The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXIII, NO. 2 • APRIL 2018

Option Mispricing around Nontrading Periods

CHRISTOPHER S. JONES and JOSHUA SHEMESH*

ABSTRACT

We find that option returns are significantly lower over nontrading periods, the vast majority of which are weekends. Our evidence suggests that nontrading returns cannot be explained by risk, but rather are the result of widespread and highly persistent option mispricing driven by the incorrect treatment of stock return variance during periods of market closure. The size of the effect implies that the broad spectrum of finance research involving option prices should account for nontrading effects. Our study further suggests how alternative industry practices could improve the efficiency of option markets in a meaningful way.

A VARIETY OF INSTITUTIONAL and psychological factors suggest that portfolio risk and return may differ between periods of trading and nontrading. The stock market is more volatile over trading periods than over nontrading periods, possibly due to a lower rate of private information revelation (French and Roll (1986)). It is also profoundly less liquid outside of regular trading hours, which should drive a wedge between average returns over trading and nontrading periods (French (1980), Longstaff (1995), Kelly and Clark (2011), Cliff, Cooper, and Gulen (2008)). Prolonged periods of nontrading may also give investors with limited attention a chance to process stale information, whether contained in firm earnings announcements (Dellavigna and Pollet (2009)) or news more generally (Garcia (2013)). Finally, traders may have a strong desire to exit the market over nontrading periods, with no open positions, particularly for more speculative strategies or for longer periods of nontrading (Hendershott and Seasholes (2006)).

*Christopher S. Jones is with the Marshall School of Business at the University of Southern California. Joshua Shemesh is with the Department of Banking and Finance at the Monash Business School. We are grateful for comments from Editors Cam Harvey and Ken Singleton, two anonymous referees, Mike Chernov, Dan Galai, Bruce Grundy, Larry Harris, Robert Hudson, Blair Hull, Nikunj Kapadia, Spencer Martin, Dmitriy Muravyev, Yigal Newman, Oguz Ozbas, Nick Polson, Allen Poteshman, Eric Renault, David Solomon, Selale Tuzel, Mark Weinstein, Fernando Zapatero, and seminar participants at USC, the University of Queensland, Hebrew University, the 2009 Montreal Financial Econometrics Conference, the 2009 Financial Management Association meetings, and the 2010 American Finance Association meetings. Most of the research was done while Shemesh was affiliated with the University of Melbourne. Shemesh thanks the Sanger Chair of Banking and Risk Management for financial support. Both authors state that their work on this research was not supported financially by any nonacademic sources and that no organization or individual has had any right to review this work.

DOI: 10.1111/jofi.12603

While the above effects have been documented to varying degrees in the market for individual stocks or stock indexes, one might expect most if not all to be more pronounced in the market for options. Predictable differences in volatility should have first-order effects on option prices (see also Dubinsky and Johannes (2006)), while varying liquidity should affect not only the option end user but also the market maker who must delta-hedge his underlying exposure in the stock market (e.g., Cetin et al. (2006)). Limited attention may be more problematic in option markets, where small changes in fundamentals can have large valuation effects. Finally, option trades can be extremely risky, particularly for net written option positions, pushing some traders to withdraw completely during overnight or weekend periods.

Among these effects, it is the desire to close positions over the weekend that Chen and Singal (2003) conclude was responsible for the weekend effect in stocks. The weekend effect, or the tendency of stock returns to be low over the weekend, is perhaps the most widely documented nontrading effect (French (1980)). Though evidence of the effect goes back to the early 1900s (Fields (1931)), authors such as Connolly (1989) and Chang, Pinegar, and Ravichandran (1993) show that it started to dissipate in the 1970s or 1980s. Chen and Singal (2003) hypothesize that the desire to close positions before the weekend would be most pronounced for short sellers who were deterred by the unbounded downside risk inherent in their positions. When option markets opened in the 1970s and 1980s, these investors were able to take bearish positions with limited downside risk or hedge downside risk over nontrading periods using options. Consistent with their hypothesis, Chen and Singal (2003) show that the introduction of options on a firm's stock coincides with the elimination of the weekend effect for that firm.

Writing options also exposes a trader to unbounded downside risk, particularly over nontrading periods. For call options, this comes from the potentially infinite positive payoff that accrues to the option buyer as the stock price rises. While delta-hedging reduces this exposure, the inability to perfectly hedge due to transactions costs, price discontinuities, and nontrading periods means that this downside risk cannot be eliminated. That is, a call writer can experience an arbitrarily large loss if the stock price rises high enough before the hedge can be rebalanced. When writing put options, downside risk is bounded but often extremely large, since the maximum payoff (which occurs when the stock price drops to zero) can be orders of magnitude greater than the option's price. Paradoxically, delta-hedging a put causes this downside to become unbounded, since hedging requires taking a short position in the stock, which itself is unbounded.

Given the above discussion, if investors are generally averse to holding positions with extreme downside risk over the weekend, option writers may have an incentive to cover their positions by Friday close. While this should have no effect in a Black and Scholes (1973) world in which options can be replicated via continuous trading, Bollen and Whaley (2004) show empirically that demand pressure does, in fact, have a substantial impact on option prices. Garleanu, Pedersen, and Poteshman (2009) add to this evidence and show theoretically that the effect is likely driven by the inability to perfectly replicate. Since covering net written option positions can be regarded as positive demand pressure, any tendency to do so prior to Friday close should lead to temporarily higher option prices that return to normal by the following Monday.

Alternatively, those same investors should be willing to keep their short positions open given some additional compensation for the substantial downside risk they face. Since compensation to option writers comes in the form of negative option returns, this argument suggests that we should also see a weekend effect in call and put options, with weekend option returns being significantly lower than returns over the rest of the week.

Using a large sample of daily returns on equity options over the period from January 1996 to August 2014, we investigate nontrading effects in the equity option market. We find pervasive, large, and highly statistically significant evidence that nontrading returns on options are lower than trading returns. We focus on average returns on delta-neutral positions, which are highly negative over periods of nontrading (i.e., weekends and midweek holidays), but essentially zero on average on other days. There is no evidence of any nontrading effect in underlying stock returns in our sample.

We find strong effects in puts and calls and across almost all levels of maturity and moneyness, with the results robust to different sampling methods and weighting schemes. The nontrading effect is negative in each year of our 19-year sample and is statistically significant in most. We also find strong evidence of a nontrading effect in S&P 500 Index puts and weaker evidence of nontrading effects in S&P 500 Index calls. We find significant nontrading effects for regular weekends, long weekends, and midweek holidays.

Nontrading effects are also evident in implied volatilities. As noted by French and Roll (1986), the variance of the return over a weekend (Friday close to Monday close) is just slightly greater than the variance over a regular weekday, implying little variance during periods of market closure. However, implied volatilities seem to embed the expectation that stock price variance will remain sizable even when the market is closed. This discrepancy is large and significant for both equity options and S&P 500 Index options.

We explore three possible explanations for these findings. The first is that nontrading returns are lower because of differential levels of risk between trading and nontrading periods. The second is that aversion to unbounded downside risk rises over nontrading periods, as hypothesized by Chen and Singal (2003). Since this risk is experienced by option writers but not buyers, option writers will require a higher risk premium to keep positions open. Since writers are short options, this means that option returns must be negative over the weekend. The third explanation we consider is that the nontrading effect is the result of widespread and highly persistent market mispricing.

We find that portfolios of delta-neutral positions do show somewhat greater risk over nontrading periods relative to trading periods. For puts, the effect is modest: the put portfolio's standard deviation is about 10% higher over nontrading periods (mostly weekends) than it is over regular trading intervals, and there is no significant change in return skewness or kurtosis. In contrast, the call portfolio's standard deviation is about 50% higher over nontrading periods, and this is accompanied by a significant increase in kurtosis.

Despite the increase in risk over nontrading periods, writing options remains attractive according to a number of performance measures. Short option positions display positive and significant alphas over nontrading periods that far exceed those of trading periods. Sharpe ratios and information ratios are also much higher over nontrading periods. We account for nonnormality in the distribution of option returns by computing certainty-equivalent rates under power utility and by computing the performance indexes of Kadan and Liu (2014), which are designed to account for the effects of higher moments. For each of the option portfolios considered, these measures are generally higher for nontrading periods, in most cases by an amount that is both substantial and significant, with the strongest results for delta-hedged puts.

To examine whether the nontrading effect is associated with greater *aversion* to risk, rather than with risk by itself, we examine cross-sectional relations between portfolio risk and the nontrading effect. We find that the magnitude of the effect is highly related to the Black and Scholes (1973) "gamma," the second derivative of the option price with respect to the stock price. We verify that gammas have a strong positive association with the risk of delta-hedged option portfolios, over both trading and nontrading periods. This is expected given that higher convexity reduces the effectiveness of delta-hedging, particularly when underlying price movements are large. Thus, if option writers have a greater aversion to downside risk over nontrading periods, as Chen and Singal's (2003) hypothesis suggests, then higher gamma options should experience more negative nontrading returns.

Likewise, the same logic would seemingly apply to "vega," the derivative of the option price with respect to volatility. Higher vega makes an option more sensitive to changes in the volatility of the underlying asset and therefore increases the risk of the option position, even if it is delta-hedged. Indeed, we find that the risk in delta-hedged option positions is even more strongly related to vega than it is to gamma. Surprisingly, however, greater vega does not seem to result in any additional return over nontrading intervals. Thus, greater risk aversion by option writers can only explain the nontrading effect given an explanation for why one type of risk exposure (gamma) is undesirable, while an even more important type (vega) is not.

The explanation that is more consistent with our findings is that investors are pricing options using a method that does not properly account for the difference in the behavior of stock prices when the market is open or closed. As French (1984) demonstrated over 30 years ago, if market participants do not account for the deterministic link between trading time and volatility, then option implied volatilities will tend to be too high just prior to nontrading periods. Furthermore, option prices will embed a rate of time decay that is too large over nontrading periods and too small over trading periods. This would result in a nontrading effect like the one we document, in particular, one that is driven by gamma risk but not vega risk, as it is gamma rather than vega that is closely tied to time decay. The consequence is widespread mispricing in the option market, with abnormal rates of return on all options—but particularly those with high gamma—negative over weekends and midweek holidays.

Our results imply that studies analyzing equity options data would benefit from accounting for nontrading effects, either by controlling for them explicitly or by analyzing data at a frequency (e.g., weekly or intraday) at which the nontrading effect can safely be ignored. Given that the effect of nontrading on expected returns is at least as large as the risk premia generated by traditional option pricing models, our findings also suggest that mispricing may play a larger role in the option market than has been emphasized in the literature thus far. Since option prices are used in many types of research (volatility estimation, reaction to corporate events, and understanding the risk-return relation) to infer the representative investor's information set, we believe that our findings should be of broad interest.

While a stand-alone trading strategy that captures the entire nontrading effect would be difficult to achieve due to transactions costs, our results remain useful for improving option-based trading strategies. Most obviously, nontrading effects have clear implications for when option traders should enter and exit their positions. Our findings should also be useful in setting quotes, whether by market makers or other traders using limit orders.

On a practical level, our study suggests that market efficiency may be improved by greater awareness of the problems associated with option pricing in calendar time. Currently, option pricing models are often implemented in calendar time, rather than in trading time, both in archival data sets (e.g., IvyDB) and in market data feeds (e.g., the VIX Index). We have seen little discussion of the shortcomings of this approach. If market participants (traders and data providers) were to begin implementing models in trading time, we believe that market efficiency could be improved measurably.[1]

We know of only one other paper that investigates weekly patterns in options markets. In examining short-term at-the-money options on 30 stocks over a period of just 21 months, Sheikh and Ronn (1994) find some evidence of a weekly seasonal pattern in which call returns are highest on Wednesdays, but they do not find a weekend effect in calls or puts. After adjusting option returns by subtracting Black and Scholes (1973) model-implied returns, the pattern in calls disappears, though a significant weekend effect arises in adjusted put returns. Our paper resolves the ambiguity of their findings, showing that weekend effects are present in both calls and puts over a far more extensive sample.

The paper is organized as follows. In Section I, we describe our data, sampling methods, and portfolio construction. Section II contains our main results documenting nontrading effects in returns and implied volatilities. Section III examines the relationship between nontrading effects and various measures of risk. Section IV concludes.

[1] As French (1984) discusses, option pricing in trading time is also incorrect given that interest accrues on a calendar-time basis. When interest rates are low, the use of trading time is highly preferable even if imperfect.

I. Data and Methods

Our primary data source is the IvyDB data set from OptionMetrics. This data set includes all U.S.-listed options on equities, indexes, and exchange traded funds. In our paper, we generally restrict attention to individual equity options on stocks that were members of the S&P 500 as of the end of the prior year, though for some results, we examine all equity options or options on the S&P 500 Index itself.

Throughout our analysis, we form portfolios of call and put options by taking the equal- or value-weighted average of the returns, either hedged or unhedged, of each option in that portfolio. Our data set includes closing bid-ask quotes rather than transaction prices, so we compute values and returns from quote midpoints. Since options have zero net supply, the concept of value-weighting must be reinterpreted—what we call value-weighted portfolios are actually weighted by the dollar value of open interest for each option. Excess returns are computed using the shortest maturity yield provided in the IvyDB zero curve file.

