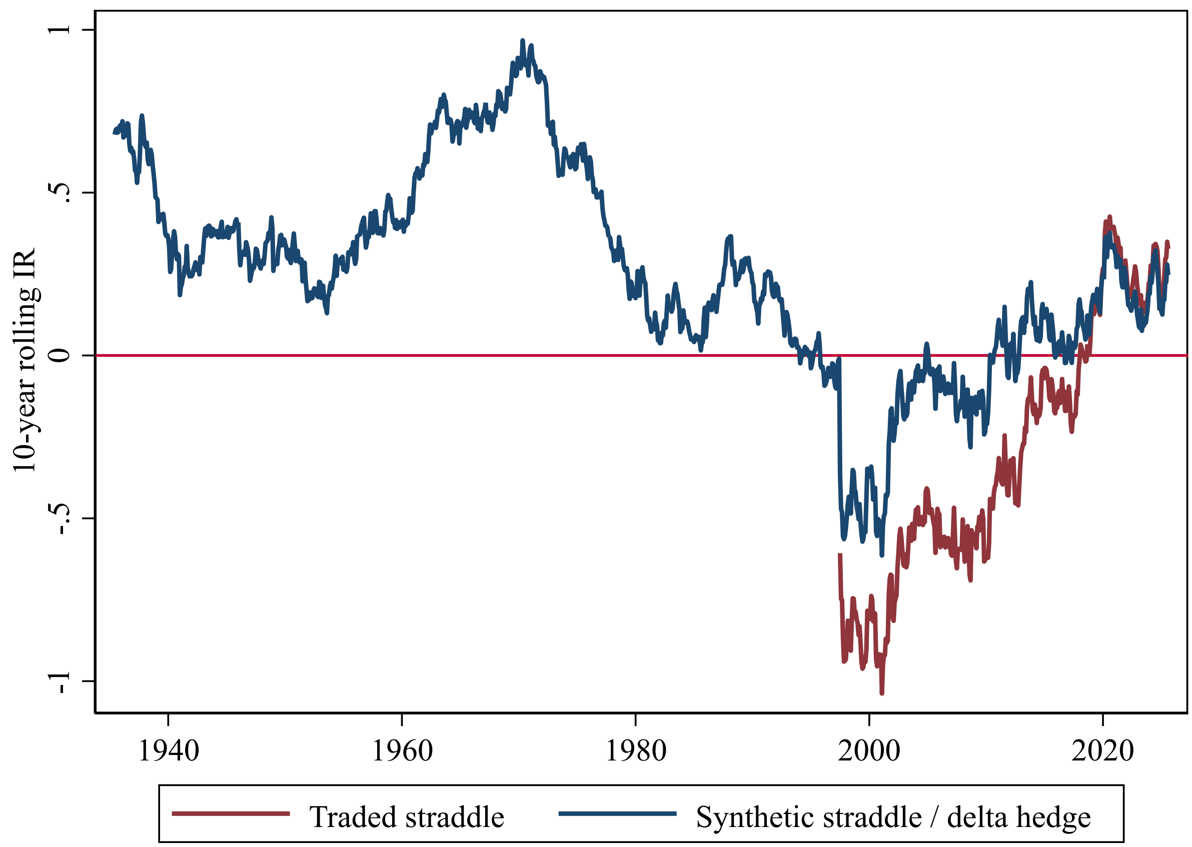

Description: Rolling 10-year CAPM information ratios for the traded ATM straddle (maroon) and the delta-hedge replicating portfolio (navy), with returns scaled by the initial option price. The replicating portfolio is constructed using implied volatilities from traded options starting when they are available (delta-hedge), and using forecasted volatilities before then (synthetic straddles). The first point in the figure using IVs is 1997:07, using the 10-years since 1987:08 to compute the IR. For each series, market betas are estimated using the full sample available and used to compute beta-adjusted returns; information ratios are the ratio of the rolling 10-year mean to the rolling 10-year standard deviation of beta-adjusted returns, annualized. Sample period and observations: synthetic ATM straddles are computed from daily S&P 500 returns from December 1926 through August 2025 (N = 1,189 monthly end-of-month rolling-window observations, of which the first valid 10-year window ends in December 1936); traded ATM straddle returns are monthly from 1987:08 through 2025:08 (N = 457, with the rolling 10-year IR available from 1997:07 onwards, N = 339 rolling-window points). **Interpretation:** The two information ratios were sharply negative through the early 2010s but have converged toward zero in recent years; the long synthetic series shows that the delta-hedge replicating portfolio never had consistently negative average alphas since 1926.

Dealer gamma was significantly negative, as has been noted widely in previous work, until around the global financial crisis. However, since the GFC, net dealer gamma has trended to zero or even positive. The timing of that switch is highly similar to the timing of the shift in the premium on options. As discussed above, we take August 2012 as the time of the break because it is the last month where the centered 12-month moving average of dealer gamma was negative. The specific timing of the date is not central, and obviously there is no clear one-time break. Indeed some of the tests in table 2 date the break to 2008 or 2010, which both also could plausibly represent the shift point in dealer gamma. The aim is simply to find a date that represents the point where dealer gamma was no longer negative in the way that it was in the early parts of the sample.

Table 4: Dealer gamma and option risk premia (ATM straddle)

	Straddle	Hedge leg	Delta-hedged straddle
Dealer net gamma	0.44	-0.31	0.61**
	(0.87)	(0.83)	(0.28)
Spot-future basis volatility	-2.12***	-1.07	-1.03***
	(0.78)	(0.76)	(0.28)

Description: Table reports the coefficient of a regression of information ratios (annualized rolling beta-hedged means scaled by rolling residual standard deviations) of traded ATM straddles, the delta-hedge leg, and delta-hedged ATM straddles (with returns scaled by the initial option price) on the lagged net gamma of dealers and on the lagged volatility of the spot-future basis. Both regressors are filtered through an exponentially-weighted moving average. Estimation is by maximum likelihood with a Student-t error distribution to limit the influence of outliers; standard errors are in parentheses (*, **, *** denote significance at 10%, 5%, 1%). Sample period: monthly, 1996:01–2025:08, $N = 356$ monthly observations. **Interpretation:** The information ratio of the delta-hedged straddle loads positively and significantly on lagged net dealer gamma, exactly as predicted by the model: as dealers' net gamma rose toward zero, the delta-hedged option premium also rose toward zero. Information ratios of traded and delta-hedged straddles also load negatively on the lagged volatility of the spot–futures basis, consistent with basis risk being an additional friction priced into options.

To further understand the implications of the changes in dealer positions, table 4 reports the coefficients from a regression of straddle information ratios onto lagged dealer gamma (exponentially-weighted to remove the very high-frequency fluctuations). To reduce the influence of outliers, the estimation is performed via maximum likelihood with a student-t

Poteshman (2008), as well as Chen, Joslin and Ni (2019) and Constantinides and Lian (2015). The dataset classifies, for each option, the total daily buy and sell orders by type of entity (customer, firm, and broker dealer). We compute the total gamma bought and sold each day by intermediaries (defined, as in Garleanu, Pedersen, and Poteshman (2008), as entities that are neither customers or firms) by combining this data with the gamma of each option from Optionmetrics. We use options with 10 to 180 days maturity.

error distribution, but results are similar using OLS regression. The table shows that the delta-hedged straddle information ratio loads positively on the intermediaries' net gamma, as predicted by the model: as intermediaries started bearing less risk over time (gamma became less negative), the delta-hedged straddle alpha became less negative as well (generating a positive coefficient in this regression). This table therefore provides direct evidence in support of the model presented in this section.

3.4.3 Additional potential drivers of option premia from a richer model

The analysis so far simply uses dealer positions as a way to measure asymmetry in retail options demand under the assumption that the dealers' supply curve is effectively vertical at zero. However, there is a literature following Garleanu, Pedersen, and Poteshman (GPP; 2008) that takes the opposite approach – it treats the excess demand from retail investors as exogenous and instead focuses on the determinants of the slope of dealers' supply curve.

In order to generate further implications of excess retail options demand that we can take to the data, appendix B extends GPP's model to add three additional realistic frictions that dealers may face: unspanned risk, basis risk, and hedging costs (see proposition 2). Unspanned risk is any risk left after discrete hedging with the underlying (e.g. jump risk and unspanned volatility). Basis risk represents the deviation between the hedging instrument – e.g. S&P 500 futures – and the actual underlying index. And hedging costs represent the cost to dealers due to the actual cost of synthesizing a hedge, such as transaction costs or price pressure. Figure 6 examines how each of these may have changed over time.

In the GPP model, all three of those factors increase the price that dealers demand in order to supply options. Changes in those frictions over time are therefore another possible explanation of the decline in the overpricing of options.

The top two panels of figure 6 show that trading costs, measured both by posted and effective spreads (Roll (1984)) declined as option premia shrunk. So, in addition to dealers bearing less option risk over time, they also pay smaller costs to hedge the risk they do bear.

We measure basis risk empirically from the gap between the level of the S&P 500 index and the futures price. The middle-left panel of figure 6 plots the three-month rolling standard deviation of that gap. The y-axis is again on a $\log_{10}$ scale. Over time, basis risk has fallen by about an order of magnitude. While there is a large decline early in the sample, similar to trading frictions, basis risk seems to settle at its current level around the early 2000's.

