Very Noisy Option Prices and Inference Regarding the Volatility Risk Premium

Jefferson Duarte, Christopher S. Jones, and Junbo L. Wang*

January 2022

Abstract

We show that volatility is negatively priced in individual equity options. Consistent with a negative volatility risk premium, the average return of the most heavily traded deep out-of-the-money call options on stocks is negative and economically significant at -73 basis points per day. After adjusting for microstructure biases, a Fama-MacBeth regression indicates that the volatility risk premium in individual equity options is about the same as the premium in S&P 500 Index call options. Our results highlight the importance of addressing microstructure biases when estimating expected returns and the risk premia of options.

Keywords: Options; Volatility risk premium; Microstructure bias

*We thank Torben Andersen, Hank Bessembinder, Greg Eaton, Bjørn Eraker, Mathieu Fournier, Ruslan Goyenko, Dmitriy Muravyev, Mobina Shafaati, Aurelio Vasquez, and seminar participants at the 2019 Latin America FMA Conference, 2019 CICF, 2019 FMA Conference, 2019 Conference on Derivatives and Volatility, 2020 SFS Cavalcade, and 2020 WFA. An earlier version of this paper was titled "Very Noise Option Prices and Inference Regarding Option Returns". Duarte is with the Jesse H. Jones School of Business at Rice University (jefferson.duarte@rice.edu). Jones is with the Marshall School of Business at the University of Southern California (christopher.jones@marshall.usc.edu). Wang is with the Ourso College of Business at Louisiana State University (junbowang@lsu.edu).

Determining whether volatility is priced has important consequences for option pricing, interpreting implied volatilities, and understanding the purpose of equity option markets. Stochastic volatility models are the norm in the option pricing literature, and numerous papers have demonstrated that a nonzero price of volatility risk can improve model fit substantially. A volatility risk premium also means that implied volatilities, whether they are from the Black and Scholes (1973) model or the model-free approach of Britten-Jones and Neuberger (2000), cannot be interpreted as unbiased forecasts of future realized volatility. Finally, knowing whether a volatility risk premium exists has important implications for understanding the purpose of options markets, namely the extent to which these can be seen as markets for volatility risk.

Though understanding whether volatility is priced is a central issue in the option literature, the evidence is somewhat mixed. While both Coval and Shumway (2001) and Bakshi and Kapadia (2003a) find evidence consistent with volatility being negatively priced in index equity options, such evidence for individual stock options is weak or non-existent (e.g. Bakshi and Kapadia (2003b) and Driessen, Maenhout, and Vilkov (2009)). This latter finding is puzzling because the volatilities of individual stocks are positively correlated with index volatility. Therefore, standard economic theory implies that individual equity options are, to some extent, substitutes for options written on the index.

We address this puzzle in this paper. Specifically, we explore two implications of a no-arbitrage stochastic volatility model to examine the volatility risk premium implied in equity and index options. First, a nonmonotonic relation between expected returns of calls and their moneyness is possible in a stochastic volatility model with a negative volatility risk premium. This differs from the Black and Scholes model, in which the expected returns of call options on stocks are positive and increasing with moneyness. In contrast, in a stochastic volatility model, option expected returns are driven by both the expected return of the underlying asset and the volatility risk premium. The balance between these two effects may result in a nonmonotonic relation between expected call returns and their moneyness or even in negative expected returns for deep out-of-the-money (OTM) call options. We examine this implication with portfolios of call options sorted by their moneyness. Second, a negative volatility risk premium implies that the expected return of delta-hedged options decreases with the sensitivity to underlying volatility (β_σ). We examine this implication using Fama and MacBeth (1973) (FM) regressions of delta-hedged option returns on β_σ.

Most importantly, we examine these standard implications of stochastic volatility models while systematically accounting for the biases that measurement errors in prices can cause in the estimation of volatility risk premia. Asset pricing researchers have long noted that measurement errors in prices can result in biased estimates of expected returns and risk premia (Blume and Stambaugh (1983), Fama (1984), Stambaugh (1988), Asparouhova, Bessembinder, and Kalcheva (2010, 2013)). These biases lead to practices that are now commonplace in the estimation of expected returns and risk premia in equity markets. For instance, it is now common in the asset pricing literature to use value weighting in addition to equal weighting because value-weighted mean returns result in unbiased estimates of expected returns (see Blume and Stambaugh (1983)). The literature on options, on the other hand, has not adopted these common practices. Instead, researchers typically attempt to address these potential biases by discarding a large portion of the data (e.g., call options with $\Delta < 0.15$) or focusing on only a small subset of data, such as at-the-money (ATM) options. Discarding a large part of the data for empirical studies is problematic, not only because it violates one of statistics' first principles — "Thou shall not throw data away" (Zhang, Mykland, and Aït-Sahalia (2005)) — but also because leads to conclusions that may be applicable only to a subset of the data.

Instead of discarding data, we adapt the bias-correction practices used in the equity literature to options. Our prescription is based on practices in Blume and Stambaugh (1983), Fama (1984), Stambaugh (1988), and Asparouhova, Bessembinder, and Kalcheva (2010, 2013). The bias-adjustment prescription can be easily implemented and comprises three simple procedures: First, estimate expected returns of options and delta-hedged options using weighted averages, where weights are proportional to one-day-lagged gross option returns. For instance, estimate the expected return of a portfolio with N calls as $\sum_{i=1}^{N} w_{C,i,t-1} R_{C,i,t} / \sum_{j=1}^{N} w_{C,i,t-1}$, where $R_{C,i,t}$ is the return of the i^{th} call at time t and $w_{i,t-1}$ is the gross return of call i at time $t-1$ ($C_{i,t-1}/C_{i,t-2}$). Second, estimate FM regressions with weighted least squares (WLS) instead of ordinary least squares (OLS) using one-day-lagged gross option returns as weights. Third, skip a day between signals and returns. That is, use doubly lagged variables (e.g., $\beta_{\sigma,t-2}$ or moneyness) in portfolio sorts, FM regressions, and sample selection.[1]

We find strong evidence that volatility is priced at the individual equity option level. Specifically, we strongly reject the assumption that call option returns monotonically increase with option moneyness. In fact, contrary to the assumption that call option returns are driven only by the expected return of the underlying asset, we find that the bias-adjusted mean return of heavily traded deep OTM call options is negative, at around -73 bps per day.[2] Moreover, our FM results indicate that the volatility risk premium in individual equity options is remarkably consistent with that in S&P 500 Index call options. Indeed, we find a volatility risk premium of about -5 bps per day for high-open-interest call and put options written on the stocks in the S&P 500 and for call options written on the S&P 500 Index.[3] Importantly, we find qualitatively different results for options on individual equity if we do not correct for microstructure biases. Specifically, without bias corrections, the FM regressions indicate that volatility is not negatively priced in individual equity options.

We use simulations to gauge the size of microstructure biases in the estimation of the volatility risk premium and to examine our proposed bias-adjustment method. Our simulations indicate that the biases in FM regressions estimating the volatility risk premium can completely change the conclusion regarding volatility being priced in individual equity options. More importantly, the simulation results indicate that the bias-adjusted method substantially reduces or completely eliminates biases due to measurement errors in prices. Specifically, an FM regression without (with) bias adjustment estimated with simulated data from a model with a volatility risk premium of -5.5 bps per day results in a estimated volatility risk premium of zero (-5.4 bps per day).