In some cases, we form call and put portfolios on the basis of maturity and/or delta. We consider three maturity ranges and six delta ranges so that all options in a given portfolio are roughly comparable and all portfolios have a reasonable number of options included. When we pool calls and puts into a single portfolio, calls and puts are each given a 50% weight, which is a way to make the resulting portfolio approximately delta-neutral and yet model-free.[2]

Because many option contracts are rarely traded, we choose our sample to alleviate the concern that our results are driven by untraded contracts. We therefore focus on a relatively liquid subsample, namely, options that have traded with positive volume for five consecutive days as of the formation date. These contracts exhibit much higher trading volume on average during our holding period, and they have much smaller bid-ask spreads.

Because individual equity options are American, early exercise can be optimal for in-the-money puts and for in-the-money calls that are about to pay dividends, and the investment return based on optimal exercise could be greater than the no-exercise return we compute. To alleviate the problem for calls, we exclude all observations that occur in the five days leading up to and including a stock's ex-dividend date, a conservative range that includes all dates on which early call exercise is optimal. Early exercise remains a minor issue for deep-in-the-money puts.

We impose several additional filters to eliminate a small number of observations that appear to be data errors, that violate arbitrage conditions, or that appear to represent noncompetitive quotes. First, we eliminate observations for which there is a large reversal (2,000% followed by −95% or vice versa) in option returns or delta-hedged returns. We also eliminate options that violate arbitrage bounds. For options on equities, which are American, we require

[2] In the Internet Appendix, which is available in the online version of the article on the *Journal of Finance* website, we consider an alternative portfolio in which calls and puts are matched on the basis of maturity, delta, and open interest. We obtain similar results for this portfolio.

Table I
Summary Statistics

This table reports basic summary statistics for our sample and for the larger data set from which our sample is constructed. Panel A shows the average number of observations (contracts observed) per day, average percentage spread, and average volume (in terms of 100-share contracts) for four different samples, each of which is a subset of the previous one. The first sample is the full IvyDB data set. The second only keeps options on firms that are members of the S&P 500. The third only keeps options with five consecutive days of positive trading volume. The final sample excludes observations that fail one of the remaining filters described in Section I. Panel B reports moments and betas for hedged and unhedged option portfolios constructed from the final sample, where portfolio weights are proportional to lagged dollar open interest. Data are daily from January 4, 1996 through August 28, 2014.

Panel A: Sample Characteristics

	Full Data Set	S&P 500 Constituents	Five Days of Positive Volume	Final Sample with Additional Filters
Average number of observations per day				
Puts	61,860.0	19,588.5	1,797.9	1,635.0
Calls	60,063.7	19,322.2	2,780.4	2,532.9
Average percentage bid-ask spread				
Puts	52.06	51.38	15.15	14.56
Calls	42.92	25.19	14.69	14.07
Average number of contracts traded				
Puts	21.67	48.45	380.51	386.56
Calls	35.91	77.34	435.84	422.28

Panel B: Return Moments

	Mean	*SD*	Skewness	Kurtosis	Beta
Unhedged excess returns					
Puts and calls	−0.0023	0.0238	1.6841	10.8287	−0.4698
Puts	−0.0057	0.1060	0.7548	5.6799	−7.0761
Calls	0.0011	0.0872	0.3924	4.9799	6.1364
Delta-hedged excess returns					
Puts and calls	−0.0012	0.0202	2.6735	29.7915	−0.9342
Puts	−0.0015	0.0204	2.1975	17.9576	−0.5344
Calls	−0.0008	0.0246	1.9672	51.2985	−1.3339

that the option price be no less than the current exercise value. For calls, the price must be less than the current stock price. For puts, it must be less than the strike. We also eliminate observations for which the bid price exceeds the ask price or the bid-ask spread is more than $10 or more than the price of the underlying stock, as these observations might be data recording errors, noncompetitive "stub" quotes posted by a market maker who does not want to trade, or undocumented missing value codes (e.g., 999). We also require, at the date of portfolio formation, that the bid price be above zero and that the bid-ask spread be no more than twice the bid.

The top panel of Table I contains statistics that describe the average cross section in our sample. On average, the IvyDB data set includes around 60,000 calls and puts per day. About one-third of those are on S&P 500 stocks, and about one-tenth are what we consider highly liquid. A small number of those observations are filtered out, leaving us with around 1,800 puts and 2,800 calls per day. The contracts we include have relatively small bid-ask spreads, around 15% on average, compared with 51% and 25% for puts and calls, respectively, for the entire S&P 500 sample. They also trade much more, at roughly 400 contracts per day on average, which is six to eight times more than the full S&P 500 sample.[3]

Taken together, the filters above eliminate the largest and most suspicious outliers in our sample and ensure that we focus primarily on more liquid options. We note that all of our main results are robust to dropping any or all of these filters, and that further tightening the requirements on trading volume increases the estimated magnitude of the nontrading effect. These robustness results are reported in Section II.B.

Our analysis requires the use of implied volatilities, and unless otherwise noted, we use those provided by IvyDB, which are computed using a binomial tree approach that accounts for dividends. Because they are equivalent to Black and Scholes (1973) values for stocks that pay no dividends, we refer to the implied volatilities and "Greeks" as "Black-Scholes" values. The option price sensitivities of each option, that is, delta, gamma, vega, and theta, are also computed with the Black and Scholes (1973) model using the implied volatility of the same option.

The IvyDB data set does not include an implied volatility for about 3% of the observations included in our analysis. Following Duarte and Jones (2008), we fill in missing implied volatilities with those of similar contracts. Specifically, if a call option's implied volatility is missing, we use the implied volatility of the put contract written on the same underlying firm with the same maturity and strike price. If both put and call implied volatilities are missing, we use the value from the same contract on the previous day. It is worth noting that this procedure for filling in missing implied volatilities relies only on current and lagged information. A portfolio strategy that uses the implied volatilities from this procedure for portfolio formation and to compute hedge ratios is therefore fully implementable. Furthermore, our results are robust to filtering out these observations.

We compute option returns based on closing bid-ask midpoints, and excess returns account for the number of calendar days within the return holding period.[4] The change in value of a delta-hedged portfolio is

$$C_t - C_{t-1} - \Delta_t (S_t - S_{t-1}),$$

[3] Trading volumes reported in Table I are those during the holding period, while the filter requiring five consecutive days of positive volume is applied to the five days prior to the holding period. The spreads reported in the table are at the start of the holding period. Spreads on the end of the holding period are almost identical.

[4] As highlighted by French (1984), interest accrues on a calendar-time basis.

where C_t is the option price, S_t is the stock price, and Δ_t is the Black and Scholes (1973) delta. Delta-hedged excess returns are therefore defined as

$$\frac{C_t - C_{t-1}}{C_{t-1}} - r_{t-1}\mathrm{ND}_{t-1,t} - \frac{\Delta_t S_{t-1}}{C_{t-1}}\left(\frac{S_t - S_{t-1}}{S_{t-1}} - r_{t-1}\mathrm{ND}_{t-1,t}\right),$$

where r_t is the riskless return per day and $\mathrm{ND}_{t-1,t}$ is the number of calendar days between date $t-1$ and t. This may be viewed as the excess return on a portfolio that combines one option contract with a zero-cost position in Δ_t shares worth of single-stock futures.

For each delta-hedged option portfolio, there is a corresponding hedge portfolio that consists of all the equity positions taken to eliminate delta. This portfolio has weights proportional to

$$-\frac{\Delta_t S_{t-1}}{C_{t-1}},$$

which is negative for calls and positive for puts. We examine these portfolios briefly to rule out the possibility that nontrading effects arise from the underlying stock returns.

The bottom panel of Table I reports moments on value-weighted portfolios of puts and calls, both separately and combined in a 50/50 portfolio. Several stylized facts are immediately apparent. First, portfolios of unhedged call options have positive average returns, while unhedged puts have negative average returns, which is consistent with calls having positive market betas and puts having negative betas. Delta-hedging makes all average returns negative, consistent with the vast literature (e.g., Coval and Shumway (2001) and Bakshi and Kapadia (2003)) documenting a negative variance risk premium.

A second observation is that delta-hedging dramatically reduces the standard deviations and betas of put and call portfolios. Delta-hedging reduces the standard deviation of the put portfolio by over 80% and reduces the beta of that portfolio from −7.08 to −0.53. Slightly smaller reductions are observed for the call portfolio. In contrast, little reduction in standard deviation is observed for the combined portfolio of puts and calls, and the beta of that portfolio is actually larger in absolute value as a result of hedging. This is because a 50/50 portfolio of puts and calls is already approximately delta-neutral, so additional delta-hedging is not particularly effective.

Third, portfolios of unhedged option positions have returns that display positive skewness and substantial excess kurtosis. Delta-hedged option returns are even more fat-tailed, the result of delta-hedging being more effective for small changes in stock prices. Large changes in stock prices, which are often the source of extreme option returns, cannot be hedged due to the convexity of option payoffs. Tails thicken as a result of a relative decrease in the low-volatility returns, for which delta-hedging is effective.

In the remainder of our paper, we focus on the delta-hedged call portfolio, the delta-hedged put portfolio, and the unhedged 50/50 portfolio of puts and calls. Because all three of these portfolios are approximately delta-neutral, our analysis focuses on the component of option prices that is orthogonal to the underlying stocks. This is desirable given that prior literature shows that the nontrading effect has vanished in equities during our sample period, a result that we confirm below in our own sample.

We view our use of the Black and Scholes (1973) model as relatively benign. Even if the delta we use to compute hedged returns is somewhat misspecified, the hedged returns nevertheless represent the returns on a feasible investment strategy (abstracting from transaction costs). Delta-hedging, while not perfect, can be expected to remove at least the majority of an option's exposure to the underlying stock. Indeeed, Hull and Suo (2002) find that Black-Scholes work about as well as any other model in this regard. Any remaining concerns should be lessened when we combine unhedged calls and puts into a single portfolio in a way that is model-free.

II. Nontrading Effects in Option Markets

A. Main Results

Our main finding is that returns in delta-neutral option positions are significantly lower over nontrading periods than trading periods. We define a nontrading period as any period in which the interval between subsequent closing prices is longer than one calendar day. The vast majority of these nontrading periods are regular weekends, with a smaller number consisting of midweek holidays and long weekends. We begin by pooling all types of nontrading periods together. Later, we examine them separately.

We present our main results in Tables II and III. Table II reports results for value-weighted portfolios of all puts and all calls, while Table III reports results for value-weighted portfolios of puts and calls formed on the basis of maturity and delta.

When we examine the unhedged put-call portfolio in Table II, we observe an average nontrading excess return of −0.76% and an average trading return of −0.08%. The difference between the two, which we define as the nontrading effect, is equal to −0.68% and under any convention is highly significant.[5] For delta-hedged puts, the average nontrading excess return is −0.82%, relative to an average trading return of 0.03%. The difference, −0.86%, is even more highly significant. For calls, the difference between nontrading and trading returns is −0.51% and again statistically significant. The smaller effect for calls relative to puts is notable but is driven largely by the fact that put and

[5] Here and elsewhere in the paper, standard errors for mean and difference-in-mean estimates are obtained by regressing returns on a constant and a nontrading dummy, then applying the Newey-West (1987) procedure with 22 lags to account for possible autocorrelation. We have also computed p-values for all of our key results using the block bootstrap, and we find these results to be very close.

Table II

Average Returns for Trading and Nontrading Periods and for Each Day of the Week

This table reports average excess returns of portfolios of unhedged and delta-hedged equity options. The first three columns of the table show close-to-close returns in trading and nontrading periods as well as the difference between them. The remaining columns show average returns by day of the week. The unhedged put and call portfolio is constructed as the 50/50 average of put and call portfolios that are each weighted according to lagged dollar open interest. Hedged puts and hedged calls are portfolios of delta-hedged options, weighted according to lagged dollar open interest. The stock hedge portfolios for puts and calls are the portfolios of stocks that are used to delta-hedge the put and call options, respectively. Values in parentheses are t-statistics computed from Newey-West (1987) standard errors with 22 lags. Data are daily from January 4, 1996 through August 28, 2014.