Finally, figure 6 plots three measures of unhedgeable risk. The first is the 10-year rolling standard deviation of the delta-hedged ATM straddle return – i.e., the difference between the traded straddle return and the synthetic straddle return, which measures the gap between the return on the option and the hedge. It shows no clear trend, particularly once the

1987 crash is no longer included in the moving average. Second, following Bollerslev et al. (2009), figure 6 plots the difference between quadratic and bipower variation, which is a measure of realized jump variation. It rose during the 2008 financial crisis, and has been lower subsequently, but again does not have a clear trend. Finally, unhedgeable risk is, more broadly, driven by higher moments in returns, so the bottom-right panel of figure 6 plots the measure of S&P 500 return skewness developed in Neuberger (2012). Realized skewness has, over time, trended consistently more negative (implied skewness does the same; see the CBOE's SKEW index and Dew-Becker (2024)). So if higher moments drove the options premium, it should have grown instead of shrunk over time. The decline that we observe in the overpricing of options cannot be driven by declines in jump or higher moment risk, as measured here, since none of those sources of risk have declined. Instead, consistent with the model, the decline in the premium can be driven by a reduction in how much of this risk is borne by intermediaries, and in the frictions they face.

Overall, figure 6 shows that hedging costs and basis risk have declined – indicating a reduction in frictions faced by intermediaries – while unhedgeable risk does not appear to have shrunk.

The second row of table 4 shows that information ratios of traded straddles and of delta-hedged straddles both moved negatively relative to basis risk, measured as the lagged exponentially-weighted average of the standard deviation of the basis pictured above: as the basis on average became less volatile over time, traded and delta-hedged straddle information ratios became less negative, again in line with the predictions of the theory.

4 Broader implications

This section discusses three additional aspects of the results: their relationship with changes over time in performance of other investment strategies; implications for portfolio choice; and implications for the magnitude of the variance risk premium.

4.1 The variance risk premium

The most direct implication of the results is that the variance risk premium has declined. That is really the literal interpretation of the results using the VIX portfolio as a synthetic variance swap. They say that the gap between realized and implied variance has shrunk.

To quantify that, consider a simple CAPM type regression:

$$\left(\mathrm{RV}_t - \mathrm{VIX}_{t-1}^2\right)/\mathrm{VIX}_{t-1}^2 = \alpha_{t-1} + \beta_{\mathrm{RV}} r_{m,t} + \varepsilon_t \tag{12}$$

where RV_t here is realized variance and $r_{m,t}$ is the market return, and the dependent variable is the variance-swap return per dollar invested. The top panel of figure 7 decomposes rolling 10-year averages of $\left(\text{RV}_t - \text{VIX}_{t-1}^2\right)/\text{VIX}_{t-1}^2$ into contributions coming from $\beta_{\text{RV}} r_{m,t}$ and $\alpha_{t-1} + \varepsilon_t$. Since ε_t has a conditional mean of zero by definition, we refer to the average of $\alpha_{t-1} + \varepsilon_t$ as just the rolling alpha.

Not surprisingly, the beta component of $\left(\text{RV}_t - \text{VIX}_{t-1}^2\right)/\text{VIX}_{t-1}^2$ has been fairly stable over time, whereas there has been a huge shift in the alpha component, to the point that it is now close to zero. To see what that implies for the relationship between the VIX and realized volatility, suppose annual realized variance is $(15\%)^2$. Inserting the first rolling 10-year mean for $\left(\text{RV}_t - \text{VIX}_{t-1}^2\right)/\text{VIX}_{t-1}^2$, the corresponding average VIX would be 18.7. That is, the VIX overstates realized volatility by 25%. But with the alpha near zero in the final rolling 10-year window, the corresponding value of the VIX is only 16.4, overstating volatility instead by only 10%. In other words, just from the decline in the variance risk premium, we would expect the VIX to now average 2.2 points lower than in the period when there was a large variance risk premium.

A simpler way to see that is to simply compare the VIX to realized volatility, which the bottom panel of figure 7 does. It plots rolling 10-year averages of the VIX against the rolling 10-year *standard deviation* of market returns. There will be a slight bias down in the realized standard deviation since it is a square root of a sample statistic. Nevertheless, what we see is not just that the gap between the two series has shrunk, but that it has actually been near zero for 17 years now. The VIX over the past 17 years has been nearly an unbiased predictor of realized volatility in the S&P 500, at least on average.

These results are consistent with those in Heston, Jacobs, and Kim (2022), who also find that there is a negative variance risk premium, but that it cannot be distinguished from simple market (beta) risk.[17]

4.2 Declining factor premia

On some level, finding that a factor risk premium has declined is possibly the least surprising outcome. In addition to the well known evidence that individual predictive relationships tend to underperform out of sample (McLean and Pontiff (2016), Linnainmaa and Roberts (2018), and Marrow and Nagel (2024), among many others), Green, Hand, and Zhang (2017) and Chen and Velikov (2023) also show that broad averages of equity risk factors have also significantly declined in performance since the early 2000's. A simple exercise is to just download data from Kenneth French on the most prominent equity anomalies – size, value

[17]See also Gao, He, and Hu (2026), who show that demand-driven variation in firm-level IVs has declined over time.

and momentum – and note that none of the three has had positive returns in over 20 years at this point.

The fact that the premia on these other strategies decline approximately contemporaneously with those for options is telling. Agarwal and Naik (2004) and Jurek and Stafford (2015), among others, provide evidence that hedge fund strategies are in the end very similar to simple S&P 500 put selling. At the same time, value and momentum, among other strategies, are known to have been important strategies for quantitative equity hedge funds (Akbas et al. (2015) and Chen, Da, and Huang (2019)). The rise in hedge funds, as measured by their assets under management, is approximately contemporaneous with the decline in these anomalies (e.g. see Chordia, Subrahmanyam, and Tong (2014) and Hanson and Sunderam (2014)).

Figure 8 plots cumulative returns on the beta-hedged short delta-hedged ATM straddle, value, momentum, and currency carry.[18] All four series visibly flatten out – due to the monthly returns moving to zero – in the early 2000's.

4.3 Portfolio construction

Driessen and Maenhout (2007) analyze the benefits to portfolio optimization from giving an investor access to options strategies (see also Santa-Clara and Saretto (2009)). Their data is for 1987–2001, the main period in which options premia were large. They find Sharpe ratios ranging from -0.2 to -0.4, but do not report betas or alphas, which are what is relevant when considering adding options to a position in just the market. That said, straddles have a beta fairly close to zero, so their Sharpe ratio is reasonably close to the information ratio. Driessen and Maenhout (2007) find a Sharpe ratio for straddles of -0.21, whereas we find -0.74 over the same sample. Subsequent to 2001, though, the Sharpe ratio of straddles in our data is -0.09, and after 2012 it is +0.12. The alpha also shifts from negative to positive.

Depending on the coefficient of relative risk aversion, Driessen and Maenhout (2007) find that a CRRA investor would put a weight of up to -0.58 (-0.47 for a log investor) on at-the-money straddles in their calibration of the returns process. If straddles now have an alpha of zero, then their weight will be approximately zero (precisely zero for a mean-variance optimizer), and if the alpha is actually positive, the weight will likely be positive (again, that is guaranteed for a mean-variance optimizer).

Part of the reason to focus on alphas is that they determine portfolio exposures for mean-variance optimizers. The past literature finding a large variance risk premium was important partly because it had major implications for portfolio choice. The large VRP implied that

[18] Value and momentum are from Ken French and currency carry from AQR (https://www.aqr.com/Insights/Datasets/Century-of-Factor-Premia-Monthly).

investors would get a huge benefit from being short volatility (e.g. by selling straddles). And in fact there was a large growth in short volatility exchange traded products. If the alpha on volatility is now zero, those products no longer have a role to play in a typical investor's portfolio.

5 Conclusion

The conventional wisdom in asset pricing is that there is a large variance risk premium for the S&P 500 – larger than can be accounted for by exposure to the market alone, and that options earn large negative alphas because they are exposed to volatility and jump risk. This paper replicates past findings that options historically earned negative returns and CAPM alphas, but since around 2012 there is no longer evidence for those negative premia.

Though the paper does not go into great depth on the theory side, it does offer an explanation, which is that the decline in SPX options premia can be explained by declining asymmetry in dealer portfolios and the hedging costs they face.

More broadly, though, the disappearance of option premia occurred contemporaneously with the decline in many other major anomalies in asset prices. In that sense, the results documented here should perhaps not be puzzling at all – they are just one more example of the performance of a trading strategy decaying after it is announced in a finance journal. That is a good thing: markets are getting more efficient. At a deep level, the implication of the results is that now it is much less expensive for investors to hedge deep losses in the aggregate stock market than it used to be.

References

Agarwal, Vikas and Narayan Y Naik, "Risks and portfolio decisions involving hedge funds," *The Review of Financial Studies*, 2004, *17* (1), 63–98.