[1]The simplicity of the prescription above does not come without a cost. Errors in stock and option prices not only lead to direct biases in returns, but also to indirect biases because hedge ratios (deltas) are calculated with stock and option prices, which are plagued with measurement errors. Although this indirect bias is small in the case of delta-hedged option mean returns, it is relatively large in some other cases such as straddles and cannot be avoided by simple applications of the methods described in the equity literature. In the appendix, we describe procedures that are appropriate for straddles, β_S-adjusted returns (Constantinides, Jackwerth, and Savov (2013)), and leverage-adjusted returns (Fournier, Jacobs, and Orlowski (2021)).

[2]Individual equity option trading is highly concentrated, with 20% accounting for about 90% of the dollar open interest. Because more actively traded options may have return patterns different from those of thinly traded options, we report results both unweighted and weighted by dollar open interest.

[3]The volatility risk premium is about -11 basis points per day on put options written on the index.

We contribute to the literature by showing that volatility is, in fact, priced at the individual stock level once we address the biases in the estimation of the volatility risk premium. The negative average returns of deep OTM call options we find are consistent with those found by Ni (2008) using monthly call option returns. We extend Ni (2008) by attributing these negative average returns to a negative volatility risk premium. Driessen, Maenhout, and Vilkov (2009) is the most comprehensive study of the pricing of volatility in individual equity options.[4] They examine the difference between the realized variance (RV) and the model-free implied variance ($MFIV$) derived from options prices. A negative volatility risk premium implies that the RV of the underlying asset is smaller than the $MFIV$ on average. Using options on stocks in the S&P 100 Index, they show that the average across stocks of the difference between the RV and the $MFIV$ is not different from zero, indicating that volatility is not priced in individual equity options. We show that the average difference between the $MFIV$ and the RV is severely affected by the choice of Driessen, Maenhout, and Vilkov (2009) to only focus on options that are close to ATM. The most prominent empirical characteristic of equity options is the implied volatility smile. That is, the implied volatility of OTM options is larger than the implied volatility of ATM options. Hence, excluding OTM options from the $MFIV$ calculation can result in low estimates of the $MFIV$, which in turn results in unreliable estimates of the difference between $MFIV$ and RV.[5]

We also extend the literature that analyzes microstructure biases in the estimation of expected returns. Specifically, our study is the first to document that these biases are large and important for option portfolios. Moreover, we show that the usual procedure in the options literature of deleting a large part of the options data may not address the microstructure biases and instead may result in estimates that are applicable to only a subset of the data. Our approach builds on Blume and Stambaugh (1983) and Asparouhova,

[4] Other studies examining this issue do not use a sample as large as the one in Driessen, Maenhout, and Vilkov (2009). For instance, Bakshi and Kapadia (2003b) examine a sample of options written on 25 stocks.

[5] Carr and Wu (2009) estimate the difference between $MFIV$ and RV for a sample of 35 stocks. Differently from Driessen, Maenhout, and Vilkov (2009), they do not restrict their sample to close-to-ATM options, and they do not compute the average difference between the $MFIV$ and the RV across all stocks, focusing instead on stock-by-stock differences. Carr and Wu (2009) point out that the variance risk premium is very noisy at the individual stock level, resulting in only seven stocks in their sample with a statistically significant negative variance risk premium. In contrast, the standard deviations on the log variance risk premium are much smaller, which leads to a statistically significant negative log-variance risk premium for 23 out of the 35 stocks in their sample.

Bessembinder, and Kalcheva (2010, 2013), who study the effect of microstructure noise in stock prices. Dennis and Mayhew (2009) examine the effect of noisy option prices on the estimation of option pricing models, while Hentschel (2003) studies the effect of noise on the estimation of the implied volatility smile. Our study differs from these papers because we focus on the effects of noisy prices on inferences about the expected returns and risk premia of options. Moreover, two recent studies explore biases that are related to differences between the closing and intraday option prices. Eraker and Osterrieder (2018) show that the VIX is a biased estimate of implied volatility because the midpoint of the closing bid-ask prices is biased due to order flow pressure, while Goyenko and Zhang (2019) document biases on mean returns computed with closing prices due to differences between the closing and intraday option prices. We differ from these studies because the biases we document are not at all related to the differences between intraday and closing prices. In fact, the biases we explore depend only on the assumption that the prices (intraday or closing) used to compute returns are observed with zero-mean measurement errors.

Our results indicate that the option literature needs to adopt procedures to deal with biases due to measurement errors in prices in the same way that the empirical equity literature has dealt with these biases at least since Blume and Stambaugh (1983) and Fama (1984). Naturally, our bias estimates are based on simulations that rely on a series of assumptions. For instance, we make assumptions about the distribution of the measurement errors. As a result, our estimates of the size of the biases are solely indicative, and they are not intended to be exact estimates. Despite this, our simulation results, together with the simple observation that option relative bid-ask spreads are enormous (on average about 12%), strongly indicate that estimation biases due to measurement errors in prices are economically significant when working with options. Our paper therefore contributes to the empirical literature on options by proposing a method to deal with these biases in the estimation of expected option returns and risk premia.

The remainder of the paper is as follows. Section 1 describes our data. Section 2 describes our methodology, and Section 3 presents the empirical results. In Section 4, we use Monte Carlo simulations to show that microstructure biases affect the estimation of volatility risk premia and that these biases are addressed with our simple methodology. Section 5 concludes.

1 Data and Data Filters

We focus on options written on the S&P 500 Index and on stocks that are in the S&P 500. Our sample period is from January 1996 to June 2019. Our primary data come from the IvyDB dataset of OptionMetrics, which contains closing bid-ask quotes on options, closing stock prices, stock returns, implied volatilities, option deltas (Δ), vegas (ν), and open interest.[6] The Δ and ν of a given option in IvyDB are calculated with that option's own implied volatility, which depends on the observed price of the option and on the underlying asset. We compute option returns based on the midpoints of the option bid-ask quotes. Excess returns are equal to returns minus the risk-free rate.[7] As in Driessen, Maenhout, and Vilkov (2009), we examine options with maturities from 14 to 60 calendar days (or 10 to 44 trading days).[8] In addition to OptionMetrics, we use the volume-weighted effective bid-ask spreads of stocks from WRDS Intraday Indicators. We calibrate the bid-ask spreads of underlying stocks in our simulations to those observed in the data.

As is standard in the literature, we remove option prices that appear to be data errors. Removing these prices has some potential to induce look-ahead bias. To understand why, note that a trader cannot impose a filter based on the option price at time t since a trader does not know at time of portfolio formation $(t-1)$ the option price at time t. Because removing these observations may induce look-ahead bias, we remove them only as a last resort. Deleting these prices is necessary, however, because they are clearly erroneous and potentially very influential. For instance, we eliminate an observation with a call ask price of $9,999 on a stock with a price of $40. These clearly erroneous observations are not used in any of our empirical analyses, and they are only about 0.3% of the observations, a small percentage of the sample. The appendix contains a detailed description of this procedure.

We then impose *baseline* filters on the sample. It is common in the literature to impose fil-

[6]Implied volatilities, deltas, and vegas are computed using the binomial tree approach of Cox and Rubinstein (1979), which accounts for dividends and for the potential early exercise of the American options on individual equities. They are calculated with the Black and Scholes model in the case of the European options on the S&P 500 Index.