	Trading	Nontrading	Difference	Monday	Tuesday	Wednesday	Thursday	Friday
Unhedged puts and calls	−0.0008	−0.0076	−0.0068	−0.0079	−0.0026	−0.0008	0.0017	−0.0021
	(−1.76)	(−9.30)	(−7.86)	(−8.18)	(−4.16)	(−1.14)	(1.88)	(−2.81)
Hedged puts	0.0003	−0.0082	−0.0086	−0.0085	−0.0016	0.0019	0.0025	−0.0023
	(0.81)	(−11.91)	(−11.93)	(−10.37)	(−2.78)	(2.72)	(3.30)	(−2.99)
Hedged calls	0.0003	−0.0048	−0.0051	−0.0049	−0.0022	0.0012	0.0031	−0.0017
	(0.69)	(−5.08)	(−5.11)	(−5.34)	(−2.69)	(1.26)	(3.46)	(−2.38)
Stock hedge for puts	0.0037	0.0058	0.0020	0.0045	0.0073	0.0084	0.0018	−0.0012
	(2.25)	(1.84)	(0.55)	(1.25)	(2.28)	(2.47)	(0.56)	(−0.43)
Stock hedge for calls	−0.0015	−0.0036	−0.0020	−0.0021	−0.0059	−0.0037	0.0005	0.0015
	(−0.97)	(−1.11)	(−0.55)	(−0.59)	(−1.72)	(−1.04)	(0.17)	(0.52)
Unhedged puts	−0.0034	−0.0140	−0.0106	−0.0130	−0.0089	−0.0066	0.0007	−0.0011
	(−1.88)	(−4.04)	(−2.72)	(−3.18)	(−2.71)	(−1.79)	(0.21)	(−0.34)
Unhedged calls	0.0018	−0.0013	−0.0031	−0.0028	0.0037	0.0049	0.0026	−0.0032
	(1.30)	(−0.47)	(−0.97)	(−0.89)	(1.23)	(1.69)	(1.06)	(−1.30)

Table III

Differences between Returns Following Nontrading Days and Trading Days for Portfolios Sorted by Delta and Maturity

This table reports the average excess returns on nontrading days minus the average excess returns on trading days for single- and double-sorted portfolios formed on the basis of delta and maturity. Corresponding Newey-West (1987) t-statistics are also shown in parentheses. Portfolios of both unhedged and delta-hedged equity options are considered, though the former are only single-sorted by maturity. All methods are identical to Table II except for the classification into more disaggregated portfolios.

	All Maturities	11 to 53 Days	54 to 118 Days	119 to 252 Days
Unhedged puts and calls				
All deltas	−0.0068	−0.0103	−0.0059	−0.0026
	(−7.86)	(−9.55)	(−8.25)	(−3.34)
Hedged puts				
All deltas	−0.0086	−0.0117	−0.0069	−0.0043
	(−11.93)	(−12.66)	(−11.34)	(−7.06)
$-0.90 > \text{delta} \geq -0.99$	−0.0020	−0.0020	−0.0007	−0.0007
	(−5.90)	(−5.51)	(−2.09)	(−1.93)
$-0.75 > \text{delta} \geq -0.90$	−0.0040	−0.0041	−0.0020	−0.0015
	(−9.39)	(−9.28)	(−4.05)	(−3.82)
$-0.50 > \text{delta} \geq -0.75$	−0.0066	−0.0083	−0.0047	−0.0026
	(−10.81)	(−11.50)	(−9.43)	(−5.50)
$-0.25 > \text{delta} \geq -0.50$	−0.0104	−0.0158	−0.0080	−0.0049
	(−11.75)	(−11.60)	(−11.62)	(−7.01)
$-0.10 > \text{delta} \geq -0.25$	−0.0173	−0.0273	−0.0121	−0.0074
	(−10.59)	(−10.21)	(−10.03)	(−6.96)
$-0.01 > \text{delta} \geq -0.10$	−0.0245	−0.0328	−0.0163	−0.0065
	(−7.04)	(−7.30)	(−6.20)	(−2.82)
Hedged calls				
All deltas	−0.0051	−0.0076	−0.0051	−0.0031
	(−5.11)	(−5.57)	(−5.57)	(−3.44)
$0.01 < \text{delta} \leq 0.10$	−0.0358	−0.0436	−0.0151	−0.0160
	(−4.94)	(−5.56)	(−2.70)	(−2.85)
$0.10 < \text{delta} \leq 0.25$	−0.0230	−0.0312	−0.0162	−0.0097
	(−7.08)	(−7.48)	(−6.44)	(−3.28)
$0.25 < \text{delta} \leq 0.50$	−0.0106	−0.0159	−0.0087	−0.0051
	(−7.25)	(−7.54)	(−7.15)	(−4.50)
$0.50 < \text{delta} \leq 0.75$	−0.0043	−0.0065	−0.0034	−0.0023
	(−5.63)	(−5.93)	(−5.43)	(−3.60)
$0.75 < \text{delta} \leq 0.90$	−0.0014	−0.0017	−0.0011	0.0000
	(−3.19)	(−3.02)	(−2.96)	(−0.01)
$0.90 < \text{delta} \leq 0.99$	0.0000	−0.0002	0.0001	−0.0001
	(−0.06)	(−0.59)	(0.25)	(−0.35)

call portfolios weight in-the-money and out-of-the-money options differently, as we show below.[6]

[6] Even in an equal-weighted portfolio, differences between average call and put returns are possible and do not suggest any violation of put-call parity, which would imply similar effects on the *prices* of puts and calls rather than on their returns.

To gauge the extent to which the nontrading portfolio is present in the underlying stocks, Table II reports average excess returns over trading and nontrading periods for the portfolios of underlying stocks used to delta-hedge the put and call portfolios. The put portfolio is hedged by going long the underlying stocks, and we see that the long stock portfolio experienced somewhat higher returns during nontrading periods, though the difference is not significant. Calls are hedged by shorting stocks, so the hedge portfolio for calls earned lower returns over nontrading periods. Given that the difference between trading and nontrading returns in stocks is insignificant and that any difference has opposite effects on puts and calls, it is unlikely that the nontrading effect in options is driven by a nontrading effect in stocks.

Table II also reports averages of unhedged put and call portfolios separately. From Table I, we know that these portfolios are several times more volatile than the corresponding hedged portfolios given their large betas. Because of this, the nontrading effect is much more difficult to detect and is statistically significant only in puts.

The right side of the table summarizes average returns by day of the week, where the day given is the day at the end of the holding period. Mondays naturally represent the vast majority of nontrading-period ends, with Tuesdays and Fridays being the next most frequent (typically following three-day weekends and Thanksgiving holidays, respectively). Overall, delta-hedged excess returns are lowest on Mondays, consistent with previous results. Smaller negative returns on other days are consistent with a variance risk premium, though the positive average returns observed on Wednesdays and Thursdays are somewhat anomalous. In the Internet Appendix, we show that robust estimators, which downweight large outliers, eliminate these positive average returns, while the nontrading effect is robust to a variety of alternative estimators and to the use of block bootstrap standard errors.

Table III examines the nontrading effect, or the difference between nontrading and trading returns, for disaggregated portfolios. We form unhedged call/put portfolios on the basis of maturity, measured as the number of trading days until option expiration, and delta-hedged call and put portfolios based on maturity and moneyness, where we use lagged option deltas to measure moneyness.[7] In 59 of the 60 portfolios, average nontrading returns are less than average trading returns, with 54 of these differences statistically significant. One difference is positive but insignificant.

Several patterns emerge from the table. First, the nontrading effect is stronger for shorter term options, being several times higher for options in the 11- to 53-day range relative to the 119- to 252-day range. Second, the effect is much larger for out-of-the-money options (deltas close to zero) than it is for in-the-money options (deltas close to -1 or $+1$). Both of these patterns suggest a possible relation between the nontrading effect and one or more of the Black

[7] We do not double-sort the unhedged put-call portfolio because this would result in portfolios with large negative or positive deltas whose returns were dominated by their underlying stock exposures.

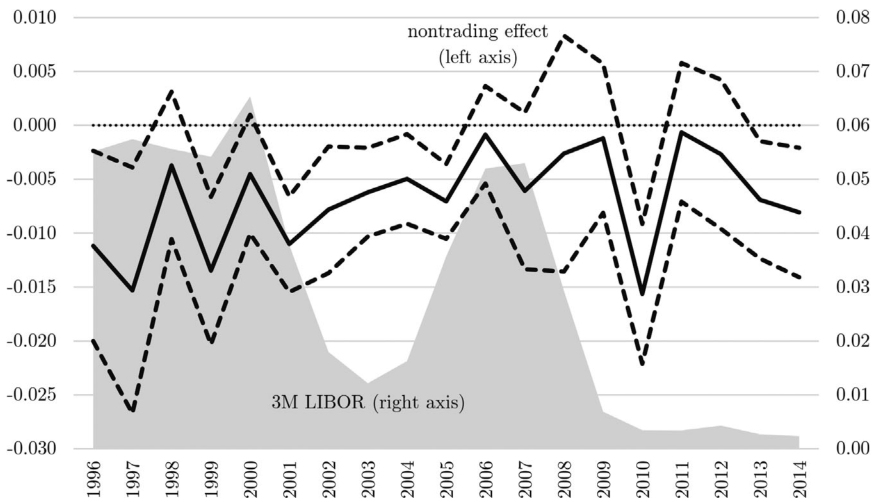

Figure 1. Average unhedged put and call portfolio returns in each calendar year. This graph depicts the average return on the unhedged put and call portfolio in each year of our sample. The solid line denotes the average, and the two dashed lines denote the upper and lower bounds of a 95% confidence interval constructed using Newey-West (1987) standard errors. The shaded area shows the average three-month LIBOR rate within each calendar year.

and Scholes (1973) Greeks, which measure the sensitivity of option values to various factors and are highly affected by maturity and moneyness. We return to this issue in Section III.C, where we directly examine the relation between option characteristics and the nontrading effect.

A final result from Table III is that, for the same absolute delta, which is a measure of option moneyness, there is little tendency for the nontrading effect in puts to exceed that of calls. This confirms that the greater nontrading effect found for the aggregated put portfolio in Table II was primarily a composition effect, driven by the fact that open interest in out-of-the-money options is larger for puts than calls.

To gauge the consistency of our results, we recompute the nontrading effect for each year of our sample. Figure 1 plots the yearly estimates for the unhedged call/put portfolio along with asymptotic 95% confidence intervals. In each of the 19 years of our sample, the nontrading effect is negative. In 11 out of the 19 years, the effect is statistically significant. We observe similar results in the delta-hedged portfolios, where 18 out of 19 years are negative for puts and 16 out of 19 are negative for calls.

We do see that the nontrading effect in the second half of the sample is substantially more volatile and slightly smaller on average relative to that in the first half of the sample. Some of these changes are likely driven by the extreme volatility of the recent financial crisis. It is also possible that the attenuation of the effect is related to the rise of algorithmic trading and the replacement of traditional market makers by high-frequency trading firms, both of which followed the 2003 creation of Options Linkage, which connected all option exchanges in the United States (see Muravyev and Pearson (2016)). It is notable, however, that the last year of the sample displays an above-average effect, indicating that the effect has survived these changes to market structure.

The figure also shows the average three-month London Interbank Offered Rate (LIBOR) in each year of the sample to alleviate the concern that nontrading effects arise through an interest rate channel. Ex ante, such an effect is unlikely for several reasons. First, our returns are in excess of the riskless rate, with the correct adjustment for the number of calendar days in the holding period. Any residual effect of interest rates can therefore arise only through risk premia. One might be concerned, for instance, that unhedged risks (vega and gamma) become larger over nontrading periods when interest rates are high. However, at least under the Black and Scholes (1973) model, the effects of interest rates on Greeks such as vega and gamma are small for most options, and negligible for options with one month till expiration, where our results are strongest. Second, prior literature, including French (1984), Scott (1997), and Bakshi, Cao, and Chen (1997, 2015), generally finds a limited role for interest rates in the pricing of equity or equity index options, in that the precise specification of the interest rate process has little effect on option prices or returns. In any case, it is evident from Figure 1 that there is no relation between the two series plotted, which confirms that nontrading effects do not arise through an interest rate channel.

As discussed in Section I, a potential concern is that our returns are calculated assuming no early exercise. Since we eliminated observations on and immediately prior to dividend ex-days, early exercise will never be optimal for calls, though it could be optimal for some in-the-money puts. For such puts, the value of early exercise comes from accelerating the fixed option premium forward in time. While this can have a significant effect on option value, the loss in value that results from delaying exercise by one day should be small, as it is at most the loss of one day of interest on the strike price. Thus, for these options, our computed returns should at most slightly understate the returns of a strategy in which exercise decisions are made optimally.

More importantly, early exercise cannot explain our results because it is never optimal for call and put options that are out of the money. Since we find large nontrading effects across all moneyness levels, early exercise due to dividends or any other factor cannot explain our findings.

B. Robustness to Portfolio Construction and Subsamples

In Table IV, we examine nontrading effects in different samples to determine whether our main results are sensitive to the choice of empirical approach. Panel A reports results using a number of alternative methods for computing option portfolio returns. We first replicate our main results using ask-to-bid and bid-to-ask returns in addition to the midpoint-based returns that we consider throughout the paper. All three methods for calculating returns lead

Table IV

Nontrading Effects in Other Portfolios and Samples

This table reports nontrading effects, defined as a nontrading mean minus a trading mean, for alternative portfolio constructions and for changes in the implied volatility of at-the-money options. Unless noted otherwise, all portfolios are constructed from the main sample (S&P 500 stocks, five days of positive volume); are weighted by dollar open interest; are based on quote midpoints; and impose filters for dividends, arbitrage violations, price reversals, and maximum spreads. In Panel A, we compute option returns as either midpoint-to-midpoint, ask-to-bid, or bid-to-ask. Results based on midpoints are identical to those in Table II. We also compute equal-weighted portfolios and portfolios weighted by open interest measured in terms of contracts rather than dollars. Panel B shows results for the following alternative data samples: options on all stocks (not just S&P 500 stocks, no trading volume requirement); the main sample but without filters; a more restricted sample of options with five consecutive days of trading 100 or more contracts; and a more restricted sample of options with relative bid-ask spreads of 5% or less. Panel C shows results for S&P 500 Index options, including deep-out-of-the-money options with Black-Scholes deltas below 0.1 in absolute value. Panel D reports nontrading effects for different types of nontrading periods, namely, regular weekends, midweek holidays, and long weekends with three or more days of market closure. Newey-West (1987) t-statistics are in parentheses. Data are daily from January 4, 1996 through August 28, 2014.