Akbas, Ferhat, Will J Armstrong, Sorin Sorescu, and Avanidhar Subrahmanyam, "Smart money, dumb money, and capital market anomalies," *Journal of Financial Economics*, 2015, *118* (2), 355–382.

Andersen, Torben G, Nicola Fusari, and Viktor Todorov, "The pricing of tail risk and the equity premium: Evidence from international option markets," *Journal of Business & Economic Statistics*, 2020, *38* (3), 662–678.

Andrews, Donald W.K., "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 1993, *61(4)*, 821–856.

Andrews, Donald WK and Werner Ploberger, "Optimal tests when a nuisance parameter is present only under the alternative," *Econometrica: Journal of the Econometric Society*, 1994, pp. 1383–1414.

Backus, David, Mikhail Chernov, and Ian Martin, "Disasters implied by equity index options," *The Journal of Finance*, 2011, *66* (6), 1969–2012.

Bakshi, Gurdip and Dilip Madan, "Spanning and derivative-security valuation," *Journal of financial economics*, 2000, *55* (2), 205–238.

___ **and Nikunj Kapadia**, "Delta-Hedge Gains and the Negative Market Volatility Risk Premium," *The Review of Financial Studies*, 2003, *16(2)*, 527–566.

Bao, Jack, Jun Pan, and Jiang Wang, "The illiquidity of corporate bonds," *The Journal of Finance*, 2011, *66* (3), 911–946.

Barbon, Andrea and Andrea Buraschi, "Gamma Fragility," 2021. Working paper.

Bates, David S, "The market for crash risk," *Journal of Economic Dynamics and Control*, 2008, *32* (7), 2291–2321.

___ , "Empirical Option Pricing Models," *Annual Review of Financial Economics*, 2022, *14*, 369–389.

Bidder, Rhys and Ian Dew-Becker, "Long-Run Risk is the Worst-Case Scenario," *The American Economic Review*, September 2016, *106* (9), 2494–2527.

Bollen, Nicolas PB and Robert E Whaley, "Does net buying pressure affect the shape of implied volatility functions?," *The Journal of Finance*, 2004, *59* (2), 711–753.

Bollerslev, Tim and Viktor Todorov, "Tails, fears, and risk premia," *Journal of Finance*, 2011, *66(6)*, 2165–2211.

___ **, Uta Kretschmer, Christian Pigorsch, and George Tauchen**, "A discrete-time model for daily S & P500 returns and realized variations: Jumps and leverage effects," *Journal of Econometrics*, 2009, *150* (2), 151–166.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, "Model Specification and Risk Premia: Evidence from Futures Options," *The Journal of Finance*, 2007, *62(3)*, 1453–1490.

———, ———, and ———, "Understanding Index Option Returns," *The Review of Financial Studies*, 2009, *22(11)*, 4493–4529.

Brunnermeier, Markus K, Stefan Nagel, and Lasse H Pedersen, "Carry trades and currency crashes," *NBER macroeconomics annual*, 2008, *23* (1), 313–348.

Büchner, Matthias and Bryan Kelly, "A factor model for option returns," *Journal of Financial Economics*, 2022, *143* (3), 1140–1161.

Campbell, John Y. and John H. Cochrane, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 1999, *107* (2), 205–251.

Carr, Peter and Liuren Wu, "Variance Risk Premiums," *Review of Financial Studies*, 2009, *22(3)*, 1311–1341.

Chambers, Donald R, Matthew Foy, Jeffrey Liebner, and Qin Lu, "Index option returns: Still puzzling," *The Review of Financial Studies*, 2014, *27* (6), 1915–1928.

Chen, Andrew Y and Mihail Velikov, "Zeroing in on the expected returns of anomalies," *Journal of Financial and Quantitative Analysis*, 2023, *58* (3), 968–1004.

Chen, Hui, Scott Joslin, and Sophie Xiaoyan Ni, "Demand for crash insurance, intermediary constraints, and risk premia in financial markets," *The Review of Financial Studies*, 2019, *32* (1), 228–265.

———, Winston Wei Dou, and Leonid Kogan, "Measuring "dark matter" in asset pricing models," *The Journal of Finance*, 2024, *79* (2), 843–902.

Chen, Yong, Zhi Da, and Dayong Huang, "Arbitrage trading: The long and the short of it," *The Review of Financial Studies*, 2019, *32* (4), 1608–1646.

Chordia, Tarun, Avanidhar Subrahmanyam, and Qing Tong, "Have capital market anomalies attenuated in the recent era of high liquidity and trading activity?," *Journal of Accounting and Economics*, 2014, *58* (1), 41–58.

Constantinides, George M and Lei Lian, "The supply and demand of S&P 500 put options," Technical Report, National Bureau of Economic Research 2015.

———, Jens Carsten Jackwerth, and Alexi Savov, "The puzzle of index option returns," *Review of Asset Pricing Studies*, 2013, *3* (2), 229–257.

Coval, Joshua D. and Tyler Shumway, "Expected Option Returns," *The Journal of Finance*, 2001, *56(3)*, 983–1009.

Darling, D. A., "The Kolmogorov-Smirnov, Cramér-von Mises Tests," *The Annals of Mathematical Statistics*, 1957.

Dehling, Herold, Roland Fried, and Martin Wendler, "A robust method for shift detection in time series," *Biometrika*, 2020, *107* (3), 647–660.

Dew-Becker, Ian, "Real-time forward-looking skewness over the business cycle," *Review of Economic Dynamics*, 2024, *54*, 101233.

____, Stefano Giglio, Anh Le, and Marius Rodriguez, "The price of variance risk," *Journal of Financial Economics*, 2017, *123* (2), 225 – 250.

Drechsler, Itamar, "Uncertainty, Time-Varying Fear, and Asset Prices," *The Journal of Finance*, 2013, *68* (5), 1843–1889.

____ and Amir Yaron, "What's Vol Got to Do with it?," *The Review of Financial Studies*, 2011, *24(1)*, 1–45.

Driessen, Joost and Pascal Maenhout, "An empirical portfolio perspective on option pricing anomalies," *Review of Finance*, 2007, *11* (4), 561–603.

Fournier, Mathieu, Kris Jacobs, and Piotr Orłowski, "Modeling conditional factor risk premia implied by index option returns," *The Journal of Finance*, 2024, *79* (3), 2289–2338.

Frazzini, Andrea and Lasse Heje Pedersen, "Embedded leverage," *The Review of Asset Pricing Studies*, 2022, *12* (1), 1–52.

Gabaix, Xavier, "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 2012, *127(2)*, 645–700.

Gao, Chao, Jia He, and Grace Xing Hu, "Maturity Structure and Its Impact on Option Prices," 2026. Working paper.

Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M Poteshman, "Demand-based option pricing," *The Review of Financial Studies*, 2008, *22* (10), 4259–4299.

Green, Jeremiah, John RM Hand, and X Frank Zhang, "The characteristics that provide independent information about average US monthly stock returns," *The Review of Financial Studies*, 2017, *30* (12), 4389–4436.

Haddad, Valentin and Tyler Muir, "Do intermediaries matter for aggregate asset prices?," *The Journal of Finance*, 2021, *76* (6), 2719–2761.

Han, Bing, "Investor sentiment and option prices," *The Review of Financial Studies*, 2008, *21* (1), 387–414.

Hanson, Samuel G and Adi Sunderam, "The growth and limits of arbitrage: Evidence from short interest," *The Review of Financial Studies*, 2014, *27* (4), 1238–1286.

Heston, Steven, Kris Jacobs, and Hyung Joo Kim, "Exploring Risk Premia, Pricing Kernels, and No-Arbitrage Restrictions in Option Pricing Models," 2022. Working paper.

Hodges, J.L. and E.L. Lehmann, "Estimates of Location Based on Rank Tests," *The Annals of Mathematical Statistics*, 1963, *34* (2), 598–611.

Jackwerth, Jens Carsten, "Recovering risk aversion from option prices and realized returns," *The Review of Financial Studies*, 2000, *13* (2), 433–451.

Jurek, Jakub W and Erik Stafford, "The cost of capital for alternative investments," *The Journal of Finance*, 2015, *70* (5), 2185–2226.

Konstantinidi, Eirini and George Skiadopoulos, "How does the market variance risk premium vary over time? Evidence from S&P 500 variance swap investment returns," *Journal of Banking & Finance*, 2016, *62*, 62–75.

Linnainmaa, Juhani T and Michael R Roberts, "The history of the cross-section of stock returns," *The Review of Financial Studies*, 2018, *31* (7), 2606–2649.

Longstaff, Francis A, Jun Pan, Lasse H Pedersen, and Kenneth J Singleton, "How sovereign is sovereign credit risk?," *American Economic Journal: Macroeconomics*, 2011, *3* (2), 75–103.

Maenhout, Pascal J, Andrea Vedolin, and Hao Xing, "Robustness and dynamic sentiment," *Journal of Financial Economics*, 2025, *163*, 103953.

Marrow, Benjamin and Stefan Nagel, "Real-Time Discovery and Tracking of Return-Based Anomalies," 2024. Working paper.

McLean, R David and Jeffrey Pontiff, "Does academic research destroy stock return predictability?," *The Journal of Finance*, 2016, *71* (1), 5–32.