[7]The risk-free rate is based on the shortest maturity yield in the IvyDB zero-coupon term structure file, and accounts for the number of calendar days in the return holding period. As discussed by French (1983), interest accrues on a calendar-time basis.

[8]Our results regarding volatility pricing are robust to extending maturities to 180 calendar days.

ters based on bid-ask spread percentages, non-missing implied volatilities, option midpoints, option open interest, and deltas.[9] We impose these filters in our baseline sample to analyze whether they may, in fact, create biases in the estimation of the volatility risk premium. The standard practice in the literature, which we follow in our baseline sample, is to impose filters like these at the start of the holding period ($t-1$). Doing so prevents look-ahead biases since a trader can impose these filters at the beginning of the holding period ($t-1$). This standard practice, however, may result in sample selection biases when prices are observed with error. We discuss these biases in detail in Section 2.

Table 1 presents the following summary statistics of our baseline sample: the mean, median, standard deviation, first and third quartiles (Q1 and Q3), first and 99^{th} percentiles (P1 and P99), and skewness of option returns, stock returns, spreads, moneyness, implied volatilities, and dollar open interest ($\$OI$). We define the moneyness of an option as $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where K is the option strike price, σ_t is the implied volatility of the option at time t, $T-t$ is time-to-expiration in years, r_t is the risk-free rate, and S_t is the observed price of the stock at time t. Option spreads are expressed as percentages relative to quote midpoints. The $\$OI$ of an option is its open interest multiplied by the midpoint of its closing quote.

Table 1 indicates that the distribution of $\$OI$ is highly skewed. Indeed, Figure 1 shows that the options in the top quintile of the $\$OI$ distribution account for around 90% of the total $\$OI$. This extreme skewness in $\$OI$ suggests that empirical results can be driven by a majority of options that are very thinly traded. To address this possibility, we also present results weighted by $\$OI$.

Table 1 also shows differences between options written on individual stocks and on the index. For call options written on individual stocks (Panel A), the median relative bid-ask spread is about 8%, and for put options it is approximately 9% (Panel B). Median moneyness

[9] Specifically, as in Goyal and Saretto (2009), Cao and Wei (2010), Muravyev (2016), and Christoffersen, Goyenko, Jacobs, and Karoui (2018), we require that the midpoint price must be at least \$0.10 and that the bid-ask spread percentage not exceed 50% of the midpoint (which rules out zero-bid prices). We also eliminate options with zero-open interest. We require the absolute value of delta to be between 0.01 and 0.99. Our constraints on delta are less severe than those of Driessen, Maenhout, and Vilkov (2009), who examine puts (calls) with delta smaller (larger) than -0.05 (0.15). We discuss the consequences of this difference in the constraints related to the deltas in Section 3. We also require that the option have a computed implied volatility by IvyDB.

is negative for call options, indicating more in-the-money call options in the sample, and it is near zero for put options. The median implied volatility for options written on individual stocks is about 31%. The mean return of calls (puts) is about 53 (-62) bps per day. Perhaps the most interesting statistic about returns is not tabulated, which is that the average autocovariance in daily call (put) returns is -12 (-13) bps. Similar values for stock returns in our sample are economically negligible.[10] These autocovariances, after multiplication by -1, provide estimates of the variance in proportional measurement errors in prices (Roll (1984)) and the size of the MR bias (Blume and Stambaugh (1983)). The mean and the median bid-ask spreads for options written on the index (Panels C and D) are substantially smaller than those of options written on individual stocks. The median moneyness is negative for put options written on the index, indicating more OTM index puts in the sample. For call options written on the S&P 500 Index, the median relative bid-ask spread is about 3%, and for put options it is approximately 6%. The median implied volatility for call (put) options written on the index is about 14% (20%), which is substantially smaller than the median implied volatility of options on individual stocks. The mean return of calls (puts) is about 84 (-221) bps per day. For index options, average daily autocovariances are about -9 bps for both calls and puts. The fact that the absolute values of these autocovariances are smaller than those in equity options indicates that the variance of proportional measurement errors in index option prices is smaller than that in individual equity options, which is unsurprising given the lower spreads of index options.

Our goal is to estimate the volatility risk premium taking into account measurement errors in prices, not to estimate the transaction costs that investors face when trading options. As it is traditional in the asset pricing literature, we consider transaction costs to be a separate issue. To a large extent, this is because transaction costs are investor-specific, depending on the investment strategy and the quality of investor's trade execution (e.g. Muravyev and Pearson (2020)). Nevertheless, Figure 2 plots the time series of the average closing bid-ask spread for calls and puts written on individual stocks and on the S&P 500 Index. We show these spreads because it is reasonable to assume that they are related to the variance of

[10] To compute average autocovariance, we calculate the cross-sectional covariance of option returns with their own lag on each day of the sample and then average across days. Newey-West t-statistics are -11.0 for calls and -8.6 for puts.

the measurement errors that plague the estimation of risk premia.[11] To the extent that measurement error variance is related to the width of bid-ask spreads, the figure shows that the biases we identify in this paper may still be strong even at the end of our sample.

Figure 3 displays average loadings (βs) as a function of moneyness in our sample. The sensitivity of a call option return with respect to the return of the underlying asset (β_S^C) is $\Delta_t^C \times S_t/C_t$, where Δ_t^C and C_t are the delta of the call option and its price, respectively, both from IvyDB. The sensitivity of a call option return with respect to the volatility change in the underlying asset (β_σ^C) is defined as ν_t^C/C_t, where ν_t^C is the vega of the call option, also from IvyDB. The definitions of put option βs are similar. As we explain in Section 2, these option loadings drive the expected returns of options in a stochastic volatility model.

2 Empirical Methodology

The instantaneous expected excess return of a derivative in a standard no-arbitrage stochastic volatility diffusion model is expressed as

$$E_t\left[\frac{d\tilde{f}_t}{\tilde{f}_t}\right] - rdt = \tilde{\beta}^f_{S,t} E_t\left[\frac{d\tilde{S}_t}{\tilde{S}_t} + (q-r)dt\right] + \tilde{\beta}^f_{\sigma,t}\tilde{\lambda}_\sigma dt, \tag{1}$$

where $\tilde{f}_t$ is the price of the derivative (a call or a put), r is the instantaneous risk-free rate, q is the dividend yield, and dt is the instantaneous time change. The sensitivity of the derivative's return with respect to the return on the underlying asset is $\tilde{\beta}^f_S$, and the sensitivity of the derivative's return with respect to the volatility of the underlying asset is $\tilde{\beta}^f_\sigma$. $\tilde{\lambda}_\sigma$ is the volatility risk premium.[12]

We start the empirical examination of a negative price of volatility risk in options by analyzing whether the mean return of unhedged calls increases monotonically with moneyness. Specifically, Equation 1 reveals that the expected returns of call options do not necessarily monotonically increase with strike prices or moneyness when $\tilde{\lambda}_\sigma < 0$. To see this, note that

[11] Consistent with the findings of Battalio, Hatch, and Jennings (2004), the top panels show a substantial decrease in the closing bid-ask spreads of options on stocks in late 1999. However, closing bid-ask spreads increased substantially in the 2013–2014 period and by the end of the sample period are close to their values at the beginning of the sample.