Panel A: Nontrading Effects Using Alternative Return Constructions and Portfolio Weights

	Dollar Open Interest Weighted			Equal Weighted	Contract Weighted
	Midpoints	Ask to Bid	Bid to Ask		
Unhedged puts and calls	−0.0068	−0.0060	−0.0079	−0.0197	−0.0184
	(−7.86)	(6.99)	(8.38)	(−14.75)	(−13.32)
Hedged puts	−0.0086	−0.0075	−0.0100	−0.0213	−0.0173
	(−11.93)	(9.95)	(11.02)	(−14.54)	(−11.06)
Hedged calls	−0.0051	−0.0045	−0.0059	−0.0162	−0.0160
	(−5.11)	(3.98)	(5.91)	(−7.25)	(−5.80)

Panel B: Alternative Data Samples

	Options on All Stocks	Zero Volume Contracts Only	Main Sample No Filters	Main Sample Five Days of Volume>100	Main Sample Maximum 5% Spread
Unhedged puts and calls	−0.0043	−0.0027	−0.0072	−0.0101	−0.0044
	(−8.50)	(−8.08)	(−8.52)	(−6.72)	(−4.69)
Hedged puts	−0.0061	−0.0042	−0.0085	−0.0112	−0.0063
	(−15.38)	(−16.31)	(−12.45)	(−8.42)	(−9.24)
Hedged calls	−0.0037	−0.0018	−0.0053	−0.0074	−0.0029
	(−5.32)	(−4.49)	(−5.04)	(−4.73)	(−4.00)

Panel C: Nontrading Effects in S&P 500 Index Options

	Value Weighted	Equal Weighted	Equal Weighted, Deep OTM
Puts and calls	−0.0073	−0.0250	−0.0471
	(−3.87)	(−4.93)	(−4.70)
Puts	−0.0107	−0.0315	−0.0519
	(−4.56)	(−5.34)	(−4.95)
Calls	−0.0030	−0.0155	−0.0453
	(−0.91)	(−1.66)	(−2.41)

(*Continued*)

Table IV—*Continued*

Panel D: Nontrading Effects in Different Nontrading Periods

	Regular Weekends (A)	Midweek Holidays (B)	Long Weekends (C)	(B)–(A)	(C)–(A)
Puts and calls	−0.0079	−0.0032	−0.0071	0.0048	0.0009
	(−8.06)	(−0.83)	(−3.39)	(1.21)	(0.40)
Puts	−0.0087	−0.0062	−0.0056	0.0025	0.0031
	(−10.20)	(−2.96)	(−2.79)	(1.04)	(1.45)
Calls	−0.0052	−0.0065	−0.0016	−0.0013	0.0036
	(−5.03)	(−1.82)	(−0.53)	(−0.33)	(1.11)
Number of observations	851	44	122		

to similar conclusions about the strength and significance of the nontrading effect.

We next compute portfolio returns using equal weighting instead of weighting by the dollar value of open interest. Equal weighting has several effects. First, it puts relatively more weight on less liquid contracts that do not trade as frequently and that hold little open interest. Second, it likely increases the measurement error bias identified by Blume and Stambaugh (1983), though this should affect both trading and nontrading returns and hence possibly wash out when computing their difference. Finally, and most importantly, it puts relatively more emphasis on lower priced contracts, which are typically shorter term and deeper out of the money. Since we previously found that the nontrading effect is most pronounced for short-term deep-out-of-the-money contracts, equal weighting should increase the size of the nontrading effect. Consistent with this conjecture, we find that equal weighting causes the nontrading effect to more than double.

When we instead weight by the open interest in terms of the number of contracts rather than the dollar value, we obtain results that are very similar to those for the equal-weighted portfolios. This suggests that the larger nontrading effects in the equal-weighted portfolio are not a result of overweighting thinly traded options, but are instead due to the equal-weighted portfolio's greater emphasis on low-priced deep out-of-the-money contracts.

In Panel B of Table IV, we consider different option samples. We start by examining options on all U.S.-listed equities rather than limiting attention to those that are members of the S&P 500, and we do not require that contracts trade for five consecutive days. The resulting sample is much larger but substantially less liquid. In this sample, the nontrading effect is reduced by roughly one-third across the three portfolios, but it remains highly statistically significant. To see if the effect holds in even the most illiquid contracts, we further reduce liquidity by restricting the sample to include only those options with zero trading volume on the portfolio formation day. The effect is again slightly smaller but still highly significant.

We also examine whether our use of filters (dividends, reversals, etc.) has any effect on our results. We find that dropping these filters has almost no effect on our main results, though they are likely more important when we analyze less aggregated portfolios.

In addition, we examine whether samples that are even more liquid than the one we focus on also show nontrading effects. We construct one such sample by requiring option contracts to have five consecutive days of trading at least 100 contracts, which reduces the sample by about 85% relative to our primary sample (S&P 500 firms, five days of positive volume).[8] The result is that the nontrading effect is amplified, roughly 40% larger for each of the three portfolios considered. Overall, it appears that the effect we document in this paper is perhaps twice as large for the most heavily traded options as it is for the least traded ones.

We also analyze a subsample of options with bid-ask spreads no greater than 5% of the quote midpoint, which should further minimize concerns regarding our reliance on bid-ask midpoints in computing returns. For these portfolios, the size of the nontrading effect is about two-thirds that of the baseline portfolios, which would appear to go against other results suggesting that the effect is larger in more liquid contracts. The explanation is that low-percentage spreads are more common among in-the-money contracts, where the nontrading effect is weaker. Liquidity, in this case, is proxying for moneyness.

Panel C examines the presence of a nontrading effect in S&P 500 Index options, where we impose all filters used to construct our primary sample except for the one eliminating dividends. Intuitively, one might expect index options to be driven by the same systematic risk factors that affect individual equity options. However, the two differ significantly in that index options are options on a portfolio rather than a portfolio of options. They are therefore affected by correlations as well as volatilities. Risk factors that drive correlations, as in the model of Driessen, Maenhout, and Vilkov (2009), should therefore affect only options on the index.

In contrast, systematic movements in average idiosyncratic volatility will affect portfolios of individual options but will have little effect on options on the index. Such systematic movements have been documented by Campbell et al. (2001) and have been found by Goyal and Santa-Clara (2003) to drive aggregate risk premia. If the risk premium on this factor has a weekly seasonal pattern, then this will lead to a weekend effect only in individual equity options.

Table IV shows that the average nontrading effects for the value-weighted unhedged put-call portfolio and the delta-hedged put portfolio are both slightly larger than the effects measured from equity options, though the t-statistics are somewhat lower. We find no significant effect in delta-hedged calls, though the point estimate in index call options is negative and sizable. The reason for the weak results on index calls appears to be related to differences in the levels of open interest of calls and puts across various levels of maturity and moneyness.

[8] Requiring 1,000 contracts per day reduces the sample much further still. Doing so has little impact on the nontrading effect, raising the effect slightly for calls and lowering it slightly for puts. However, standard errors are larger due to a smaller sample, reducing the significance of the effect.

For example, deep-out-of-the-money puts (deltas greater than −0.1), which are commonly used to obtain portfolio insurance against market risk, comprise about 3% of the total dollar value of puts traded, on average. Deep-out-of-the-money calls (deltas below 0.1) comprise only 0.6% of the total amount of calls. Since out-of-the-money contracts exhibit the largest nontrading effects, the average nontrading effect for calls is reduced.[9]

If we instead form portfolios that are equal weighted, which gives relatively more weight to out-of-the-money contracts, all nontrading effects rise substantially, but the effect is still insignificant for S&P 500 Index calls. However, if we focus solely on calls with deltas less than 0.1, the effect is statistically significant and, at around −4.5%, extremely large in economic magnitude.[10]

Panel D of Table IV reports estimates of the nontrading effect for different types of nontrading periods, all of which have been included in our analysis thus far. For each type, the values reported are the average returns on the nontrading period described (e.g., midweek holidays) minus the average returns over all trading days. If what we are documenting is truly a nontrading effect rather than a weekend effect, then we should also see negative option returns over midweek holidays, and we might also expect the effect to be stronger over long weekends.

In our sample, there are only 44 midweek holidays, so the average returns over these periods are estimated imprecisely. Nevertheless, this small subsample shows a highly significant nontrading effect for hedged puts and a marginally significant result for hedged calls, which supports the view that our findings represent a *nontrading* effect rather than *weekend* effect. On the other hand, the nontrading effect for unhedged options in this small sample is negative but insignificant. The table also reports results on the difference between the nontrading effect over midweek holidays and that of regular weekends. No significant difference is observed, but this may again be attributable to the small sample of midweek holidays.

Long weekends, with three or more days of nontrading, are somewhat more common. During these periods, we find significant nontrading effects for unhedged options and hedged puts, though not for hedged calls. As with midweek holidays, there is no significant difference between the nontrading effect over long weekends and that of regular weekends. Still, the effects appear weaker rather than stronger over long weekends relative to regular weekends. While the weaker nontrading effect over these periods should be taken with a grain of salt given the small number of long weekends, it may also be the case that some options traders are aware of the nontrading effect and try to write options to benefit from it in more conspicuous cases, such as long weekends. As indirect

[9] We explore the relation between option moneyness and the strength of the nontrading effect in Section III.C, where we provide evidence that the relation is due to a greater amount of excess time decay in out-of-the-money options.

[10] The Internet Appendix examines S&P 500 Index options in more detail. In short, the evidence is statistically weaker than that for the equity option sample that we focus on in the paper, but point estimates of the effect for different maturities and deltas are generally similar to those found in equity options.

evidence of this interpretation, we find that average hedged put and call returns are −0.53% and −0.37%, respectively, on the day before a long weekend, compared to just −0.04% and −0.07% on the day before a regular weekend. These differences are statistically significant at the 5% level.

Taken together, the results in Table IV show that the nontrading effect is extremely robust but more pronounced in liquid contracts. While the effect is insignificant in S&P 500 Index calls, it is large and significant in puts. This is notable because the markets for index and equity options are different in many respects. For instance, as Garleanu, Pedersen, and Poteshman (2009) observe, the "end users" presumed responsible for variation in option demand are on average buyers of index options but writers of equity options. The fact that both types of options exhibit nontrading effects of the same sign and similar magnitude appears to challenge demand-based option pricing as an explanation for our findings. Furthermore, we find a strong effect even when there is no apparent change in demand, namely, in contracts with zero trading volume. This does not necessarily rule out the demand-based option pricing model as an explanation for the nontrading effect, considering that Garleanu, Pedersen, and Poteshman show that option prices should be affected by trades in other contracts. Nevertheless, the strength of the effect in untraded contracts is surprising.

C. Implied and Realized Variances

In this section, we assess the significance of nontrading effects in implied variances. Specifically, we ask whether the variance of returns over nontrading periods that is implicit in option prices matches the realized variance from stock returns.

Define the quadratic variation between times u and v as

$$QV_{u,v} = \int_{s=u}^{v} \sigma_s^2 ds,$$

where σ_t is the instantaneous volatility of the log stock price process. Now define the time-t implied variance of an option with fixed expiration date T as

$$IV_{t,T} = \mathrm{E}_t^Q\left[QV_{t,T}\right] = \mathrm{E}_t^Q\left[QV_{t,t'} + QV_{t',T}\right].$$

Note that we are measuring the total amount of variance remaining over the option's lifetime rather than dividing by the time to expiration, which would result in a variance rate.

The change, from time t to time t', in implied variance remaining until expiration is

$$IV_{t',T} - IV_{t,T} = \mathrm{E}_{t'}^Q\left[QV_{t',T}\right] - \mathrm{E}_t^Q\left[QV_{t,T}\right].$$

This change is on average negative as a result of moving closer to expiration, as the quadratic variation between times t and t' is eliminated from implied variance at time t'. In studies of earnings announcements, Patell and Wolfson

(1981) and Dubinsky and Johannes (2006) use this change to estimate the announcement-day variance for a sample of stocks.

If we add the realized quadratic variation of the underlying asset over the same period to offset the effect of moving closer to expiration, we get a variance "differential,"

$$D_{t,t'} \equiv IV_{t',T} - IV_{t,T} + QV_{t,t'}, \tag{1}$$

which measures the difference between the actual variance and the decline in implied variance remaining until expiration. This differential can be reexpressed as

$$D_{t,t'} = \left(\mathrm{E}_{t'}^{Q}\left[QV_{t',T}\right] - \mathrm{E}_{t}^{Q}\left[QV_{t',T}\right]\right) + \left(QV_{t,t'} - \mathrm{E}_{t}^{Q}\left[QV_{t,t'}\right]\right). \tag{2}$$

The first term represents the change in the risk-neutral expectation of future variance over the interval (t', T), which is equal to the price appreciation on a forward start variance swap. Without a variance risk premium, forward prices are martingales, so this term should be zero in expectation. The second term represents the difference between actual and implied variance. Again, the absence of a variance risk premium should make this term zero in expectation. Thus, if there is no variance risk premium, that is, if implied variance is an unbiased forecast of future-realized variance, then the expectation of $D_{t,t'}$ should be zero. This statement holds regardless of the length of time between t and t' and regardless of whether that interval includes a nontrading period.