Muravyev, Dmitriy and Xuechuan Charles Ni, "Why do option returns change sign from day to night?," *Journal of Financial Economics*, 2020, *136* (1), 219–238.

Nagel, Stefan, "Evaporating liquidity," *The Review of Financial Studies*, 2012, *25* (7), 2005–2039.

Neuberger, Anthony, "Realized skewness," *The Review of Financial Studies*, 2012, *25* (11), 3423–3455.

Ni, Sophie X, Neil D Pearson, Allen M Poteshman, and Joshua White, "Does option trading have a pervasive impact on underlying stock prices?," *The Review of Financial Studies*, 2021, *34* (4), 1952–1986.

Roll, Richard, "A simple implicit measure of the effective bid-ask spread in an efficient market," *The Journal of finance*, 1984, *39* (4), 1127–1139.

Santa-Clara, Pedro and Alessio Saretto, "Option Strategies: Good Deals and Margin Calls," *Journal of Financial Markets*, 2009, *12*, 391–417.

Schreindorfer, David, "Macroeconomic tail risks and asset prices," *The Review of Financial Studies*, 2020, *33* (8), 3541–3582.

Seo, Sang Byung and Jessica A Wachter, "Option prices in a model with stochastic disaster risk," *Management Science*, 2019, *65* (8), 3449–3469.

A Simulation results for break tests

B Theoretical results for a richer intermediary model

This section studies a simple extension of the model of Garleanu, Pedersen, and Poteshman (GPP; 2008) to help clarify how various frictions can affect option prices when markets are segmented. The only addition to their framework is to allow for transaction costs and index-futures basis risk.

B.1 Setup

The dealers in GPP are assumed to have time-additive CARA preferences over consumption with risk aversion γ. There is a constant gross risk-free rate R_f. The underlying index has an exogenous excess return R^I_{t+1}. We consider a simplified version of the model where there is a single option traded that has some price P_t. Its excess dollar payoff is then $R^O_{t+1} = P_{t+1} - R_f P_t$.

The key equation in the model is the dynamic budget constraint,

$$W_{t+1} = (W_t - C_t) R_f + D_t R^O_{t+1} + F_t R^F_{t+1} - \frac{\kappa}{2} F_t^2 \tag{13}$$

W_t is wealth and C_t consumption. The risk-free rate, R_f is constant for simplicity, R^F_{t+1} is the excess payoff on index futures, and R^O_{t+1} is the excess dollar payoff on the options. The dealers endogenously choose consumption and the allocations to derivatives and futures D_t and F_t, respectively.

We add two frictions: a quadratic trading cost, $\frac{\kappa}{2} F_t^2$, and a wedge between the futures payoff, R^F_{t+1}, and the underlying index return, R^I_{t+1},

$$R^F_{t+1} = R^I_{t+1} + z_{t+1} \tag{14}$$

z_{t+1} represents *basis risk*. Ideally the dealers would like to hedge the options they trade with the underlying, like the S&P 500. But the S&P 500 is not itself directly tradable (except at significant cost by buying 500 stocks). Instead, dealers must buy futures (or ETFs or other instruments) whose price is not guaranteed to perfectly track the index. z_{t+1} captures the risk associated with imperfect tracking.[19]

[19] The S&P 500 index is the underlying for CBOE options, but not CME futures options. For the futures options, an interpretation of basis risk would be that intermediaries price options based on a model for the underlying, meaning that deviations of the futures price from the index create risk.

While the dealers choose D_t, markets must clear, meaning that in equilibrium their choice of D_t must perfectly offset the demand from retail investors, which GPP take to be exogenous.

Dealers/intermediaries maximize discounted utility over consumption C_t,

$$E_t \sum_{j=0}^{\infty} \rho^j \left(-\gamma^{-1}\right) \exp\left(-\gamma C_{t+j}\right) \tag{15}$$

subject to a transversality condition and budget constraint, which is

$$W_{t+1} = \left(W_t - C_t\right) R_f + D_t R_{t+1}^O + F_t R_{t+1}^F - \frac{\kappa}{2} F_t^2 \tag{16}$$

where W_t is wealth. The intermediaries optimize over D_t, F_t, and C_t subject to the budget constraint and taking the returns as given.

It is assumed that the futures contract on the underlying that the dealers trade is available in infinite supply. For the options, there is some exogenous demand from outside investors, d_t, and the market clearing condition is $D_t + d_t = 0$.

B.2 Predictions

In the model, intermediaries hedge their options each period with a position in the underlying. The optimal position, in the absence of any frictions, is denoted by β_t^I (which is simply the local sensitivity of option payoffs to the underlying index). The unhedgeable risk is defined as

$$\sigma_{\varepsilon,t}^2 \equiv \text{var}_t^d \left(R_{t+1}^O - \beta_t^I R_{t+1}^I\right) \tag{17}$$

where var_t^d is a variance taken under the intermediaries' pricing measure d based on date-t information.

The model's key prediction is for the sensitivity of option prices, P_t^O, to demand:

Proposition 2 *Up to first order in the transaction cost κ and the index-futures basis risk* $\text{var}_t^d\left(z_{t+1}\right)$,

$$\frac{\partial P_t^O}{\partial D_t} = -\frac{\gamma\left(R_f - 1\right)}{R_f^2} \left(\underbrace{\sigma_{\varepsilon,t}^2}_{\textit{Unhedgeable risk}} + \underbrace{\left(\beta_t^I\right)^2 \text{var}_t^d\left(z_{t+1}\right)}_{\textit{Basis risk}} \right) - \underbrace{\frac{\kappa}{R_f}\left(\beta_t^I\right)^2}_{\textit{Imperfect hedging}} \tag{18}$$

The sensitivity is proportional to risk aversion, γ, and has three terms.

The first component, $\sigma^2_{\varepsilon,t}$, is the unhedgeable risk from (17). Dealers hedge by taking positions exposed to the underlying, but since options have nonlinear exposure, that hedge is inevitably imperfect, due to discrete hedging, jumps, and unspanned volatility. The synthetic options studied in our empirical analysis exactly map into the hedge that the dealers use here – they are updated discretely and inherit risk from deviations between the discrete replication and the traded option payoff.

The second term represents basis risk. When there are larger random gaps between the hedging instrument and the true underlying index, dealers face greater risk and thus demand larger premia. Finally, the third term arises due to the quadratic trading cost, κ, which causes dealers to hedge incompletely, further raising their risk from holding derivatives.

In the context of the general theoretical analysis in section 3, this is a model in which traded options are not priced by the marginal utility of retail equity investors, but instead by that of dealers. And the exogenous option demand, since it must be borne only by dealers, drives option prices up, creating negative CAPM alphas.

B.3 Proof of proposition 2

Lemma 3 *In this model, assets are priced under a probability measure d which is equal to the measure P multiplied by the factor $\frac{\exp(-k(W_{t+1}+G(d_{t+1},X_{t+1})))}{E_t[\exp(-k(W_{t+1}+G(d_{t+1},X_{t+1})))]}$. In addition,*

$$\kappa F_t = E_t^d \left[R^F_{t+1}\right] \quad (19)$$

$$P_t = R_f^{-1} E_t^d P_{t+1} \quad (20)$$

where P_t is the price of the option (equivalently, $0 = E_t^d R^O_{t+1}$).

Proof. The value function and budget constraint satisfy

$$V_t = \max_{C_t, D_t, F_t} -\gamma^{-1} \exp\left(-\gamma C_t\right) + \rho E_t V_{t+1} \quad (21)$$

$$W_{t+1} = \left(W_t - C_t\right) R_f + D_t \left(P_{t+1} - R_f P_t\right) + F_t R^F_{t+1} - \frac{\kappa}{2} F_t^2 \quad (22)$$

Now guess that

$$V_t = -k^{-1} \exp\left(-k\left(W_t + G_t\right)\right)$$

for some variable G_t that is exogenous to the dealers, and where

$$k = \gamma \frac{R_f - 1}{R_f} \quad (23)$$

We have

$$\begin{aligned}\frac{\partial}{\partial W_t} V_t &= -kV_t \quad (24)\\ \text{and } \frac{dW_{t+1}}{dC_t} &= -R_f \quad (25)\end{aligned}$$

So then the FOC for consumption under this guess is

$$0 = \exp(-\gamma C_t) + kR_f \rho E_t V_{t+1} \quad (26)$$

Noting that

$$\begin{aligned}V_t &= -\gamma^{-1}\exp(-\gamma C_t) + \rho E_t V_{t+1} \quad (27)\\ \rho E_t V_{t+1} &= V_t + \gamma^{-1}\exp(-\gamma C_t) \quad (28)\end{aligned}$$

We have

$$\exp(-\gamma C_t) = \exp(-k(W_t + G_t)) \quad (29)$$

Now consider the FOC with respect to F_t. First,

$$\frac{dW_{t+1}}{dF_t} = R^F_{t+1} - \kappa F_t \quad (30)$$

And hence the FOC is

$$\begin{aligned}0 &= \rho E_t\left[\exp(-k(W_{t+1} + G_{t+1}))\left(R^F_{t+1} - \kappa F_t\right)\right] \quad (31)\\ \kappa F_t &= E^d_t\left[R^F_{t+1}\right] \quad (32)\end{aligned}$$

where E^d is the expectation under the dealers' pricing measure, which is the physical measure distorted by the factor

$$\frac{\exp(-k(W_{t+1} + G_{t+1}))}{E_t[\exp(-k(W_{t+1} + G_{t+1}))]} \quad (33)$$