[12] Specifically, $\tilde{\lambda}_\sigma$ is the difference between the drifts of the volatility diffusion under the physical and risk-neutral probability measures. The appendix contains a proof of this equation, following the results in Duarte and Jones (2010).

both $\tilde{\beta}_S^C$ and $\tilde{\beta}_\sigma^C$ are positive, increasing, and nonlinear functions of moneyness (see Figure 3). Moreover, the $\tilde{\beta}_\sigma$ of an in-the-money (ITM) call option is close to zero, while the $\tilde{\beta}_\sigma$ of an OTM call option is of the same order of magnitude as its $\tilde{\beta}_S$. That is, the expected return of ITM calls is largely driven by the expected return of the underlying asset, which is positive, while the expected return of OTM calls is driven by both the positive expected return of the underlying asset and the negative volatility risk premium when $\tilde{\lambda}_\sigma < 0$. As a result, $\tilde{\beta}_{\sigma,t}^f \tilde{\lambda}_\sigma$ in Equation 1 is negative and may be the dominant term driving the expected returns of OTM call options. In fact, the expected returns of deep OTM call options may be negative when $\tilde{\lambda}_\sigma < 0$. On the other hand, when $\tilde{\lambda}_\sigma = 0$, the expected return of a call option is driven only by its exposure to the underlying asset, and it increases with moneyness as $\tilde{\beta}_S^C$ does. We analyze the relation between call option expected returns and moneyness with portfolios of call options sorted by moneyness.

Equation 1 also indicates that FM regressions of delta-hedged excess returns on an option's $\tilde{\beta}_\sigma$ can be used to estimate $\tilde{\lambda}_\sigma$. Indeed, Equation 1 implies that expected delta-hedged excess returns are $E_{t-1}[\tilde{R}_{f,t}] - \tilde{\beta}_{S,t-1}^f E_{t-1}[\tilde{R}_{S,t}] \approx \tilde{\beta}_{\sigma,t-1}^f \tilde{\lambda}_\sigma$. Hence, we can estimate $\tilde{\lambda}_\sigma$ with FM regressions of delta-hedged option returns on option volatility betas. The linear relation in Equation 1 holds in continuous time, while the relation between options and underlying excess returns is nonlinear for discrete holding periods. To alleviate possible problems due to this nonlinearity, we estimate the FM regressions with high-frequency data (using daily instead of monthly returns).

Naturally, the observed derivative and underlying asset prices, as well as the returns and βs, contain measurement errors. Following Blume and Stambaugh (1983), we assume that the observed price of a derivative is $f_t = \tilde{f}_t(1+\epsilon_{f,t})$, where $\tilde{f}_t$ represents the true price of the derivative, and $\epsilon_{f,t}$ is the zero-mean derivative price error term. The measurement errors in prices affect the observed returns. To see this, note that a Taylor expansion implies that

$$\left(\frac{f_t}{f_{t-1}} - 1\right) \approx \left(\frac{\tilde{f}_t}{\tilde{f}_{t-1}} - 1\right) + \frac{\tilde{f}_t}{\tilde{f}_{t-1}} \times \epsilon_{f,t} - \frac{\tilde{f}_t}{\tilde{f}_{t-1}} \times \epsilon_{f,t-1} - \frac{\tilde{f}_t}{\tilde{f}_{t-1}} \times \epsilon_{f,t} \times \epsilon_{f,t-1} + \frac{\tilde{f}_t}{\tilde{f}_{t-1}} \times \epsilon_{f,t-1}^2. \quad (2)$$

Analogously, measurement errors in the prices of the derivative and of the underlying asset also infect the observed option βs because the βs of an option are calculated with that option's own implied volatility, which depends on the observed price of the option and of the underlying asset.

Because we do not observe true prices, our call portfolio sorts and FM regressions rely on observed returns and βs, which have measurement errors. We use an upper tilde ($\tilde{}$) for the variables without measurement errors, while the variables infected by measurement errors in prices do not have an upper tilde. For example, the observed one-day excess return and sensitivities of a derivative are represented as $R_{f,t}$, β_S^f, and β_σ^f, while the unobserved variables without measurement errors are represented by $\tilde{R}_{f,t}$, $\tilde{\beta}_S^f$ and $\tilde{\beta}_\sigma^f$ respectively.[13] Hence, the baseline FM regression suggested by Equation 1 is

$$R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t} = \lambda_0 + \lambda_\sigma \times \beta^f_{\sigma,i,t-1} + \eta_{i,t}. \tag{3}$$

Measurement errors in prices lead to four types of bias in the estimated average return of sorted option portfolios and in the volatility risk premia estimated with the baseline FM regression in Equation 3:

First is the mean return (MR) bias, which affects portfolio sorts. To understand the MR bias, note that Equation 2 leads to $E[R_{f,t} - \tilde{R}_{f,t}] \approx E[\tilde{f}_t/\tilde{f}_{t-1}]E[\epsilon^2_{f,t-1}]$.[14] The MR bias was first described by Blume and Stambaugh (1983) in the context of stocks. The variance of the measurement error in option prices is likely to be much larger than that of stock prices because the relative bid-ask spreads of options are much larger than those of stocks. Hence, the MR bias in mean option returns is likely to be much larger than that in stock returns.

Second is the regression coefficient (RC) bias, which affects our FM regressions. To understand this bias, assume a simple model in which the true (unobserved) expected return of a derivative satisfies the univariate relation, $E_{t-1}[\tilde{R}_{f,i,t}] = \tilde{\lambda}\tilde{\beta}^f_{i,t-1}$. For simplicity, assume that we observe $\tilde{\beta}^f_{i,t-1}$ and we estimate the regression $E_{t-1}[R_{f,i,t}] = \lambda_0 + \lambda\tilde{\beta}^f_{i,t}$. Under some simplifying assumptions, the OLS estimate of λ is approximately $\tilde{\lambda}(1 + E[\epsilon^2_{f,t-1}])$.[15] Intu-

[13] For a call option, its observed price is represented as $C_t = \tilde{C}_t(1 + \epsilon_{C,t})$. Put prices ($P_t$) and stock prices ($S_t$) also contain measurement errors ($\epsilon_{P,t}$ and $\epsilon_{S,t}$). The observed call, put, and underlying asset excess returns are represented by R_C, R_P, and R_S, respectively. In our empirical analysis, the observed call and put prices are the closing midpoint of the bid and ask prices. Our analysis, however, is general, and the observed price does not need to be the closing midpoint of the bid and ask prices.

[14] To arrive at this expression we use the standard assumption in this literature that measurement errors have zero mean, are uncorrelated with the actual return ($\tilde{f}_t/\tilde{f}_{t-1}$), and are IID. Note we focus on a conditional option pricing model in our empirical work. However, we follow the literature and define the bias as the unconditional expectation of the difference between observed and actual returns ($E[R_{f,t} - \tilde{R}_{f,t}]$).