To determine whether options embed an expectation of future variance over nontrading periods that is larger than the actual variance over nontrading periods, we measure the variance differential $D_{t,t+1}$ for each stock and for the S&P 500 Index on each day in our sample. For the equity option sample, we average these values across stocks to obtain a single value on each day. If the nontrading variance implicit in option prices exceeds that present in actual returns, then $D_{t,t+1}$ should be negative on average over nontrading periods.

We measure the variance differential using two different methods for computing implied variances, which are based on either the Black and Scholes (1973) formula applied to at-the-money options or model-free implied variances. In the former case, the implied variance is the average of the values computed from the call and put with deltas closest to 0.5 and −0.5, respectively. In the latter case, we apply the interpolation method of Hansis, Schlag, and Vilkov (2010) to the calculations proposed in Bakshi, Kapadia, and Madan (2003). In both cases, the implied variances we compute are "to term," meaning that they are not annualized by dividing by the time remaining until expiration.

As is standard (e.g., Andersen et al. (2001)), we proxy for quadratic variation with a realized variance measure computed as the daily sum of all squared five-minute returns along with the squared overnight return.[11] Intraday returns

[11] We have also examined robustness to using squared daily returns to proxy for quadratic variation. Doing so produces equivalent results for the equity option sample but somewhat weaker results for the S&P 500 Index option sample, though most significant coefficients remain so.

on individual equities are computed using the NYSE Trades and Quotes (TAQ) database, where we use the last recorded transaction price within each five-minute interval between 9:30 am and 4:30 pm EST, a total of 84 observations per day for each firm.[12] Realized variances for the S&P 500 Index are from the Oxford-Man Realized Library, to which we add the squared overnight return.

We assess the impact of nontrading and risk premia using the following regression:

$$D_{t,t+1} = a + b\,\mathrm{NT}_{t+1} + c\,\mathrm{VRP}_t + d\,\mathrm{NT}_{t+1}\mathrm{VRP}_t + e_{t,t+1}, \tag{3}$$

where $D_{t,t+1}$ now represents the *average* volatility differential computed from time t to $t+1$. In this regression, NT represents a nontrading indicator, which is equal to 1 on nontrading days and 0 otherwise, and VRP represents a variance risk premium proxy similar to that proposed by Bollerslev, Tauchen, and Zhou (2009). Specifically, VRP represents the demeaned difference between the square of the VIX Index, which is a model-free measure of implied variance, and the 22-day moving average of daily realized variances computed from five-minute index returns. Demeaning has no effect on the c or d coefficients, but it does affect the estimates of a and b in the specifications in which VRP is included.[13]

The results of these regressions are in Table V. In short, we observe a highly significant nontrading effect (the b coefficient) for both equity options and index options. This indicates that the decrease in implied volatility remaining until expiration is too large over nontrading periods, implying that option prices embed an assumption of nontrading-period variance that is significantly higher than what is actually realized. To gauge how much higher, we can compare the estimated b coefficient, which measures the size of the nontrading effect on variance differentials, to the average stock-level realized variance.

In our sample, the average stock-level realized variance is 0.00082, which is equivalent to an average daily volatility of 2.87%. For the same equities, estimates of the nontrading effect for variance differentials (b) ranges from around 0.0006 to 0.0009 (note that the b coefficients in Table V have been multiplied by 1,000), numbers that are approximately equal to the average realized variance. This indicates that, during the average nontrading period, the excess decline in implied variance is around one full day of realized variance. Similar results obtain for index options. While the b coefficient is smaller for those options, the average realized variance for the S&P 500 Index is only 0.00012 (an average volatility of 1.08%). Estimated b coefficients from index options are all larger

[12] TAQ data are filtered before the five-minute returns are computed. Specifically, we exclude observations with zero price or zero size, corrected orders, and trades with condition codes B, G, J, K, L, O, T, W, or Z. Trades on all exchanges are included. We also eliminate observations that result in transaction-to-transaction return reversals of 25% or more and observations that are outside the CRSP daily high-low range. Finally, we compute size-weighted median prices for all transactions with the same time stamp.

[13] Demeaning the variance risk premia (VRP) is done to eliminate collinearity between the nontrading indicator and its interaction with the VRP. Without detrending, neither regressor would be significant in the unrestricted specifications.

Table V

The Differential between Elapsed Actual and Implied Variance

This table examines the daily differentials (D_{it}) between elapsed realized variance (RV_{it}) and elapsed implied variance ($-\Delta IV_{it}$). RV_{it} is the sum of all squared returns computed over five-minute intervals within a day plus the squared overnight return. $-\Delta IV_{it}$, which measures the one-day reduction in risk-neutral variance remaining until the option's expiration, is computed as the τ-day implied variance on day $t-1$ minus the $\tau-1$-day implied variance on day t, where implied variance is computed on a cumulative nonannualized basis. As described in Section II.C, we use both the Black-Scholes model (BSIV) and the model-free approach (MFIV) of Bakshi, Kapadia, and Madan (2003) with the interpolation scheme of Hansis, Schlag, and Vilkov (2010) to compute these implied variances. In the top panel, we regress the average D_{it} across all stocks in the S&P 500 Index on a nontrading indicator, a measure of the variance risk premium, and an interaction term. In the lower panel, we regress the S&P 500 Index differential on the same variables. The variance risk premium is equal to the square of the VIX Index, converted to represent a daily value, minus the 22-day moving average of realized variances on the S&P 500 Index, which are themselves the sum of all squared five-minute returns plus the overnight return. Estimates of the intercept and the nontrading indicator have been multiplied by 1,000. Values in parentheses are t-statistics computed from Newey-West (1987) standard errors with 22 lags. Data are daily from February 2, 1996 through August 28, 2014.

Panel A: Individual Equities						
	D_{it} Based on BSIV			D_{it} Based on MFIV		
Intercept	0.1100	0.1091	0.1078	0.1520	0.1505	0.1498
	(2.53)	(2.56)	(2.58)	(3.53)	(3.60)	(3.62)
Nontrading indicator (NT Ind)	−0.9128	−0.9086	−0.9253	−0.5858	−0.5792	−0.5887
	(−8.76)	(−8.73)	(−8.39)	(−7.37)	(−7.28)	(−7.07)
Variance risk premium (VRP)		0.8451	2.0608		1.3423	2.0399
		(0.84)	(1.36)		(1.30)	(1.43)
NT Ind × VRP			−5.8610			−3.3646
			(−1.36)			(−1.16)
R^2	0.0326	0.0342	0.0463	0.0166	0.0214	0.0264
Panel B: S&P 500 Index						
Intercept	−0.4428	−0.4400	−0.4406	−0.0611	−0.0586	−0.0590
	(−10.99)	(−10.94)	(−11.06)	(−2.85)	(−2.58)	(−2.68)
Nontrading indicator (NT Ind)	−0.2964	−0.3092	−0.3172	−0.1443	−0.1555	−0.1608
	(−3.46)	(−3.68)	(−3.64)	(−2.38)	(−2.63)	(−2.57)
Variance risk premium (VRP)		−2.7161	−2.0983		−2.4489	−2.0367
		(−1.78)	(−1.25)		(−3.75)	(−2.83)
NT Ind × VRP			−2.9820			−1.9906
			(−1.21)			(−0.87)
R^2	0.0033	0.0180	0.0210	0.0023	0.0387	0.0426

than this value, meaning that the decline in index implied volatilities is at least as excessive on a relative basis. Thus, it appears as if option prices embed a full extra day of variance over each nontrading period.

Controlling for variance risk premia has little effect on the nontrading coefficient, though in some cases, it is a highly significant predictor of variance differentials for the S&P 500 Index. The addition of an interaction term also has very little effect. Were the nontrading effect compensation for variance risk, one might expect that it would increase when risk premia were larger. The insignificance of the interaction term therefore suggests that nontrading effects do not stem from an aversion to variance risk. Nevertheless, we explore this possibility more thoroughly in the following section.

III. Risk and Risk Aversion

Having established a nontrading effect in option portfolios of all types, we now address whether the differential between trading and nontrading periods is likely the result of differences in risk or risk aversion. In the results to follow, we document that risk is indeed somewhat higher over nontrading periods, which raises the possibility that the nontrading effect is attributable to risk. Unfortunately, risk adjustment for option returns is notoriously difficult, mainly due to the fact that option returns are highly non-Gaussian. We therefore employ a number of techniques for risk adjustment that, while individually imperfect, together make a compelling case that risk cannot ultimately explain our findings.

We proceed by first analyzing the behavior of option trading volume and open interest around nontrading periods, motivated by the idea that differences in risk or risk aversion over nontrading periods should lead to predictable patterns in trading and open interest as traders rebalance their positions. We then assess the risk of option returns, first by analyzing the time series, then by examining the cross section. We find no patterns in volume or open interest that would indicate that traders or market makers wish to reduce exposure prior to nontrading periods. In the time series, we document that, while risk is higher over nontrading periods, various measures of risk-adjusted performance are in agreement that the greater risk is not sufficient to justify the size of the nontrading returns. There is also no apparent time series variation in liquidity. Finally, in the cross section, we find that only certain types of option risk are rewarded in terms of more negative nontrading returns. These results undermine any explanation based on the Chen and Singal (2003) hypothesis that aversion to downside risk is higher over nontrading periods. Instead, we show that the nontrading effect is consistent with the particular type of mispricing considered by French (1984).

A. Patterns in Volume and Open Interest

If option writers are averse to holding their positions over nontrading periods, then it should be reasonable to expect their positions to shrink prior to those periods. Because options are in zero net supply, a reduction in written positions means a decrease in the overall open interest in the options market. Furthermore, if option writers are particularly averse to maintaining written positions in some types of contracts, say, those that are riskier in some sense, then we might expect the decline to be largest in those contracts.

To examine the behavior of open interest around nontrading periods, we compute daily open interest, in terms of the number of option contracts, for each of the portfolios analyzed in Table III.[14] In doing so, we do *not* impose data filters, such as the requirement of five consecutive days of trading activity. We find that open interest is not prone to measurement error, making these filters unnecessary. More importantly, the variables on which filters are based (i.e., volume) are too closely related to the open interest series we are tracking. The only filter we impose is that contracts have more than five days remaining until expiration. This eliminates option contracts in their final week, which would artificially induce a nontrading effect due to the fact that expiring contracts have zero open interest on the last day of their expiration week.

We compute two weekly open interest series for each portfolio formed by moneyness and maturity. One is the average open interest over all days just prior to a nontrading period, which most often is simply the open interest on the Friday of that week. The other is the average open interest over all other days (usually Monday through Thursday). We then take the ratio of these two series within each week and compute the time series median, which is not sensitive to occasional outliers in the open interest ratio that are due to denominators near zero.[15] This median tells us how open interest rises or falls for a particular portfolio prior to nontrading periods.

The top panel of Figure 2 plots these medians, on the horizontal axis, against the estimated nontrading effects from Table III. If option writers are averse to keeping positions open over nontrading periods, we would expect to see two results. One is that open interest should decline prior to nontrading periods. The other is that open interest should decline more in contracts that are particularly undesirable to write, which presumably are those with the most negative nontrading effects. However, the figure shows no evidence favoring either prediction. Open interest tends to be higher just prior to nontrading periods, by around 2% on average. Furthermore, there is no relation between open interest and the size of the nontrading effect. Thus, if the nontrading effect is the result of a desire to avoid risk over weekends and holidays, no attempt to avoid that risk by trimming positions is evident.

The bottom panel of Figure 2 plots the results of a similar examination of trading volume. Since these ratios are all below one, it does not appear that investors are motivated to reduce positions prior to nontrading periods. Furthermore, there is no clear relation between volume patterns and nontrading returns.

It seems more likely that the decline in option trading is driven by the same forces that have caused equity market volume to decline prior to weekends and holidays for a number of decades (see Lakonishok and Maberly (1990)). Interestingly, the nontrading effect in equities documented by French (1980) has vanished in recent decades, while the volume pattern has not, making it somewhat implausible that there is a causal relation between volume and nontrading returns in stocks either.

[14] Results are robust to using dollar open interest instead.

[15] For most portfolios, there are no such outliers, and there is little difference between means and medians.