Next, for D_t,

$$\frac{dW_{t+1}}{dD_t} = R^O_{t+1} \quad (34)$$

So then

$$0 = \rho E_t\left[\exp\left(-k\left(W_{t+1}+G_{t+1}\right)\right)R^O_{t+1}\right] \quad (35)$$

$$0 = R_f^{-1}E_t^d R^O_{t+1} \quad (36)$$

It is straightforward to get a recursion for G_t by following the derivation in GPP. ■

Lemma 4 *The effect of options demand on prices is*

$$\frac{\partial P_t}{\partial D_t} = -\frac{\gamma\left(R_f-1\right)}{R_f^2}\operatorname{cov}_t^d\left(R^O_{t+1} - R^F_{t+1}\beta^F_t\frac{\operatorname{var}_t^d\left(R^F_{t+1}\right)}{\kappa k^{-1}+\operatorname{var}_t^d\left(R^F_{t+1}\right)}, R^O_{t+1}\right) \quad (37)$$

where

$$\hat{\beta}^F_t \equiv \beta^F_t\frac{\operatorname{var}_t^d\left(R^F_{t+1}\right)}{\kappa k^{-1}+\operatorname{var}_t^d\left(R^F_{t+1}\right)} \quad (38)$$

$$\beta^F_t \equiv \frac{\operatorname{cov}_t^d\left(R^F_{t+1}, R^O_{t+1}\right)}{\operatorname{var}_t^d\left(R^F_{t+1}\right)} \quad (39)$$

Proof. Based on the analysis from the previous proof, the pricing kernel can be written as

$$m^d_{t+1} \equiv \frac{\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)}{R_fE_t\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)} \quad (40)$$

Differentiate m^d_{t+1} with respect to D_t to get

$$\frac{\partial m^d_{t+1}}{\partial D_t} = \frac{-k\left(R^O_{t+1}+R^F_{t+1}\frac{\partial F_t}{\partial D_t}\right)\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)}{R_fE_t\left[\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)\right]} \quad (41)$$

$$-\frac{\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)}{\left(R_fE_t\left[\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)\right]\right)^2} \quad (42)$$

$$\times E_t\left[-kR_f\left(R^O_{t+1}+R^F_{t+1}\frac{\partial F_t}{\partial D_t}\right)\exp\left(-k\left(F_tR^F_{t+1}+D_tR^O_{t+1}+G_{t+1}\right)\right)\right] \quad (43)$$

$$= -km^d_{t+1}\left(A_{t+1}-E_t^d\left[A_{t+1}\right]\right)$$

where

$$A_{t+1} \equiv \left(R^O_{t+1}+R^F_{t+1}\frac{\partial F_t}{\partial D_t}\right) \quad (44)$$

and expectation operator

$$E_t^d\left[\cdot\right] \equiv R_fE_t\left[m^d_{t+1}\cdot\right] \quad (45)$$

Next, we differentiate the first-order condition for F_t with respect to D_t. Since $E_t^d[\cdot] = R_f E_t\left[m_{t+1}^d \cdot\right]$, the first-order condition is $\kappa F_t = R_f E_t\left[m_{t+1}^d R_{t+1}^F\right]$ and hence

$$\begin{aligned} \kappa \frac{\partial F_t}{\partial D_t} &= R_f E_t\left[\frac{\partial m_{t+1}^d}{\partial D_t} R_{t+1}^F\right] && (46) \\ &= R_f E_t\left[-k\left(A_{t+1} - E_t^d\left[A_{t+1}\right]\right) R_{t+1}^F m_{t+1}^d\right]. && (47) \end{aligned}$$

$$\begin{aligned} \kappa k^{-1} \frac{\partial F_t}{\partial D_t} &= R_f E_t\left[-\left[A_{t+1} - E_t^d\left[A_{t+1}\right]\right] R_{t+1}^F m_{t+1}^d\right] && (48) \\ &= R_f E_t\left[-A_{t+1} R_{t+1}^F m_{t+1}^d + E_t^d\left[A_{t+1}\right] R_{t+1}^F m_{t+1}^d\right] && (49) \\ &= -E_t^d\left[A_{t+1} R_{t+1}^F\right] + E_t^d\left[R_{t+1}^F\right] E_t^d\left[A_{t+1}\right] && (50) \\ &= -\operatorname{cov}_t^d\left(A_{t+1}, R_{t+1}^F\right) && (51) \\ &= -\operatorname{cov}_t^d\left(R_{t+1}^O + R_{t+1}^F \frac{\partial F_t}{\partial D_t}, R_{t+1}^F\right) && (52) \\ &= -\operatorname{cov}_t^d\left(R_{t+1}^O, R_{t+1}^F\right) - \frac{\partial F_t}{\partial D_t} \operatorname{var}_t^d\left(R_{t+1}^F\right). && (53) \end{aligned}$$

$$\begin{aligned} \frac{\partial F_t}{\partial D_t} &= -\frac{\operatorname{cov}_t^d\left(R_{t+1}^O, R_{t+1}^F\right)}{\left(\kappa k^{-1} + \operatorname{var}_t^d\left(R_{t+1}^F\right)\right)} && (54) \\ &= -\beta_t^F \frac{\operatorname{var}_t^d\left(R_{t+1}^F\right)}{\kappa k^{-1} + \operatorname{var}_t^d\left(R_{t+1}^F\right)} && (55) \end{aligned}$$

where

$$\beta_t^F \equiv \frac{\operatorname{cov}_t^d\left(R_{t+1}^O, R_{t+1}^F\right)}{\operatorname{var}_t^d\left(R_{t+1}^F\right)} \tag{56}$$

The price sensitivity comes from differentiating the pricing equation for the option

$$\begin{aligned} \frac{\partial P_t}{\partial D_t} &= E_t\left[\frac{\partial m_{t+1}^d}{\partial D_t} P_{t+1}\right] \\ &= -k E_t\left[m_{t+1}^d\left[A_{t+1} - E_t^d\left[A_{t+1}\right]\right] P_{t+1}\right] \\ &= -k R_f^{-1} \operatorname{cov}_t^d\left(A_{t+1}, P_{t+1}\right) \\ &= -k R_f^{-1} \operatorname{cov}_t^d\left(R_{t+1}^O - R_{t+1}^F \beta_t^F \frac{\operatorname{var}_t^d\left(R_{t+1}^F\right)}{\kappa k^{-1} + \operatorname{var}_t^d\left(R_{t+1}^F\right)}, P_{t+1}\right) \end{aligned}$$

$$= -k R_f^{-1} \operatorname{cov}_t^d\left(R_{t+1}^O - R_{t+1}^F \beta_t^F \frac{\operatorname{var}_t^d\left(R_{t+1}^F\right)}{\kappa k^{-1} + \operatorname{var}_t^d\left(R_{t+1}^F\right)}, R_{t+1}^O\right) \tag{57}$$

■

Proof. The proof of proposition 2 involves simply analyzing the expectation in 4 above. We have

$$\begin{aligned}
\mathrm{cov}_t^d\left(R_{t+1}^O - R_{t+1}^F\beta_t^F\frac{\mathrm{var}_t^d\left(R_{t+1}^F\right)}{\kappa k^{-1}+\mathrm{var}_t^d\left(R_{t+1}^F\right)}, R_{t+1}^O\right) &= \mathrm{cov}_t^d\left(R_{t+1}^O - \hat{\beta}_t^F R_{t+1}^F, R_{t+1}^O\right) \quad (58)\\
&= \mathrm{cov}_t^d\left(\varepsilon_{t+1}^F + \left(\beta_t^F - \hat{\beta}_t^F\right)R_{t+1}^F, \beta_t^F R_{t+1}^F + \varepsilon_{t+1}^F\right) \quad (59)\\
&= \mathrm{var}_t^d\left[\varepsilon_{t+1}^F\right] + \left(\beta_t^F - \hat{\beta}_t^F\right)\beta_t^F\,\mathrm{var}_t^d\left[R_{t+1}^F\right] \quad (60)\\
&= \mathrm{var}_t^d\left[\varepsilon_{t+1}^F\right] + \frac{\kappa k^{-1}}{\kappa k^{-1}+\mathrm{var}_t^d\left(R_{t+1}^F\right)}\left(\beta_t^F\right)^2\mathrm{var}_t^d\left[R_{t+1}^F\right] \quad (61)
\end{aligned}$$