[15] We obtain this expression by substituting the expression $E_{t-1}[\tilde{R}_{f,i,t}] = \tilde{\lambda}\tilde{\beta}^f_{i,t-1}$ in Equation 2 and using

itively, the RC bias is a generalization of the MR bias to an FM regression setting. To see this, note that the MR bias, $E[\tilde{f}_t/\tilde{f}_{t-1}]E[\epsilon^2_{f,t-1}]$, increases with the expected return of the derivative. Hence, the MR bias changes with $\tilde{\beta}^f_{i,t-1}$ in our simple model for expected return and λ is a biased estimate of $\tilde{\lambda}$.

The third is the sample selection (SS) bias, which potentially affects both portfolio sorts and FM regressions. To understand this bias, assume that our intent is to estimate the population expected return of a certain type of derivative (e.g., the unconditional expected return of call options, $E[R_{C,i,t}]$) with a sample that is built with some data filters (e.g., price at time $t-1$ is above \$0.10 or implied volatility is not missing at time $t-1$). Mathematically, the expected return of the derivatives in our sample is $E[R_{f,i,t}|\mathbb{1}_{Sample}(f_{i,t-1}) = 1]$, where $\mathbb{1}_{Sample}(f_{i,t-1})$ is a function indicating whether the derivative i is within our sample at time $t-1$. The measurement errors in $f_{i,t-1}$ affect both $R_{f,i,t}$ and the indicator function $\mathbb{1}_{Sample}(f_{i,t-1})$, inducing a spurious covariance between $R_{f,i,t}$ and $\mathbb{1}_{Sample}(f_{i,t-1})$ that, in turn, results in a bias. Specifically, $E[R_{f,i,t}|\mathbb{1}_{Sample}(f_{i,t-1})] - E[R_{f,i,t}] = Cov[R_{f,i,t}, \mathbb{1}_{Sample}(f_{i,t-1})]/E[\mathbb{1}_{Sample}(f_{i,t-1})]$.

An example of how the SS bias arises in our context is the common practice of discarding options without implied volatility at time $t-1$. To see how this bias works in this case, note that negative values of $\epsilon_{f,i,t-1}$ lead to low values of observed option prices ($f_{i,t-1}$) that tend to be reversed, resulting in high values of $R_{f,i,t}$. These negative measurement errors may also result in observed midpoint option prices ($f_{i,t-1}$) that are below lower-arbitrage bounds and hence do not have a computed implied volatility. As a result, deleting these observations from the sample results in $Cov[R_{f,i,t}, \mathbb{1}_{Sample,t-1}] < 0$, leading to a negative bias in the estimated expected return.

Finally, the fourth bias is the correlated error in variables (CEIV) bias, which also potentially affects both portfolio sorts and FM regressions. In the FM regression context, the CEIV bias renders the estimated λ_σ in Equation 3 spurious because $\beta^f_{\sigma,i,t-1}$ and returns ($R_{f,i,t}$ and $R_{S,t}$) are infected by the same measurement errors in option and underlying prices. In portfolio sorts, the estimated expected return is $E[R_{f,i,t}|\mathbb{1}_{Portfolio}(f_{i,t-1})]$, where $\mathbb{1}_{Portfolio}(f_{i,t-1})$

the assumptions that the zero-mean measurement errors are independent and identically distributed across securities. Asparouhova, Bessembinder, and Kalcheva (2010) describe the RC bias in more detail.

is a function indicating whether the derivative is within the portfolio of derivatives satisfying the target characteristic in time $t-1$. The measurement errors in $f_{i,t-1}$ affect both $R_{f,i,t}$ and the indicator function $\mathbb{1}_{Portfolio}(f_{i,t-1})$, inducing a spurious covariance between $R_{f,i,t}$ and $\mathbb{1}_{Portfolio}(f_{i,t-1})$ that, in turn, creates a bias.

We address these biases with a bias-adjustment procedure based on the best practices from the asset pricing literature. Our simulations in Section 4 show that the bias-adjustment procedure described below addresses the MR, CEIV, RC, and SS biases in portfolio sorts and FM regressions.

We dispense with SS bias by adjusting our sample selection to break its dependency on the option prices used to estimate returns. Specifically, we use an *adjusted* option sample that is built with the same type of filters as the baseline filters described in Section 1. However, these filters are imposed on $t-2$ prices in our adjusted sample instead of on $t-1$ prices as in our baseline sample. Moreover, when the option's delta at time $t-1$ is missing, which occurs when implied volatility cannot be calculated, we use the option's delta at time $t-2$ to compute $\beta^{f}_{S,i,t-1}$ as $\Delta^{f}_{S,i,t-2} \times S_{t-1}/f_{t-1}$. This procedure addresses the deletion of observations without implied volatility at time $t-1$ from the sample. Also note that this procedure does not imply that we are using prices that allow for arbitrage. This is because OptionMetrics computes implied volatilities based on closing stock prices and option mid-quotes, which results in missing implied volatilities when these prices do not satisfy no-arbitrage bounds. However, these no-arbitrage filters are often incorrect, since an actual arbitrage opportunity requires that the option's ask price, rather than the mid-quote, be below the no-arbitrage lower bound. Furthermore, the no-arbitrage bound depends on the stock's end-of-day bid or ask price instead of the stock closing price, which is almost always outside the bid-ask spread (Bogousslavsky and Muravyev (2021)).

Analogously, we address the CEIV bias in our portfolio sorts by adjusting our portfolio-creating criteria to break the dependency on option prices used to estimate returns. We sort unhedged call options by moneyness at time $t-2$, instead of using moneyness at time $t-1$. In the FM regressions, we avoid the CEIV bias by using independent variables that are not calculated with the same prices as the dependent variables (e.g., Fama (1984) and Stambaugh (1988)). Specifically, the independent variable in the bias-adjusted FM regression is lagged by one additional day:

$$R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t} = \lambda_0 + \lambda_\sigma \times \beta^f_{\sigma,i,t-2} + \eta_{i,t}. \tag{4}$$

We address the MR bias by estimating expected returns with the weighted average — $\sum_{i=1}^{N} w_{f,i,t-1} R_{f,i,t} / \sum_{j=1}^{N} w_{f,i,t-1}$ — where $R_{f,i,t}$ is the return or delta-hedged return of the i^{th} derivative at time t and $w_{i,t-1}$ is the gross return of the derivative i at time $t-1$ ($f_{i,t-1}/f_{i,t-2}$). Analogously, we address the RC bias with FM regressions estimated with WLS using the gross return of the derivative i at time $t-1$ ($f_{i,t-1}/f_{i,t-2}$) as the weight.[16] Asparouhova, Bessembinder, and Kalcheva (2010, 2013) show that the MR and RC biases are reduced by weighting observations proportionally by their lagged gross returns. Intuitively, the derivative prices with large $\epsilon_{f,i,t-1}$ have larger weights ($f_{i,t-1}/f_{i,t-2}$) and smaller returns ($R_{f,i,t}$). This negative covariance between weights and returns approximately offsets the MR and RC biases, resulting in estimates that are near to unbiased. To wit, in the case that observed prices were martingales, gross-weighted average returns would be similar to equal-weighted average returns, making the Asparouhova, Bessembinder, and Kalcheva (2010) adjustment unnecessary.[17]

3 Empirical Results

We present empirical results based on our bias-adjustment procedure described above. We also show bias-unadjusted results, which are based on equal-weighted averages of returns of options portfolios, and on OLS estimation of Regression 3. The bias-unadjusted results are also based on our baseline sample described in Section 1. We present results of both adjusted and unadjusted estimations to show the extent to which the MR, CEIV, RC, and