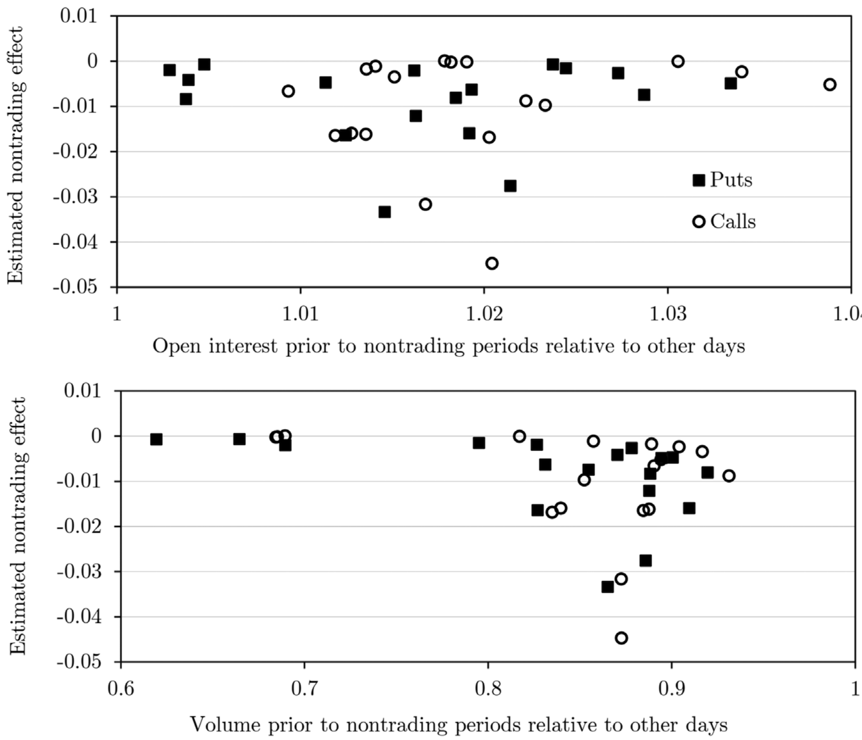

Figure 2. Volume, open interest, and the nontrading effect. The plots show the relation between the estimated nontrading effect and volume and volatility. For each of the maturity and moneyness-sorted portfolios analyzed in Table III, we compute the median ratio of either volume or open interest on days prior to nontrading periods to the same variable on all other days. These ratios appear on the horizontal axis. On the vertical axis, we show the nontrading effect estimates from Table III.

B. Risk and Liquidity of Option Strategies

Table VI reports a variety of statistics describing the risk and risk-adjusted returns corresponding to the main strategies examined thus far. One subtle change is that we now examine excess returns from the perspective of the option writer, whose returns are positive when option values decline. Flipping the sign of the option returns is necessary for the validity of some of the performance adjustments we consider, which account for asymmetries in the return distribution.

The first few rows of the table describe the first four moments of the three strategies we focus on. Since option writers are short options, the nontrading effect is shown in the table as a positive number, but the magnitudes are the same as those reported in Table II. We now see that these higher means are associated with moderately higher standard deviations, somewhat more negative skewness, and substantially higher excess kurtosis. Few of these differences are statistically significant, with the exception of the standard deviation and

Table VI

Risk and Risk-Adjusted Performance

This table reports measures of risk and return and of risk-adjusted performance of option writing strategies. The option portfolios examined are the same as in Table II except that they are from the perspective of an option writer rather than buyer. CE (gamma = γ) refers to the certainty-equivalent rate for a CRRA investor with a risk aversion of γ. P^{AS} and P^{FH} are the performance indices of Kadan and Liu (2014) using the risk measures of Aumann and Serrano (2008) and Foster and Hart (2009), respectively. Except for the means, all parameters are estimated using exactly identified GMM with a Newey-West (1987) covariance matrix. t-Statistics are in parentheses.

	Written Unhedged Puts and Calls			Written Hedged Puts			Written Hedged Calls		
	Trading	Nontrading	Difference	Trading	Nontrading	Difference	Trading	Nontrading	Difference
Average excess return	0.0008	0.0076	0.0068	−0.0003	0.0082	0.0086	−0.0003	0.0048	0.0051
	(1.76)	(9.30)	(7.86)	(−0.81)	(11.91)	(11.93)	(−0.69)	(5.08)	(5.11)
Standard deviation	0.0231	0.0253	0.0022	0.0196	0.0218	0.0023	0.0218	0.0322	0.0103
	(31.49)	(16.00)	(1.60)	(24.03)	(12.53)	(1.41)	(10.49)	(6.40)	(3.05)
Skewness	−1.5865	−2.1887	−0.6022	−2.1909	−2.7998	−0.6089	−1.2331	−2.9260	−1.6929
	(−9.07)	(−4.88)	(−1.29)	(−6.71)	(−3.77)	(−0.76)	(−2.10)	(−1.95)	(−0.84)
Kurtosis	9.4588	15.6314	6.1726	16.4390	25.7372	9.2982	32.3123	56.1377	23.8254
	(6.82)	(4.73)	(1.72)	(4.81)	(4.19)	(1.31)	(4.34)	(4.78)	(2.17)
Sharpe ratio	0.0346	0.3010	0.2664	−0.0172	0.3763	0.3935	−0.0125	0.1499	0.1624
	(1.70)	(6.73)	(6.09)	(−0.82)	(7.05)	(7.77)	(−0.71)	(3.41)	(3.64)
Alpha	0.0006	0.0076	0.0070	−0.0005	0.0083	0.0088	−0.0008	0.0049	0.0057
	(1.41)	(10.01)	(8.67)	(−1.36)	(13.22)	(13.04)	(−2.38)	(7.07)	(7.62)
Beta	0.4406	0.5598	0.1192	0.4912	0.6633	0.1721	1.1988	1.7188	0.5200
	(3.52)	(3.93)	(0.92)	(8.56)	(6.05)	(1.65)	(11.04)	(7.49)	(3.29)
Information ratio	0.0275	0.3173	0.2898	−0.0289	0.4163	0.4452	−0.0471	0.2264	0.2734
	(1.38)	(7.62)	(6.97)	(−1.39)	(8.89)	(9.74)	(−2.57)	(4.70)	(5.57)
CE (gamma = 1)	0.0006	0.0075	0.0069	−0.0005	0.0082	0.0087	−0.0004	0.0045	0.0049
	(1.38)	(8.90)	(8.02)	(−1.23)	(11.37)	(11.96)	(−1.18)	(4.35)	(4.22)
CE (gamma = 3)	0.0000	0.0068	0.0068	−0.0009	0.0077	0.0085	−0.0010	0.0031	0.0040
	(0.05)	(7.52)	(7.37)	(−2.20)	(9.79)	(10.86)	(−2.30)	(2.13)	(2.71)
CE (gamma = 10)	−0.0023	0.0035	0.0058	−0.0027	0.0050	0.0077	−0.0033	−0.0141	−0.0108
	(−4.02)	(2.46)	(4.14)	(−4.89)	(3.67)	(5.61)	(−4.24)	(−1.14)	(−0.90)
P^{AS}	0.0288	0.1526	0.1238	−0.0180	0.1763	0.1943	−0.0115	0.0634	0.0749
	(1.86)	(5.70)	(4.38)	(−0.91)	(5.49)	(5.60)	(−0.75)	(2.97)	(2.81)
P^{FH}	0.0274	0.0503	0.0229	−0.0182	0.0462	0.0644	−0.0113	0.0241	0.0354
	(1.90)	(32.09)	(1.60)	(−0.93)	(28.72)	(3.36)	(−0.80)	(14.47)	(2.47)

kurtosis of the hedged call portfolio, but they nevertheless raise the possibility of a relation between risk and return.[16]

When we undertake the kind of risk adjustment that would be appropriate in a Gaussian setting or for investors with mean-variance preferences, we see that the greater risks associated with nontrading returns have little effect on the relative attractiveness of writing options over nontrading rather than trading periods. Sharpe ratios, CAPM alphas, and information ratios are all significantly higher over nontrading periods, with magnitudes that are economically large. The Sharpe ratio of the unhedged put-call strategy, for example, is 0.301. This daily value corresponds to an annualized figure of 4.9.[17]

It is well understood, however, that measures such as the Sharpe ratio are generally inappropriate when returns are highly non-Gaussian. This is certainly the case in options, where a nonlinear dependence on the underlying asset will induce nonnormality even when the underlying is Gaussian. A number of alternatives have been proposed, and we consider two. The first involves specifying a utility function and measuring the certainty-equivalent return, which is defined as the value c such that $\mathrm{E}[\mathrm{U}(1+R)] = \mathrm{U}(1+c)$, where R is the return on the risky asset. We implement this approach using power utility with relative risk aversion γ of 1 (log utility), 3, or 10.

The second approach is based on recent work by Kadan and Liu (2014), who show how to evaluate investment performance in a way that accounts for higher moments by using the generalized risk measures of Aumann and Serrano (2008) and Foster and Hart (2009). The performance indexes they construct, P^{AS} and P^{FH}, are the solutions to

$$\mathrm{E}[\exp(-P^{AS}(1+R))] = 1$$

and

$$\mathrm{E}[\log(1 + P^{FH}(1+R))] = 0.$$

As with the other statistics in the table, we estimate them using the generalized method of moments (GMM). Higher values indicate superior performance.

The result of this analysis is that non-Gaussian risk adjustment still cannot explain the nontrading effect. For all three levels of γ, we find that risk-adjusted returns are higher over nontrading periods. The sole exception is in the hedged call portfolio when $\gamma = 10$. In this case, risk-adjusted returns are lower over nontrading periods, though by a statistically insignificant amount. Based on the P^{AS} and P^{FH} indexes, risk-adjusted performance over nontrading periods is

[16] Throughout Table VI, estimates and t-statistics are the result of exactly identified GMM estimation. t-statistics for nontrading minus trading differences are obtained by estimating a system that includes both trading and nontrading returns and that parameterizes the nontrading parameter (e.g., kurtosis or alpha) as the trading parameter plus a "delta" term.

[17] It is important to note that none of the calculations in the table factor in the effects of transaction costs, both commissions and bid-ask spreads. Therefore, the returns are not feasible as a stand-alone trading strategy. Instead, they represent the additional return that would result from holding previously written options for another day.

generally significantly better than that over trading periods. The only exception is for the P^{FH} index for unhedged options, where nontrading performance is better but not significantly so.

Overall, the results reported in Table VI suggest that no measure of risk varies enough between trading and nontrading periods to explain the differences in average returns. It is possible that, due to the existence of unobserved low-probability events, true risk over nontrading periods is higher than we have measured, and if the increase is severe enough, it might justify a substantial risk premium. As an informal check for the unhedged put-call portfolio, we randomly replace 20 nontrading returns with the worst one-day return in the full sample (-20.3%). Since the sample contains only 16 nontrading periods with returns lower than -10%, this represents an extremely large increase in the left tail of the return distribution. Yet, even with this magnified tail risk, both P^{AS} and P^{FH} remain reliably positive.

C. Risk and Return in the Cross Section

Thus far in this section, we have analyzed differences between trading and nontrading periods in highly aggregated portfolios of options. We now turn to disaggregated portfolios and further explore some of the patterns related to moneyness and maturity that we observe in Table III. Our goal is to see whether those patterns in the size of the nontrading effect are related to cross-sectional patterns in risk.

Option market participants and option researchers often identify three major varieties of price risk in an option contract. The first, delta ($\Delta \equiv \partial C/\partial S$), is the sensitivity of the option price to small movements in the underlying security. This risk can be hedged away by taking the appropriate position in the underlying. The second, vega ($\mathcal{V} \equiv \partial C/\partial \sigma$), measures the sensitivity of the option price to changes in volatility. This risk can be eliminated only by combining purchased and written options on the same underlying, though it can be diversified somewhat by buying or writing options on multiple underlyings. The last risk, gamma ($\Gamma \equiv \partial^2 C/\partial S^2$), measures the sensitivity of the option price to larger movements in the price of the underlying security. Like vega, gamma is reduced mainly by buying and writing options on the same underlying, though diversification across underlyings should also provide some benefit.

These so-called Greeks are the coefficients of the following Taylor expansion of option prices:

$$C_{t+1} \approx C_t + \Delta_t (S_{t+1} - S_t) + \frac{1}{2}\Gamma_t (S_{t+1} - S_t)^2 + \mathcal{V}_t (\sigma_{t+1} - \sigma_t) + \Theta_t,$$

where S is the price of the underlying, σ is its volatility, and $\Theta_t \equiv \partial C_t/\partial t$ is the rate of time decay, or theta. This expansion can be rearranged to produce an approximate expression for delta-hedged option returns:

$$\frac{C_{t+1} - C_t}{C_t} - \frac{\Delta_t S_t}{C_t} R_{t+1} \approx \frac{\Gamma_t S_t^2}{2C_t} R_{t+1}^2 + \frac{\mathcal{V}_t}{C_t} (\sigma_{t+1} - \sigma_t) + \frac{\Theta_t}{C_t}, \tag{4}$$

where R_{t+1} is the return on the stock.[18] Expression (4) shows how delta-hedged returns depend on "larger" movements in prices (R_{t+1}^2) and changes in volatility ($\sigma_{t+1} - \sigma_t$). Specifically, this expression shows that these variables affect delta-hedged returns through normalized versions of gamma and vega.

The risk of a diversified option portfolio therefore depends, in general, on both option contract-level characteristics (the normalized Greeks) and stock-level risk exposures (the systematic risk in squared returns and volatilities). However, in the current exercise, we examine portfolios formed only on the basis of maturity and moneyness, which are both option-level characteristics, so all portfolios will include essentially the same set of underlying firms. Thus, the cross-sectional variation in the risk of the portfolios we consider will be determined largely by option-level characteristics rather than stock-level risk exposures.

Table VII reports results of Fama-Macbeth (1973) regressions in which the dependent variable is the return, squared return, or cubed return of a portfolio of delta-hedged option positions. The independent variables include the portfolio's normalized gamma ($0.5\,\Gamma_t S_t^2/C_t$), vega ($\mathcal{V}_t/C_t$), and theta (Θ_t/C_t), which are computed by averaging the corresponding contract-level measures. As before, we weight by the lagged dollar value of open interest when computing portfolio returns. The same weights are also used to calculate portfolio-level normalized Greeks.