Next, we want to further decompose $\mathrm{var}_t^d\left[\varepsilon_{t+1}^F\right]$. Assume that, under the dealer pricing measure, z_{t+1} is orthogonal to both R_{t+1}^I and ε_{t+1}^I. We have

$$\begin{aligned}
R_{t+1}^F &= R_{t+1}^I + z_{t+1} \quad (62)\\
\beta_t^F &= \beta_t^I\frac{\sigma_{I,t}^2}{\sigma_{I,t}^2+\sigma_{z,t}^2} \quad (63)
\end{aligned}$$

where $\sigma_{I,t}^2 = \mathrm{var}_t^d\left(R_{t+1}^I\right)$. We can write

$$R_{t+1}^O = \beta_t^I R_{t+1}^I + \varepsilon_{t+1}^I \quad (64)$$

where β_t^I is the (d-measure) regression coefficient. Then

$$\begin{aligned}
\varepsilon_{t+1}^F &= R_{t+1}^O - \beta_t^F R_{t+1}^F \quad (65)\\
&= \beta_t^I R_{t+1}^I + \varepsilon_{t+1}^I - \beta_t^I\frac{\sigma_{I,t}^2}{\sigma_{I,t}^2+\sigma_{z,t}^2}\left(R_{t+1}^I + z_{t+1}\right) \quad (66)\\
&= \beta_t^I\frac{\sigma_{z,t}^2}{\sigma_{I,t}^2+\sigma_{z,t}^2}R_{t+1}^I + \varepsilon_{t+1}^I - \beta_t^I\frac{\sigma_{I,t}^2}{\sigma_{I,t}^2+\sigma_{z,t}^2}z_{t+1} \quad (67)
\end{aligned}$$

$$\mathrm{var}_t^d\left[\varepsilon_{t+1}^F\right] = \left(\beta_t^I\right)^2\frac{\sigma_{z,t}^2\sigma_{I,t}^2}{\sigma_{I,t}^2+\sigma_{z,t}^2} + \sigma_{\varepsilon,t}^2 \quad (68)$$

where $\sigma_{\varepsilon,t}^2 \equiv \mathrm{var}_t^d\left[\varepsilon_{t+1}^I\right]$.

Up to first order in κ and σ_z^2,

$$\frac{\partial P_t}{\partial D_t} = -\frac{\gamma\left(R_f-1\right)}{R_f^2}\left(\underbrace{\left(\beta_t^I\right)^2\sigma_{z,t}^2}_{\text{Basis risk}} + \underbrace{\sigma_{\varepsilon,t}^2}_{\text{Unhedgeable risk}}\right) - \underbrace{\frac{\kappa}{R_f}\left(\beta_t^I\right)^2}_{\text{Imperfect hedging}} \quad (69)$$

■

Figure 5: Net dealer gamma over time

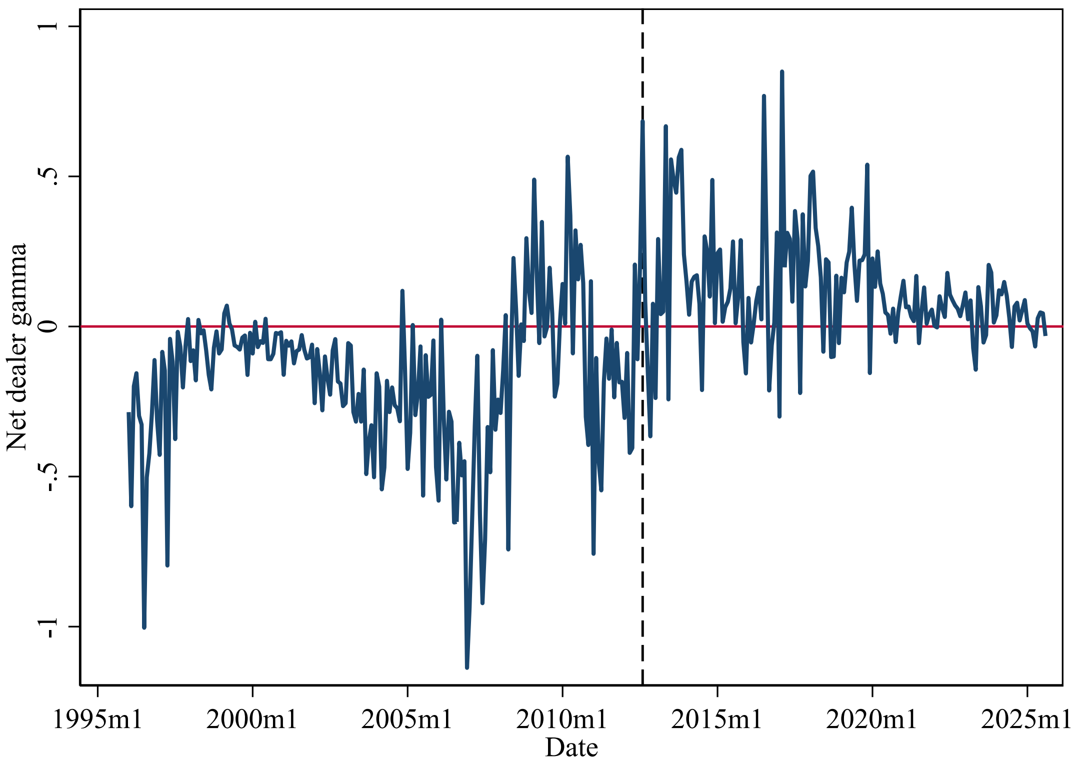

Description: Figure plots the net gamma of intermediaries over time, based on CBOE open-close and Optionmetrics data. Intermediaries are defined as in Garleanu, Pedersen, and Poteshman (2008) as entities that are neither customers nor firms, and the series uses options with 10–180 days to maturity. The dashed vertical line marks August 2012. Sample period: monthly, 1996:01–2025:08, N = 356 monthly observations. **Interpretation:** Net dealer gamma was strongly negative through the late 2000s but trended toward zero (or slightly positive) by the early 2010s, with the timing of the shift closely tracking the disappearance of the option premium documented in the previous figures.