[16] Formally, assume we have a sample with N derivatives and define the 2×1 vector $\boldsymbol{\beta}_{i,t-2} = [1, \beta^f_{\sigma,i,t-2}]'$. Define $\boldsymbol{\beta}_{t-2} = [\boldsymbol{\beta_{1,t-2}}, \cdots, \boldsymbol{\beta_{N,t-2}}]$ as a $2 \times N$ matrix. Similarly, define the delta-hedged return $R_{i,t} = R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$ and $\boldsymbol{R}_t = [R_{1,t}, \cdots, R_{N,t}]'$, which is $N \times 1$ matrix. Following Asparouhova, Bessembinder, and Kalcheva (2010), we estimate Regression 4 with the WLS estimator $(\boldsymbol{\beta}_{t-2} \boldsymbol{W}_{t-1} \boldsymbol{\beta}'_{t-2})^{-1} (\boldsymbol{\beta}_t \boldsymbol{W}_{t-1} \boldsymbol{R}_t)$, where $\boldsymbol{W}_{t-1}$ is a diagonal matrix of weights ($w_{i,t-1}$) given by the gross return of the derivative i at time $t-1$ ($f_{i,t-1}/f_{i,t-2}$). That is, the derivatives with large measurement errors, $\epsilon_{f,t-1}$, have larger weights ($f_{i,t-1}/f_{i,t-2}$) and smaller returns ($R_{i,t}$). We estimate our FM regressions with individual options. Gross return weighting also addresses the RC bias when FM regressions are estimated with portfolios of options. In this case, $R_{i,t}$ is the return of the i^{th} portfolio and its weight is $w_{i,t-1} = R_{i,t-1} + 1$.

[17] Our method to address the MR, RC, CEIV and the SS biases relies on the assumption that price errors are independent in time-series and independent among stocks, calls and puts. In addition, the price errors are independent of true prices.

SS biases affect inferences about the volatility risk premium in options written on individual stocks and on the S&P 500 Index. In addition to showing adjusted and unadjusted results, we also report results weighted by $\$OI$ to address the possibility that our results are driven by a large number of thinly traded options.[18]

We start our empirical work by testing whether the expected return of call options increases monotonically with moneyness. For this purpose, Panel A of Table 2 reports statistics of portfolios of call options written on stocks in the S&P 500 Index sorted by moneyness. The results in this panel reveal that the median β_S of deep in-the-money (ITM) call options is 4.97, which is about 33 times their median β_σ (0.15). On the other hand, the median β_S of deep out-of-the-money (OTM) call options is 22.01, which is only about 1.5 times their median β_σ (13.17).[19] Therefore, these factor sensitivities, along with Equation 1, indicate that the expected returns of deep ITM calls are mostly driven by the expected returns on the underlying asset, while the expected returns of deep OTM calls are driven by both the expected returns on the underlying asset and the volatility risk premium. As a result, if λ_σ is negative, the expected returns of calls can vary nonmonotonically with their moneyness.

Consistent with volatility being negatively priced in individual stock options, the results in Panel A of Table 2 reveal that the average return of calls does not increase monotonically with moneyness. The right column of the table displays the p-value of a Wolak (1989) test. The null hypothesis in this test is that average return of calls increases monotonically with moneyness. The p-values in Panel A indicate that, independently of the bias-adjustment, this null hypothesis is rejected at the usual significance levels. Moreover, the $\$OI$-weighted results are somewhat stronger than the unweighted returns. Indeed, the bias-adjusted $\$OI$-weighted average return of OTM options is negative, at about -73 bps per day. A stochastic volatility model in which volatility is not priced cannot explain this negative expected return. As in the case of individual stock calls, the results in Panel B reveal that the average returns

[18]Specifically, we weight our means by $\$OI_{i,t-2}$ which is the dollar open interest of option i observed at time $t-2$. We use $\$OI_{i,t-2}$ instead of $\$OI_{i,t-1}$ to separate the effect of weighting by $\$OI$ from the effect of weighting by lagged gross returns ($f_{i,t-1}/f_{i,t-2}$). This way, we assess the extent that our results are affected by bias adjustment (weighting by $f_{i,t-1}/f_{i,t-2}$) or by focusing on options that are more heavily traded (weighting by $\$OI_{i,t-2}$).

[19]It is important to note that the bias unadjusted method uses $t-1$ sorting variables while the bias adjusted method uses $t-2$ variables. The displayed median values of β_σ, β_S, relative spread, and moneyness are based on $t-2$ variables and they are qualitatively the same as the median values based on $t-1$ values.

of calls on the S&P 500 Index also vary nonmonotonically with moneyness.

The results in Table 2 give a first look at the biases in mean option returns. Under the assumption that our bias correction method is accurate, the differences between the unadjusted and adjusted estimates indicate that estimation biases are large and complex.[20] To see this, note that the biases for individual equity calls change sign, with the bias in ITM calls negative, at about -11 bps per day (1.39-12.96), positive for Portfolio 4, at 10 bps (83.50-73.64), and negative for deep-OTM calls, at about -53 bps (-14.70-38.77). In the appendix, we isolate the MR bias from the total biases suggested by Table 2. The results in the appendix show that the MR bias is likely very large for deep-OTM calls, at around 80 bps per day, but it is more than offset by negative CEIV and SS biases.[21] Similarly, the appendix shows that the MR bias for index options is also large, at around 66 bps for deep-OTM calls, but it is again offset by the CEIV and SS biases. Naturally, the fact that these biases compensate for each other may not occur in other instances. Therefore, the large economic magnitude of the these biases indicates that, even for index options, bias adjustment can result in economically significant differences.

We next examine the sign of the volatility risk premium (λ_σ) in FM Regressions 3 and 4. Table 3 shows the results of our FM regressions, which regress delta-hedged returns on β_σ. The coefficients indicate that volatility is negatively priced on options written on stocks after adjusting for biases.[22] Panel A shows the results of FM regressions estimated without bias adjustments. The λ_σs in Panel A are consistent with the stylized fact in the options literature that volatility is not negatively priced at the individual stock level. Specifically, for call options, λ_σ is positive (1.74), though not significant, when estimated with OLS, and it is negative (-3.25) when estimated with WLS using $\$OI$ as weights. For put options, the estimated λ_σs are positive and insignificant. In contrast, the coefficients estimated with bias adjustments show a different picture. Indeed, Panel B reports that λ_σs estimated with bias

[20]In Section 4, we show that our bias adjustment methodology addresses the bias well and we precisely decompose the total bias in its components (MR, SS, and CEIV) with simulations. Our simulations show that the MR, SS, and CEIV biases are often of different signs, which explains the complexity observed here.

[21]The positive MR bias is consistent with the theoretical result that the MR bias is approximately $E[\tilde{f}_t/\tilde{f}_{t-1}]E[\epsilon^2_{f,t-1}]$, and it supports our maintained hypothesis that quote midpoints contain measurement errors.

[22]The appendix shows the results for the first and second half of our sample are qualitatively similar to those in Tables 2 and 3.

adjustment are mostly around -5 bps and statistically significant.