Because of the close connection between gamma and theta, which we discuss below, these two regressors are highly collinear and are not included in the same regression. Gamma and vega are also related and, in fact, are perfectly proportional to each other for a given expiration date. Since our option portfolios include both short-dated and long-dated contracts, gamma and vega are not collinear, and we are able to distinguish between their effects as two individual predictors.

In our regressions, independent variables are included alone and are also interacted with a nontrading dummy variable that takes the value of zero during trading periods and one during nontrading periods. Since all independent variables are interacted with the nontrading dummy, the Fama-Macbeth (1973) regression coefficients are estimated by running the regression with no interaction terms separately on the trading and nontrading subsamples. The interaction term's coefficient is then estimated as the difference between the nontrading and trading coefficients.

Our first set of cross-sectional results focuses on risk. Regressions (1) to (3) in Table VII examine the cross section of risk by regressing squared delta-hedged portfolio returns, which represent price changes regardless of their sign, on normalized gammas and vegas. When we include only gamma (by itself and interacted with the nontrading dummy), we see a significant tendency of higher gamma options to exhibit greater risk, though there is no significant evidence

[18] For simplicity, in this section, we use regular returns rather than excess returns, which would introduce an additional term into the return decomposition. Results are extremely similar when excess returns are used.

Table VII

Cross-Sectional Determinants of Risk and Return

This table reports the results of Fama-Macbeth (1973) regressions of portfolio returns, squared returns, and cubed returns on various normalized Greeks and their interactions with a nontrading dummy. The portfolios considered are the double-sorted hedged portfolios analyzed in Table III except that returns are not in excess of the riskless rate. The gamma regressor is the dollar-open interest-weighted normalized gamma, defined as $\Gamma S^2/2P$, where Γ is the Black-Scholes gamma, S is the stock price, and P is the option price. The vega and theta regressors are the dollar-open interest-weighted normalized vega and theta, $\mathcal{V}/P$ and Θ/P, where $\mathcal{V}$ and Θ are the Black-Scholes vega and theta, respectively. Each Greek is included alone and interacted with a nontrading dummy variable. Coefficient estimates for the gamma and vega terms have been multiplied by 1,000. t-statistics are in parentheses.

	Squared Portfolio Returns			Cubed Portfolio Returns			Portfolio Returns				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Intercept	−0.0012	−0.0027	−0.0027	−0.0001	−0.0005	−0.0006	0.0017	0.0008	0.0008	0.0021	0.0007
	(−10.37)	(−17.03)	(−17.37)	(−0.41)	(−4.59)	(−4.56)	(6.49)	(4.10)	(4.31)	(7.74)	(4.84)
× Nontrading	−0.0008	−0.0008	−0.0008	−0.0004	0.0001	0.0000	−0.0017	−0.0006	−0.0008	−0.0022	−0.0006
	(−2.36)	(−1.61)	(−1.66)	(−0.82)	(0.15)	(−0.06)	(−3.19)	(−1.41)	(−1.95)	(−3.77)	(−1.77)
Gamma	0.1995		0.1304	0.0301		0.0095	−0.0008		−0.0432		−0.0389
	(7.99)		(8.13)	(3.57)		(0.81)	(−0.04)		(−2.27)		(−1.90)
× Nontrading	0.1495		0.1292	0.0036		0.0231	−0.2970		−0.2647		−0.2718
	(1.46)		(1.73)	(0.10)		(0.57)	(−5.97)		(−5.69)		(−5.76)
Vega		1.6765	0.8011		0.3120	0.2320		0.1140	0.4385		0.4907
		(12.25)	(8.03)		(4.83)	(2.49)		(0.71)	(3.35)		(3.73)
× Nontrading		0.7940	0.0920		0.0783	−0.1894		−2.0931	−0.3063		−0.3564
		(1.53)	(0.36)		(0.30)	(−1.12)		(−6.35)	(−1.01)		(−1.19)
Theta										0.0590	
										(2.21)	
× Nontrading										0.4893	
										(8.14)	
Theta-interest rate component only											0.6340
											(0.44)
× Nontrading											1.8710
											(0.59)
Cross-sectional R^2	0.362	0.566	0.609	0.609	0.594	0.733	0.804	0.222	0.851	0.756	0.900

that gamma-related risk rises over nontrading periods. The relation between risk and normalized vega appears to be much stronger and more significant, with a notably higher cross-sectional R^2, indicating that vega (i.e., volatility) risk is more important than gamma risk in our sample.[19] Including both gamma and vega terms in regression (3) leads to a modest improvement in overall fit relative to the vega-only specification in (2). Thus, overall, it appears that vega rather than gamma is more responsible for portfolio variance.

In columns (4) to (6) of the table, we repeat the same regressions with cubed returns as the dependent variable. The rationale for using cubed returns is that Chen and Singal's (2003) hypothesis suggests that it is specifically the unbounded downside risk that option sellers are averse to holding over nontrading periods, and this may be reflected more in third moments than in second moments. The results show a tendency of higher gamma and higher vega options to exhibit greater positive skewness, which for the option writer becomes negative skewness. However, there is no evidence that this effect changes over nontrading periods.

In regressions (4) and (5), which include either gamma or vega terms, but not both, we see approximately the same goodness of fit and similar statistical significance of the slope coefficients. When both gamma and vega terms are included in regression (6), vega is significant but gamma is not, suggesting that volatility risk may also be the more relevant determinant of the option writer's downside risk.

The last five columns of Table VII examine whether the same variables also drive average returns. Regressions (7) to (9) replace the dependent variable with signed delta-hedged portfolio returns. Distinct from the other regressions, we now see a clear difference between the importance of gamma and vega, with the former turning out to be the more important determinant by far. In regression (7), which does not include vega, we obtain a cross-sectional R^2 of 0.80. In contrast, regression (8) shows that vega by itself performs relatively poorly in terms of explaining average returns, with an R^2 of just 0.22. When gamma and vega terms are both included, as in regression (9), we continue to see a large and highly significant coefficient on the interaction between the nontrading dummy and the normalized gamma, and just a slight improvement in overall fit relative to the gamma-only specification.[20] Furthermore, the nontrading interaction with vega is insignificant.

In the demand-based option pricing theory of Garleanu, Pedersen, and Poteshman (2009), fluctuations in demand impact option prices only to the

[19] The cross-sectional R^2 is defined as $1 - [\frac{1}{N}\sum_{i=1}^{N}(\bar{Y}_i - \bar{\hat{Y}}_i)^2]/[\frac{1}{N}\sum_{i=1}^{N}(\bar{Y}_i - \bar{\bar{Y}})^2]$, where $\bar{Y}_i$ is the average value of the dependent variable for portfolio i, $\bar{\hat{Y}}_i$ is the average fitted value for that portfolio, and $\bar{\bar{Y}}$ is the cross-sectional average of $\bar{Y}_i$.

[20] The relevance of vega in explaining average returns is not new. Volatility risk has been shown by Bakshi, Cao, and Chen (1997) and others to be an important determinant of expected returns on S&P 500 Index options. However, evidence tying volatility risk to the expected returns on individual equity options is less clear-cut (see, e.g., Driessen, Maenhout, and Vilkov (2009) and Christoffersen, Fournier, and Jacobs (2016)). Controlling for a nontrading effect may be useful in uncovering this risk premia.

extent that the option contract cannot be perfectly hedged. Thus, regardless of whether the nontrading effect is driven by greater aversion to risk over nontrading periods or by an uptick in net option demand, we would expect riskier option contracts to display more negative nontrading returns. In contrast, regressions (1) to (9) show that, while the risk of a diversified portfolio of delta-hedged options on individual equities is determined more by its average vega than its average gamma, the nontrading effect in average returns is related only to gamma. This result undermines the hypothesis put forth by Chen and Singal (2003) that aversion to downside risk over nontrading periods drives nontrading effects in average returns. Nevertheless, we cannot completely rule out the possibility that, for some reason, one type of risk (gamma) is priced over nontrading periods, while another type (vega) is not.

Another interpretation of these results is that they reflect mispricing caused by an incorrect treatment of the difference between trading and nontrading volatilities. As shown in Section II.C, options appear to be priced as if nontrading periods contained, on average, an extra day's worth of variance, raising the possibility that some investors are pricing options as though variance was proportional to calendar time rather than trading time. As demonstrated by French (1984), if this is, in fact, the error being made, then its effects will depend strongly on the moneyness and maturity of the contract.

To gain some intuition about what effects this mistake might have, assume that stock prices are instantaneously lognormal with deterministic volatility. Assuming no dividends, the price C_t of an option therefore solves the following partial differential equation:

$$\frac{\sigma_t^2 S_t^2}{2}\Gamma_t + \Theta_t + rS_t\Delta_t = rC_t.$$

If the interest rate were equal to zero, this equation implies an exact relation between theta and gamma, namely,

$$\Theta_t = -\frac{\sigma_t^2 S_t^2}{2}\Gamma_t.$$

In other words, time decay would be due solely to the loss in stock price variance $(\sigma_t^2 S_t^2)$ remaining until expiration, the value of which is determined by the amount of curvature (Γ_t) in the option price.

When interest rates are nonzero, theta includes an additional component, which we discuss below. Although we do not assume zero interest rates in our empirical analysis and therefore include this component, the low levels of interest rates observed during our sample period means that this component is in general extremely small. In theory, therefore, we should see little if any relation between the theta of an option and its risk-neutral drift, since theta is offset by a gamma effect that is approximately equal in expected value but of the opposite sign.

Intuitively, theta comes largely from the fact that the total volatility remaining until expiration decreases from one trading day to the next, and with less volatility remaining over the life of the contract, there is a less of a chance to benefit from positive convexity. By definition, however, the total variance until expiration is reduced exactly by the variance that is realized between these two dates, which, due to positive gamma, tends to lead to higher option values. These effects should cancel out.

In practice, however, theta and gamma will not offset each other if the option pricing model used is inconsistent with actual stock price dynamics. This would occur, for instance, if volatility is assumed constant in calendar time. As French (1984) shows theoretically, an option trader using a calendar-time model implicitly assumes that three days of time decay occur between a given Friday and the following Monday. Because actual variance over this three-day period is much lower than the variance of three trading days, the convexity effect induced by the option's positive gamma will be muted and will only partially offset the high amount of time decay. The result is a predictable decline in option values over the weekend or any other multiday period that includes one or more nontrading days.

Because theta is driven almost entirely by gamma effects rather than interest rates, normalized gamma and theta are highly collinear in our sample. We find an average cross-sectional correlation between normalized gamma and theta of -0.95 at the portfolio level. Thus, the evidence in regressions (7) to (9), which shows that gamma appears solely responsible for the nontrading effect, is also consistent with nontrading effects being the result of excess time decay. Nevertheless, to verify that the nontrading effect can be attributed to theta, we include a regression with only the normalized theta and its interaction with the nontrading dummy as independent variables.

The results, reported in column (10) of the table, show that about 76% of the variation in average returns can be attributed to theta. The interaction term on theta is highly significant, more so than any other interaction term in the table. Thus, the nontrading effects that we document appear closely related to theta, suggesting that excess time decay over nontrading periods is a primary cause of our main findings. That is, market prices embed a rate of time decay that is too large over nontrading periods because option traders do not account sufficiently for the difference in stock volatility when the market is open or closed.

The finding that nontrading effects are closely tied to theta helps explain the patterns observed in Table III, namely, the tendency of the nontrading effect to be stronger in options that are further out of the money and closer to expiration. In the Black and Scholes (1973) model, gamma shrinks as an option moves more out of the money. However, the option's price shrinks even faster. The result is that the normalized theta (Θ_t/C_t) is largest for out-of-the-money contracts. Intuitively, these contracts only have value if the underlying price moves significantly, which rapidly becomes less likely as time to maturity declines. Normalized theta also tends to be larger for options with short maturities. This is driven by the fact that the passage of a fixed amount of time represents a larger fraction of the total time remaining until maturity.

Finally, to rule out the possibility that the results in these regressions are driven by an omitted interest rate effect, we estimate a final regression in which we include the interest rate component of theta along with our main regressors. As discussed above, when interest rates are nonzero (still assuming no dividends), time decay (Θ_t) is the sum of two components,

$$\Theta_t = -\frac{\sigma_t^2 S_t^2}{2}\Gamma_t - \frac{r_t \exp(-r_t\tau)K\Phi(d_2)}{C_t} \equiv \Theta_t^\sigma + \Theta_t^r, \tag{5}$$

where r_t is the instantaneous riskless rate. The first component, as discussed above, captures convexity effects. It depends very little on interest rates, only indirectly through the option's gamma.[21] The second component, Θ_t^r, can be interpreted as a pure interest rate component, as it reflects the cost of borrowing in the option's replicating portfolio. It is not collinear with gamma, so in regression (11), we add this "interest rate theta" to the gamma and vega terms already considered.

The results of this last regression show that including interest rate effects has little bearing on our findings. Coefficients on normalized gamma and vega, alone and interacted with the nontrading indicator, are relatively unchanged from their values in regression (9). Vega continues to play a role in explaining trading-period returns, while gamma drives all significant nontrading effects. This tells us that the role of gamma in explaining the nontrading effect is not due to a failure to account for differences in the cost of leverage between trading and nontrading periods.