Figure 6: Hedging costs, basis risk, and unhedgeable risk over time

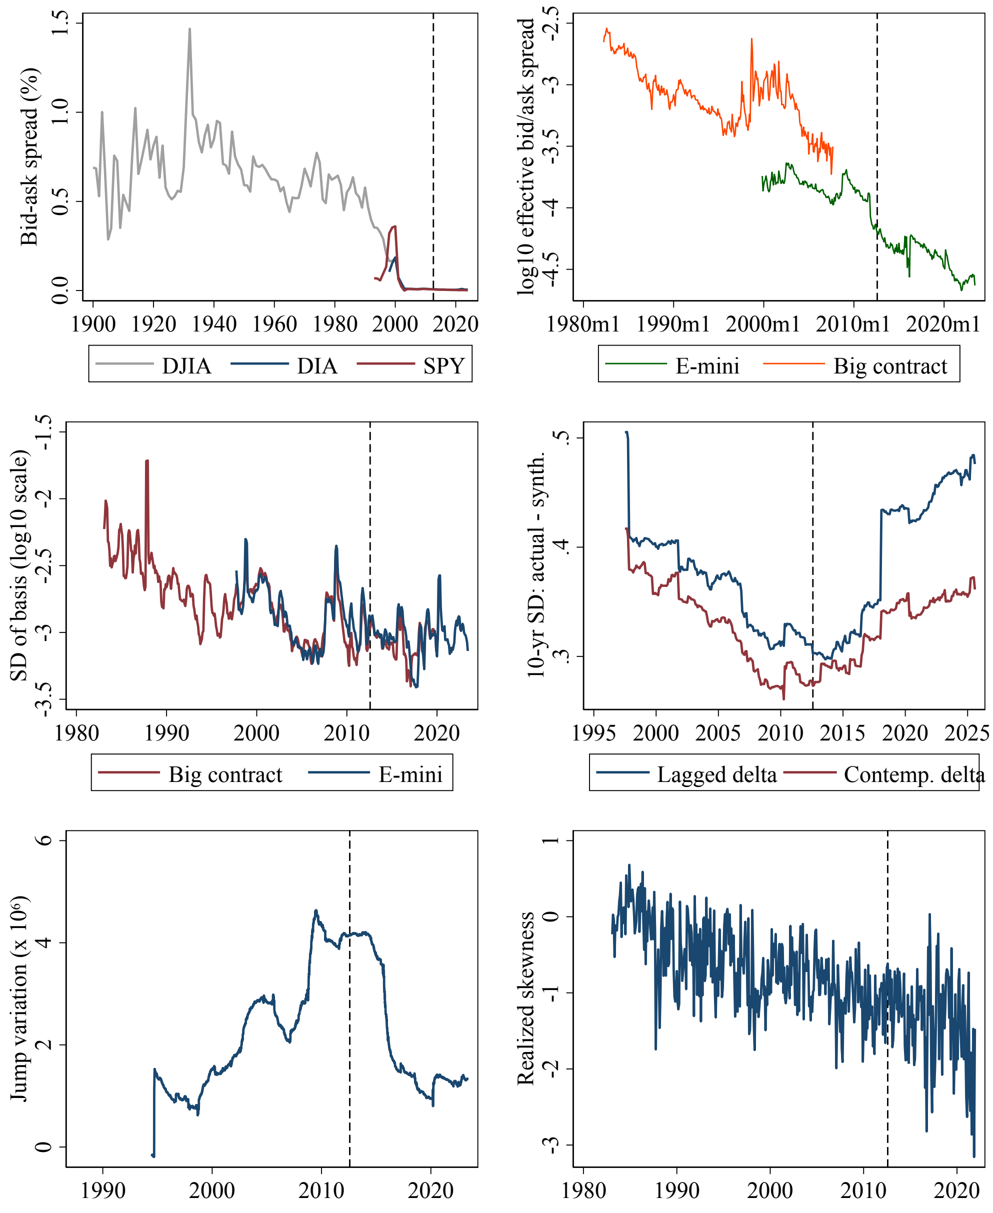

Description: Top row: posted bid-ask spreads (left) and effective bid-ask spreads (right). Middle row: standard deviation of the index-futures basis on a $\log_{10}$ scale (left) and 10-year rolling standard deviation of the difference between traded and synthetic ATM straddle returns, scaled by the initial option premium (right). Bottom row: relative jump variation (10-year moving average of realized variance minus bipower variation, left) and realized skewness (Neuberger (2012), right). The dashed vertical line in each panel marks August 2012. Sample periods and observations vary by panel: posted spreads cover 1900–2024 (DJIA stocks, $N = 101$ annual; DIA $N = 27$; SPY $N = 32$); effective spreads cover 1982:04–2023:06 ($N = 495$ monthly); index-futures basis covers 1983:02–2020:05 for the standard contract ($N = 448$ monthly) and 1997:10–2023:06 for the E-mini ($N = 309$ monthly); 10-year rolling SD of straddle hedging error covers 1987:08–2025:08 ($N = 457$ monthly, with rolling-window observations from 1997:08 onward); jump variation covers 1987:08–2023:05 ($N = 7{,}036$ daily, plotted as a 10-year moving average); realized skewness covers $N = 467$ monthly observations. **Interpretation:** Bid-ask spreads and basis risk have fallen sharply over the sample, indicating reduced hedging frictions for dealers; unhedgeable risks (jump variation and realized skewness) show no comparable decline, so the disappearance of the option premium is consistent with friction reduction rather than with declining underlying risks.

Figure 7: Variance risk premium decomposition

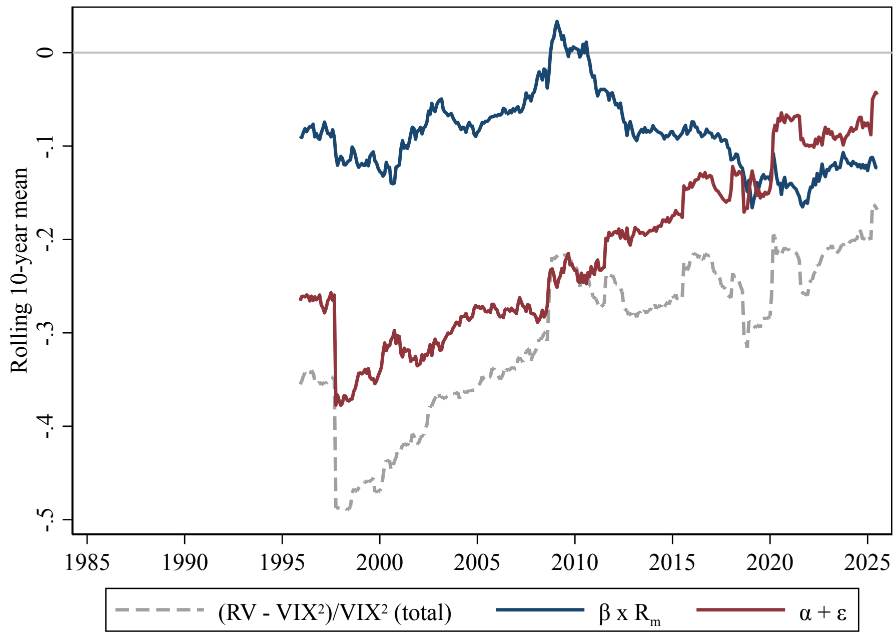

(a) Rolling 10-year mean of $\left(\mathrm{RV}_t - \mathrm{VIX}^2_{t-1}\right)/\mathrm{VIX}^2_{t-1}$: decomposition

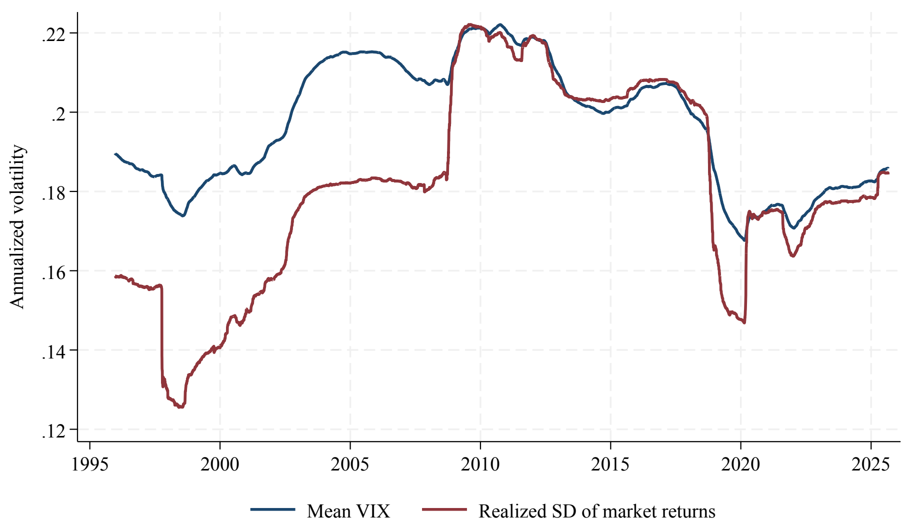

(b) Rolling 10-year VIX vs. realized volatility

Description: Panel (a) decomposes the rolling 10-year mean of $\left(\mathrm{RV}_t - \mathrm{VIX}^2_{t-1}\right)/\mathrm{VIX}^2_{t-1}$ (the variance-swap return scaled by the initial price of the swap) into a beta component ($\hat{\beta}_{\mathrm{RV}} r_{m,t}$, from a full-sample CAPM regression) and a residual alpha component ($\alpha + \varepsilon_t$). Panel (b) plots the rolling 10-year mean of the VIX against the rolling 10-year annualized standard deviation of daily market returns. Sample period: monthly variance-swap returns from 1987:08–2025:08, $N = 457$ monthly observations, of which $N = 338$ rolling 10-year-window observations are plotted in panel (a); panel (b) uses rolling 10-year windows of daily VIX and market returns over the same period. **Interpretation:** The beta component of the variance-swap return has been roughly stable while the alpha component has fallen from strongly negative to near zero; consistent with this, the rolling 10-year VIX has been close to realized volatility consistently over the past 17 years.