The difference between the results with and without $\$OI$ weights indicates that the most heavily traded call options on stocks have a negative volatility risk premium while thinly traded call options do not. The size of the estimation biases does not vary much with $\$OI$. However, the volatility risk premium is apparently more significant in options with high $\$OI$. The bias-adjusted estimate of λ_σ is -0.32 and not significant for call options. In contrast, the bias-adjusted FM regression using call options and $\$OI$ weights results in an estimated λ_σ of -5.4. These findings highlight that $\$OI$ weights address systematic differences between highly and thinly traded options. These systematic differences are not the same as the microstructure biases that we address in our bias-adjusted procedure but are nevertheless important.

As in the case of options written on individual stocks, the results in Table 3 indicate that volatility is negatively priced on options written on the S&P 500 Index. The estimated price of volatility risk with put options written on the S&P 500 Index is between -11 and -15, which is economically much larger than that estimated with call options (between -5 and -7). The larger price of volatility risk implied by put options on the index may be related to the index OTM put puzzle (e.g., Constantinides, Jackwerth, and Savov (2013)). The volatility risk premium in individual equity options is about the same as it is for S&P 500 Index call options. Both the adjusted FM regressions with individual stock options and the ones with S&P 500 Index call options result in estimates of a volatility risk premium close to -5.

In contrast to options written on individual stocks, the bias adjustment methodology is not strikingly important for index options. For both the unadjusted and adjusted (Panels A and B of Table 3 respectively) results, the estimated λ_σs are between -5 (-11) and -7 (-15) for call (put) options in the S&P 500 Index.

To better understand the effect of bias adjustment on these FM regression results, Tables 4 and 5 show the mean returns of β_σ-sorted portfolios of delta-hedged options on stocks and on the S&P 500 Index, respectively. We find that the large biases in λ_σ estimated from individual stock options result from biases in both ITM and OTM options. The results in Table 4, both for calls (Panel A) and puts (Panel B), suggest that the bias on the equal-weighted return of ITM options (low β_σ) is negative. Under the assumption that our bias adjustment procedure corrects for all the estimation biases, we infer that this bias is close to -29 bps (-30.96+2.30) for calls, and it is approximately -31 bps (-34.51+3.94) for puts. In contrast, the bias on the equal-weighted return of OTM options (high β_σ) is positive. For calls, this bias is close to 11 bps (-0.10+10.65), and it is approximately 45 bps (-20.58+65.84) for puts. The combination of a negative bias for options with low β_σ and a positive bias for options with high β_σ results in the large positive bias on the coefficient on β_σ estimated with the FM regression on individual stock options.

Panels A and B of Table 5 display the results for delta-hedged call and put options on the S&P 500 Index. The results in Table 5 clarify why the bias adjustment in the FM regressions is much more important for options on stocks than for index options. Specifically, the results for deep OTM index options are somewhat similar to those for stocks.[23] In contrast, the results for deep ITM options in this table paint a picture different from that of individual stock options in Table 4. Specifically, the bias on the equal-weighted return of ITM index options (low β_σ) is positive, at about 14 bps (11.82+2.59) for calls and 7 bps (3.12+3.87) for puts. These positive estimates contrast with the negative estimates for options on individual stocks, which are about -30 bps. This difference stems from the fact that many ITM stock options have missing implied volatility. In fact, about 23% of the deep ITM options on stocks have missing implied volatility, while this is true for deep ITM index options for only 4% of the sample. As it is common practice in the literature, our unadjusted sample does not include ITM options with missing implied volatility at time $t-1$ since it is not immediately obvious how to compute a delta-hedged return for them. Not including these options results in an economically large negative SS bias. Such a large negative SS bias reveals an important issue with the standard practice of excluding from empirical work stock options for which the midpoint price is in violation of an arbitrage bound.

Lastly, the column "All" in Tables 4 and 5 displays the mean delta-hedged return of all puts and calls. Consistent with a negative volatility risk premium for equity and index options, the displayed mean returns are in general negative. In addition, the economic significance of the mean returns on index options is larger than of the mean returns on

[23] It is interesting that bias adjustment makes the estimate of expected return of OTM put options smaller, which deepens the OTM index put puzzle (e.g. Constantinides, Jackwerth, and Savov (2013)).

individual equity options. For instance, the bias-adjusted mean return of delta-hedged puts on the index (stocks) is about -99 bps (-18 bps) per day. Note however that the values of β_σ for options on stocks are generally smaller than those of the options on the index. In particular, Figure 3 shows that the β_σs of options written on the index are about double the β_σ of the options written on individual stocks. The difference in βs arises because the median implied volatility of the options on the index is about half of the implied volatility of the options on individual stocks, and β_σ decreases with volatility. This difference in β_σs explains, to a large extent, why the mean delta-hedged returns of index options are economically more significant that those of the stocks options, while the estimated volatility risk premia (λ_σs) in Table 3 are somewhat similar.

Overall, the results in Tables 2 to 5 indicate that the volatility risk premium in individual stock options is not nearly as different from the volatility risk premium in S&P 500 Index options as previously suggested in the literature. To reconcile our results with those in the literature, we note that the methodologies we use in Tables 2 and 5 are different from those usually employed in the literature. In fact, the most comprehensive study of the volatility risk premium in individual stocks is Driessen, Maenhout, and Vilkov (2009), who use a methodology different from ours to examine the price of volatility risk. Specifically, they examine the volatility risk premium in stocks in the S&P 100 using the difference between the mean realized variance ($\overline{RV}$) and the mean model-free implied variance ($\overline{MFIV}$) of the stocks in their sample. The difference, $\overline{RV} - \overline{MFIV}$, is called the variance risk premium (VRP) and it is related to the volatility risk premium (λ_σ). Specifically, the instantaneous VRP is by definition equal to the difference between the drifts of the variance process under the physical and the risk neutral probability measures. We can show by Ito's lemma that the instantaneous VRP is $2\sigma_t\lambda_\sigma dt$.

Driessen, Maenhout, and Vilkov (2009) estimate RV from the daily returns over a one-month window and the $MFIV$ from one-month options. The $MFIV$ for a given stock is estimated by numerically solving the integral in the equation below:

$$MFIV_{i,t} = 2\int_0^\infty \frac{C_{i,t}(K) - max(S_{i,t} - K, 0)}{K^2}\, dK. \tag{5}$$

Equation 5 is the main result of Britten-Jones and Neuberger (2000) and states that we can derive the variance implied by prices of options (calls in the case above) written on stock i by integrating the ratio of their time value to the square of their strike prices. Naturally, some empirical choices are necessary to estimate the integral above. For instance, we do not observe a continuum of strike prices from zero to infinity; hence, an empirical methodology is necessary to deal with the fact that, for a given stock i on day t, we only observe data for a few strike prices between the interval $[K_{min,i,t}, K_{max,i,t}]$. Driessen, Maenhout, and Vilkov (2009) follow the methodology of Jiang and Tian (2005) by using both OTM puts and OTM calls instead of just calls as in Equation 5, interpolating implied volatilities between the observed strike prices, and assuming a constant volatility between zero and K_{min} as well as between K_{max} and infinity.