Interestingly, while the extremely limited variation in Θ_t^r makes precise inference difficult, the magnitude of the coefficients on Θ_t^r and its nontrading interaction are consistent with the treatment of interest over trading and nontrading periods. Specifically, the coefficient on the interest component over trading periods is insignificantly different from 1.0, consistent with one day of interest accruing over these periods. Over nontrading periods, the coefficient on the interest component of theta is larger by approximately 2.0. Although this estimate is insignificant, the magnitude is consistent with the fact that most nontrading periods, being regular weekends, result in the payment or receipt of two additional days of interest.

IV. Conclusion

The evolution of modern option markets has paralleled the development of option pricing models. While no existing model is capable of explaining all observed patterns in option prices, models are irreplaceable tools for options hedging and market making activities. Yet, periodically, these models turn out

[21] Under Black and Scholes (1973), the sensitivity of gamma to interest rates is fairly low for most options and negligible for short-term options, where nontrading effects are largest. In our data sample, the variation in portfolio-level Θ_t^σ is driven largely by the absolute value of option moneyness, with a smaller role for option maturity. Interest rates, in contrast, explain essentially no variation.

to be wrong, as was the case for Black and Scholes (1973) in the October 1987 crash.[22] Although the sophistication of option markets has advanced tremendously since then, there is no guarantee that the models in use today are free from such potential shortcomings.

We believe that our results suggest a somewhat less dramatic but nevertheless quantitatively important model misspecification. At the end of our sample period, the total dollar open interest in exchange-traded equity option contracts was about \$300 billion, which is just slightly larger than the average value over our entire sample. Assuming an average alpha of about 0.50% over a single nontrading period, the average dollar transfer from option buyers to option writers is likely close to \$1.5 billion over the course of a single weekend. For many option contracts, particularly those that are out of the money and relatively short dated, the alpha is likely much larger.[23]

The finding of a nontrading effect is highly significant and pervasive among puts and calls of nearly every combination of maturity and moneyness. This finding is robust to the choice of sample period, method of portfolio construction, sampling method, and weighting scheme. We find strong evidence of a nontrading effect in S&P 500 Index put returns and weaker evidence in index call returns. Implied volatilities of both equity and index options embed an excessively high expectation of variance over nontrading periods.

Option portfolio risk appears to be moderately higher over nontrading periods, but after risk adjustment, we find that written option positions offer nontrading returns that are still extremely attractive and significantly better than trading-period returns. We interpret these results as evidence against a risk-based explanation for our main findings.

In the cross section, options with higher gammas show stronger nontrading effects. Since the risk of hedging an option portfolio rises with gamma, this relation seems to suggest that greater aversion toward risk by option writers over nontrading periods could explain why nontrading returns are low. Hedging risk also rises with vega, but we find no evidence that higher vegas are related to average nontrading returns. Thus, our findings appear to contradict the Chen and Singal (2003) hypothesis that investors have greater aversion to downside risk over weekends or other nontrading periods, and also appear inconsistent with a shift in option demand, as in Garleanu, Pedersen, and Poteshman (2009).

The alternative hypothesis that we test is that some option traders rely on models that do not account for differences in the behavior of stock prices over trading and nontrading periods. Following French (1984), option prices computed in calendar time based on the assumption that volatility is constant will mechanically fall over nontrading periods given that actual volatility drops significantly when markets are closed. This will appear in the data as an

[22] See Coval, Pan, and Stafford (2014) for a compelling account of this episode.

[23] As an admittedly unfair comparison, large-cap stocks with positive momentum comprised roughly \$6 trillion in value at the start of our sample and produced an alpha of roughly one basis point per day, implying a daily dollar alpha of \$0.6 billion.

excessive rate of time decay. We find in our sample that the nontrading effect is indeed strongly related to the rate of time decay.

Since bid-ask spreads are often large, it is not feasible to capture the entire nontrading effect in a stand-alone trading strategy based on market orders. That said, Muravyev and Pearson (2016) show that transaction costs by some traders who take liquidity (i.e., trade with market rather than limit orders) may be far smaller than the posted bid-ask spread because of superior trade timing. These traders, who most likely use automated execution algorithms, pay costs only one-fifth as large as those implied by the bid-ask quotes.

Regardless of trading costs, nontrading effects have important implications for quote-setting by market makers or other traders and are a determinant of optimal trade timing. Our findings suggest that setting aggressive ask prices and conservative bid prices would be optimal prior to nontrading periods as a means to reduce inventory ahead of declining option values. For an investor who has already written an option, the nontrading effect represents an additional return that would result from keeping the position open for another day. For an option buyer, the nontrading return is the benefit from delaying the purchase by one day. Thus, our results should be useful in improving option-based trading strategies.

Our results also imply that it is important to control for nontrading effects in event studies of option returns, such as Cao, Chen, and Griffin (2005) or Xing and Zhang (2013), if events have a tendency to occur on a certain day of the week. Controlling for nontrading effects could also improve the research design of studies that seek to explain option pricing regularities such as upward-sloping term structures and smile-shaped cross sections in Black and Scholes (1973) implied volatilities. To maximize power and to best approximate continuous-time theoretical models, this literature commonly analyzes option data on a daily basis (e.g., Bakshi, Kapadia, and Madan (2003) or Broadie, Chernov, and Johannes (2007)). At a daily frequency, the effects we document are likely larger than any risk premia for stochastic volatility or jump risk, for example. It is therefore possible that controlling for nontrading effects explicitly or by analyzing data at a lower frequency could have a material effect on the results reported by previous studies.

Moreover, we believe that our results should encourage researchers to consider market inefficiency as a potential explanation for option pricing anomalies. Though there are notable exceptions, including the behavioral papers of Stein (1989), Poteshman (2001), and Han (2008), empirical shortcomings of a particular option pricing model are almost always taken as motivation for an improved stochastic process for the underlying price. While this progression makes sense in many cases, our results suggest that an alternative approach is sometimes called for.

If the effects we document are due to the widespread use of a misspecified model, there will be no quick fixes. Indeed, recognition of the issue that we empirically document appears to go back at least 30 years to the publication of French (1984). Nevertheless, in the options market, a relatively small number of highly influential entities provide a large amount of information to market participants, and if the current practices of these market leaders were changed, a more efficient market might obtain.

As an example, the current VIX Index, which is computed by the Chicago Board Options Exchange (CBOE, 2009), is based on a calendar time convention that causes a predictable drop in the level of the index on Mondays. This is despite the fact that the original VIX, now known as the VXO, was based on trading time. In fact, Whaley (2000, p. 14), who was retained by the CBOE to develop the VIX, noted in a paper that

> the VIX is based on trading days. If the time to expiration is measured in calendar days, the implied volatility is a volatility rate per calendar day. This means, among other things, that the return variance of the OEX index over the weekend (from Friday close to Monday close) should be three times higher than it is over any other pair of adjacent trading days during the week (say, Monday close to Tuesday close). Empirically, this is not true.

Surprisingly, this insight seems to have been lost in the 2003 revision of the VIX methodology.

Standard academic databases (Optionmetrics IvyDB) and industry data feeds (ISE/Hanwick) currently report option-implied volatilities and Greeks on a calendar time basis or do not specify the convention used for measuring time to expiration. French's (1984) work shows that measuring time in trading days, while not perfect, is far superior to the use of calendar days. We believe that adopting this standard and highlighting the issues involved could improve the efficiency of the options market in a meaningful way.

Initial submission: February 4, 2016; Accepted: October 14, 2016
Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

REFERENCES

Andersen, Torben G., Tim Bollerslev, Frank X. Diebold, and Heiko Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.

Aumann, Robert J., and Roberto Serrano, 2008, An economic index of riskiness, *Journal of Political Economy* 116, 810–836.

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

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 2015, Option pricing and hedging performance under stochastic volatility and stochastic interest rates, in Cheng F. Lee and John Lee, eds.: *Handbook of Financial Econometrics and Statistics* (Springer, Berlin/New York).

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

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

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

Blume, Marshall E., and Robert F. Stambaugh, 1983, Bias in computed returns: An application to the size effect, *Journal of Financial Economics* 12, 387–553.

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

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.

Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56, 1–43.

Cao, Charles, Zhiwu Chen, and John M. Griffin, 2005, Informational content of option volume prior to takeovers, *Journal of Business* 78, 1073–1109.

Cetin, Umut, Robert A. Jarrow, Philip E. Protter, and Mitch Warachka, 2006, Pricing options in an extended Black Scholes economy with illiquidity: Theory and empirical evidence, *Review of Financial Studies* 19, 493–529.

Chang, Eric C., J. Michael Pinegar, and R. Ravichandran, 1993, International evidence on the robustness of the day-of-the-week effect, *Journal of Financial and Quantitative Analysis* 28, 497–513.

Chen, Honghui, and Vijay Singal, 2003, Role of speculative short sales in price formation: The case of the weekend effect, *Journal of Finance* 58, 685–706.

Chicago Board Options Exchange (CBOE), 2009, *The CBOE Volatility Index - VIX*. Available at https://www.cboe.com/micro/vix/vixwhite.pdf.

Christoffersen, Peter, Mathieu Fournier, and Kris Jacobs, 2016, The factor structure in equity options, Rotman School of Management Working Paper No. 2224270.

Cliff, Michael, Michael J. Cooper, and Huseyin Gulen, 2008, Return differences between trading and non-trading hours: Like night and day, Working paper, University of Utah and Purdue University.

Connolly, Robert, 1989, An examination of the robustness of the weekend effects, *Journal of Financial and Quantitative Analysis* 24, 133–169.

Coval, Joshua D., Kevin Pan, and Erik Stafford, 2014, Capital market blind spots, Working paper, Harvard University.

Coval, Joshua D., and Tyler Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1010.

Dellavigna, Stefano, and Joshua M. Pollet, 2009, Investor inattention and Friday earnings announcements, *Journal of Finance* 64, 709–749.

Driessen, Joost, Pascal Maenhout, and Grigory Vilkov, 2009, The price of correlation risk: Evidence from equity options, *Journal of Finance* 64, 1377–1406.

Duarte, Jefferson, and Christopher S. Jones, 2008, The price of market volatility risk, Working paper, Rice University and University of Southern California.

Dubinsky, Andrew, and Michael Johannes, 2006, Earnings announcements and equity options, Working paper, Columbia University.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Fields, M.J., 1931, Stock prices: A problem in verification, *Journal of Business* 4, 415–438.

Foster, Dean P., and Sergiu Hart, 2009, An operational measure of riskiness, *Journal of Political Economy* 117, 785–814.

French, Dan W., 1984, The weekend effect on the distribution of stock prices: Implications for option pricing, *Journal of Financial Economics* 13, 547–559.

French, Kenneth R., 1980, Stock returns and the weekend effect, *Journal of Financial Economics* 8, 55–69.

French, Kenneth R., and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.

Garcia, Diego, 2013, Sentiment during recessions, *Journal of Finance* 68, 1267–1300.

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

Goyal, Amit, and Pedro Santa-Clara, 2003, Idiosyncratic risk matters! *Journal of Finance* 58, 975–1007.
Han, Bing, 2008, Investor sentiment and option prices, *Review of Financial Studies* 21, 387–414.
Hansis, Alexandra, Christian Schlag, and Grigory Vilkov, 2010, The dynamics of risk-neutral implied moments: Evidence from individual options, Working paper, Goethe University.
Hendershott, Terrence, and Mark S. Seasholes, 2006, Market maker inventories and stock prices, *American Economic Review Papers and Proceedings* 97, 210–214.
Hull, John, and Wulin Suo, 2002, A methodology for assessing model risk and its application to the implied volatility function model, *Journal of Financial and Quantitative Analysis* 37, 297–318.
Kadan, Ohad, and Fang Liu, 2014, Performance evaluation with high moments and disaster risk, *Journal of Financial Economics* 113, 131–155.
Kelly, Michael A., and Stephen P. Clark, 2011, Returns in trading versus non-trading hours: The difference is day and night, *Journal of Asset Management* 12, 132–145.
Lakonishok, Josef, and Edwin Maberly, 1990, The weekend effect: Trading patterns of individual and institutional investors, *Journal of Finance* 40, 231–243.
Longstaff, Francis A., 1995, How much can marketability affect security values? *Journal of Finance* 50, 1767–1774.
Muravyev, Dmitriy, and Neil D. Pearson, 2016, Option trading costs are lower than you think, Working paper, Boston College and University of Illinois at Urbana-Champaign.
Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
Patell, James M., and Mark A. Wolfson, 1981, The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices, *Journal of Accounting Research* 19, 434–458.
Poteshman, Allen M., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market, *Journal of Finance* 56, 851–876.
Scott, Louis O., 1997, Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Applications of Fourier inversion methods, *Mathematical Finance* 7, 413–426.
Sheikh, Aamir M., and Ehud I. Ronn, 1994, A characterization of the daily and intraday behavior of returns on options, *Journal of Finance* 49, 557–579.
Stein, Jeremy, 1989, Overreactions in the options market, *Journal of Finance* 44, 1011–1023.
Whaley, Robert E., 2000, The investor fear gauge, *Journal of Portfolio Management* 26, 12–17.
Xing, Yuhang, and Xiaoyan Zhang, 2013, Anticipating uncertainty: Straddles around earnings announcements, Working paper, Rice University and Purdue University.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.