Figure 8: Cumulative returns on option and factor strategies

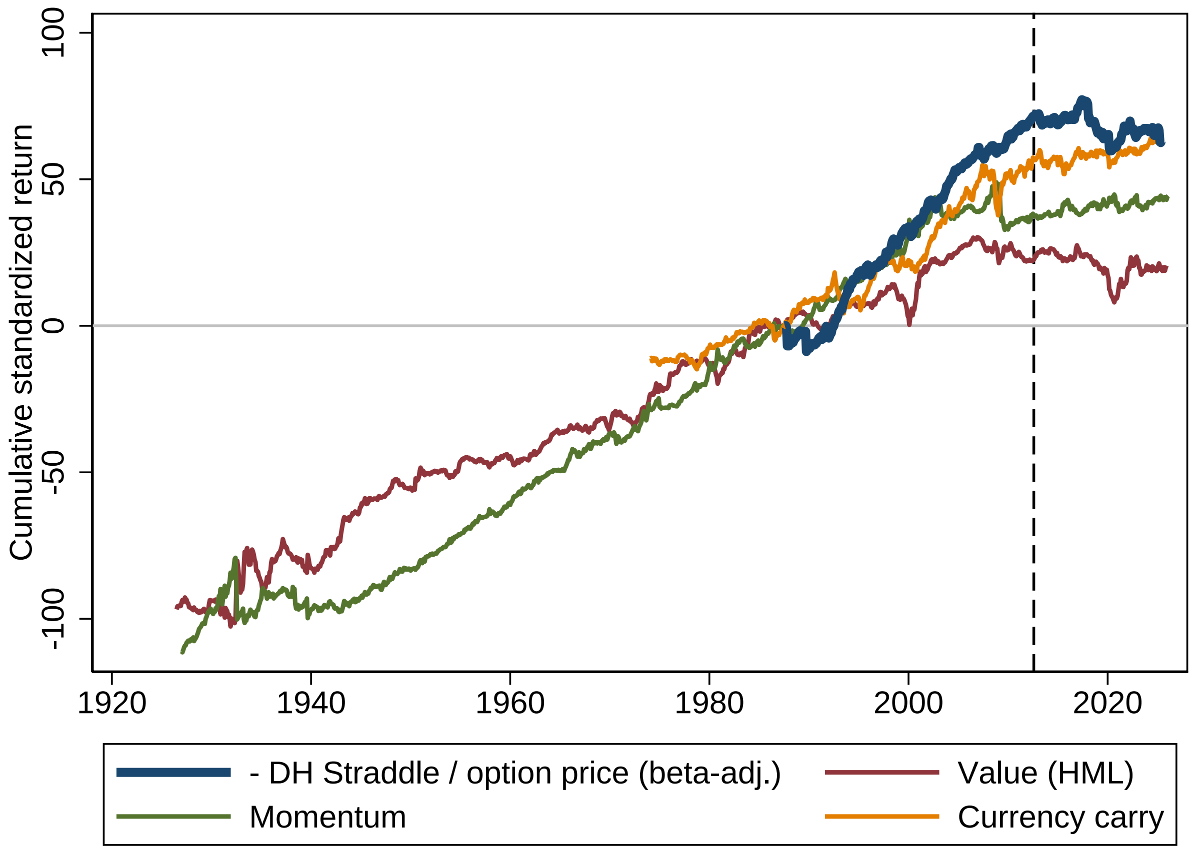

Description: Cumulative standardized monthly returns for four strategies: short delta-hedged ATM straddle (beta-hedged), HML (value), momentum, and currency carry. Each series is standardized to have unit monthly standard deviation before cumulating. Value and momentum are from Kenneth French's website; currency carry is from AQR. The dashed vertical line marks August 2012. Sample periods and observations: the short delta-hedged ATM straddle runs from 1987:08 to 2025:08 ($N = 457$ monthly observations); HML from Kenneth French covers 1926:07–2025:12 ($N = 1{,}194$); momentum covers 1927:01–2026:02 ($N = 1{,}190$); AQR currency carry covers 1974:02–2024:12 ($N = 611$). **Interpretation:** All four strategies' cumulative-return lines visibly flatten in the early 2000s, indicating that the disappearance of the option premium is part of a broader phenomenon in which several well-known factor strategies have stopped earning positive returns at roughly the same time.

Figure A.1: Power of break tests: varying shift size (full-sample bootstrap)

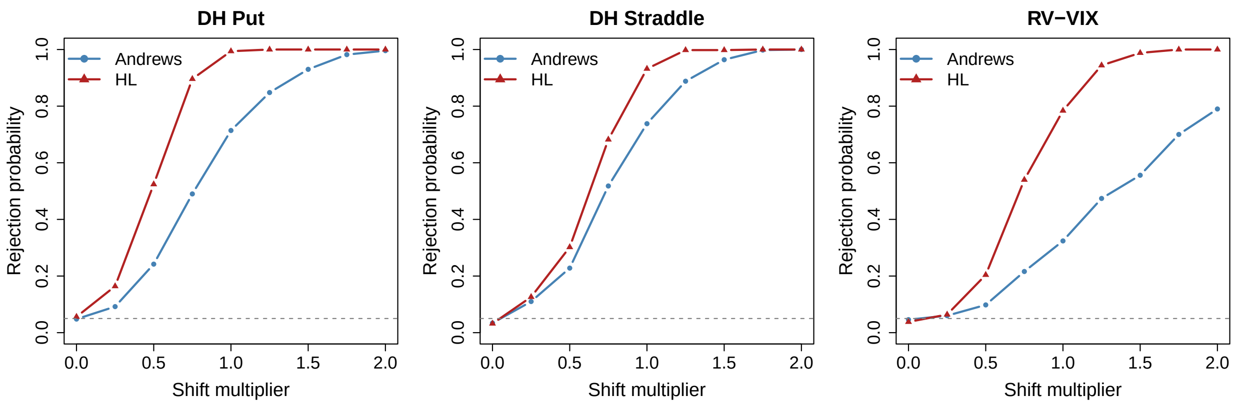

Description: Each panel shows rejection probabilities at the 5% level for the HL test (red) and the Andrews–Ploberger (1994) exponential-Wald test (blue), based on 500 simulations. The x-axis is the size of the imposed shift as a multiple of the empirically estimated shift. Returns are simulated by bootstrapping from the full sample and adding the specified shift to the post-break observations. The dashed line marks the 5% nominal size. Puts and straddles are delta-hedged, and all three series are beta-hedged. The RV-VIX series is not scaled by the VIX, since the scaled version does not satisfy the distributional stability assumption required by the HL estimator, but the scaled version does. Each simulated sample uses the empirical sample length of 457 monthly observations (1987:08–2025:08); the number of simulation draws is 500 per (shift size, series) cell. **Interpretation:** Both tests are correctly sized when the imposed shift is zero, but the HL test has substantially higher power than the AP test against shifts of the magnitude estimated empirically; with a shift equal to the empirical estimate, HL rejects nearly 100% of the time while AP rejects much less frequently.

Figure A.2: Power of break tests: varying shift size (pre/post bootstrap)

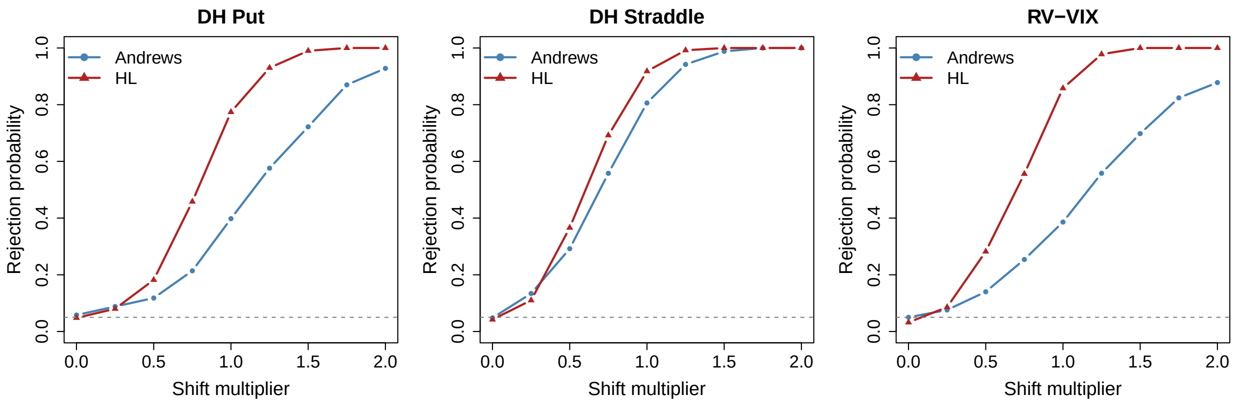

Description: Same as figure A.1 but returns are simulated by bootstrapping separately from the pre- and post-2012 subsamples (after removing the estimated mean shift), allowing for heteroskedasticity across subperiods. Empirical sample length: 457 monthly observations (301 pre-August 2012, 156 post); 500 simulation draws per (shift size, series) cell. **Interpretation:** The relative-power ranking of HL versus AP is robust to allowing different return distributions before and after 2012; the HL test continues to dominate AP against empirically plausible shift magnitudes.

Figure A.3: Power of break tests: varying sample size

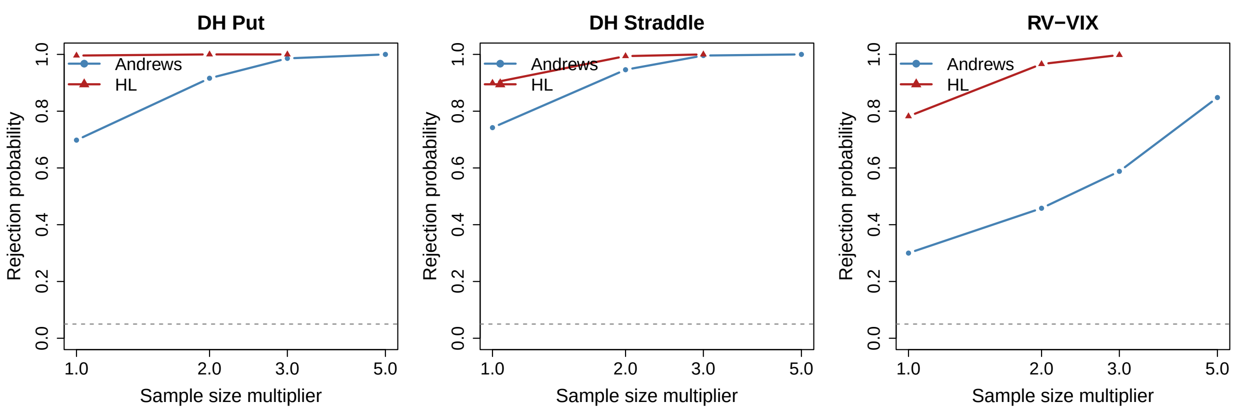

Description: Rejection probabilities at the 5% level for the HL test (red) and the Andrews–Ploberger (1994) exponential-Wald test (blue) as a function of sample size, based on 500 simulations. The x-axis is the sample size as a multiple of the actual sample. Returns are simulated by bootstrapping from the full sample and adding the empirically estimated shift to the post-break observations. The HL test is only run for scale factors up to 3× due to computational cost. Empirical sample length is 457 monthly observations (1987:08–2025:08); 500 simulation draws per (sample-size, series) cell. **Interpretation:** At the actual sample size, HL achieves rejection probabilities of 80–95%, while AP requires 2–5× more data to attain the same power; in this setting, the median-based HL estimator is the appropriate test for the magnitude of break we observe in the data.