To reconcile our results with respect to the pricing of volatility with those in Driessen, Maenhout, and Vilkov (2009), we estimate the VRP for the stocks in our sample. Specifically, on each standard expiration Friday, we compute the $MFIV$ for stocks in our sample based on options expiring in the following month. We follow Driessen, Maenhout, and Vilkov (2009) and calculate $MFIV$ only for those stocks that have at least three options with positive open interest, where at least one has strike above the stock forward price and one has strike below the forward price. RV is the annualized realized variance computed from all daily returns up to the next expiration date. The results are in Table 6.

The first column of Table 6 displays the results when $MFIV$ is computed using the methodology of Driessen, Maenhout, and Vilkov (2009), which is based on a sample of OTM calls with $\Delta^C > 0.15$ and OTM puts with $\Delta^P < -0.05$. That is, $K_{min,i,t}$ is the smallest observed strike price for which $\Delta^P_{i,t} < -0.05$, and $K_{max,i,t}$ is the largest observed strike price for which $\Delta^C_{i,t} > 0.15$. As in Driessen, Maenhout, and Vilkov (2009), the results using this methodology suggest that volatility is not priced on options written on individual stocks. Indeed, $\overline{RV} - \overline{MFIV}$ is not statistically different from zero in the first column of Panel A.

The second column of Table 6 examines the average difference between RV and $MFIV$, estimating the $MFIV$ with a sample of primarily OTM calls with $\Delta^C > 0.01$ and OTM puts with $\Delta^P < -0.01$. In contrast to Driessen, Maenhout, and Vilkov (2009), the results using this methodology suggest that volatility is indeed priced on options written on individual stocks. The sample used to estimate results in Column (2) is the baseline sample described in Section 1, and it only includes options with open interest above zero. Hence, it does not rely on options that are never traded. It differs from the first approach in that it uses options that are deeper OTM. In addition, if an OTM call (put) is unavailable, it infers the implied volatility of that strike price from the corresponding ITM put (call) with the same strike. The purpose of both of these changes is to extend the range of strikes over which Equation 5 is calculated. In fact, the extended range of Δs in the estimation of $MFIV$ increases the average number of options per stock from 5.3 in Column (1) to 7.3 in Column (2). Moreover, the number of stocks for which we calculate $MFIV$ increases from about 254 per day to 352 per day in the extended sample, which occurs because an extended Δ range increases the number of stocks with at least three options with positive open interest on the $MFIV$ calculation day.[24] The results using the extended sample indicate that volatility is negatively priced in individual equity options. Indeed, $\overline{RV} - \overline{MFIV}$ is about -1.6% in the second column of Table 6, Panel A.

The third column in Panel A presents the results using the $MFIV$ data from Rehman and Vilkov (2012), which is based on the OptionMetrics Volatility Surface file.[25] Similar to our findings and in contrast to the findings of Driessen, Maenhout, and Vilkov (2009), the results in Column (3) show a negative VRP of about -1.4%.[26]

Interestingly, the negative VRPs in Table 6 are somewhat consistent with the negative values of λ_σ estimated in Table 3. Recall that the instantaneous VRP is $2\sigma_t\lambda_\sigma dt$. Replacing σ_t with 35% (the average volatility of stocks from Table 1), λ_σ with -5 bps per day from Panel B of Table 3, and dt with 30 days, we arrive at a monthly VRP of -1 percentage point, which is reasonably close to the VRPs in Columns (2) and (3) of Table 6, Panel A.

The difference between the calculation of $MFIV$ by Driessen, Maenhout, and Vilkov (2009) and by the methodology used in Column (2) of Table 6 is represented in Figure

[24] The differences between Columns (1) and (2) almost entirely arise as the result of including additional options for each firm and not from adding more firms to the sample. Applying the Column (2) approach to the sample of firms analyzed in Column (1) results in almost the same average $RV - MFIV$ reported in Column (2).

[25] We are grateful to Greg Vilkov for providing this data. Because the Rehman and Vilkov (2012) data are winsorized, we delete observations for which the $MFIV$ is equal to the maximum or minimum value in the cross-section.

[26] For completeness, Panel B of Table 6 presents the results for the S&P 500 Index options. In all three cases, the results indicate that volatility is priced in these options.

4. This figure displays the average of the implied volatilities of the OTM options used to numerically calculate the integral in Equation 5 as a function of the strike-to-spot ratio.[27] The continuous line in the figure shows that restricting the sample to OTM call (put) options with $\Delta^C > 0.15$ ($\Delta^P < -0.05$) results in the use of options with implied volatilities that are on average lower than those observed in the extended sample, which in turn leads to a smaller average $MFIV$ and a VRP close to zero. In contrast, our sample includes a larger range of strike prices, which results in larger $MFIV$ and a VRP consistent with the values of λ_σ estimated with the bias-adjusted FM regressions in Table 3.

4 Examining the Biases with Simulations

This section uses simulations to address two questions: First, of all the biases that we have described (MR, RC, SS, and CEIV), which matter most for each particular group of options and option strategy (e.g., hedged or unhedged options)? Second, the solutions we adopt to address the biases in the estimation of volatility risk premia are direct applications of the weighting scheme suggested by Asparouhova, Bessembinder, and Kalcheva (2010, 2013). It is well known that the weighting schemes proposed in these papers work for stocks. How well do these standard solutions work for options?

Theoretically, when we consider option strategies that are more complex than simple unhedged positions (e.g., straddles and delta-hedged options), the sources of MR and RC biases are not limited to those described in Asparouhova, Bessembinder, and Kalcheva (2010, 2013) and Blume and Stambaugh (1983); hence, it is not clear a priori that the weighting scheme proposed in these papers will work for FM regressions and portfolio sorts of delta-hedged options. Specifically, hedged option strategies have a type of MR bias that has not been previously described in the literature, because the measurement errors in prices affect the MR bias of hedged options both directly through prices and indirectly through the hedge

[27] For Figure 4 only, we restrict the sample to only include firms that have at least one call with delta below 0.01 and one put with delta above -0.01. This is done in order to be able to plot implied volatilities for deep OTM options.

ratio. For instance, the MR bias in a delta-hedged derivative is approximately[28]

$$E[\epsilon^2_{f,t-1}] - E[\tilde{\beta}^f_{S,t-1}]E[\epsilon^2_{S,t-1}] + E[\frac{\partial \tilde{\beta}^f_{S,t-1}}{\partial \tilde{S}}\tilde{S}_{t-1}]E[\epsilon^2_{S,t-1}]. \tag{6}$$

We refer to the first two terms in Equation (6) as the "direct" MR bias (DMR). These terms represent the bias that would arise if $\tilde{\beta}^f_{S,t-1}$ were observed at time $t-1$. To see this, note that these terms are the direct result of Equation 2, along with the assumption that measurement errors are IID mean-zero and that the definition of delta-hedged returns is $E_{t-1}[\tilde{R}_{f,t}] - \tilde{\beta}^f_{S,t-1}E_{t-1}[\tilde{R}_{S,t}]$. In contrast, we refer to the last term in Equation (6) as the "indirect" MR bias (IMR), which is not present in simple unhedged positions and has not been previously described in the literature. The IMR bias stems from the fact that $\beta^f_{S,t-1}$ is also affected by errors in underlying stock prices.[29] Similarly, we can decompose the RC bias into a direct RC bias (DRC) and an indirect RC bias (IRC). The RC bias is a generalization of the MR bias to a regression framework; hence, as we do for the MR bias, we can decompose the RC bias into two components: The DRC bias is due to the effect of measurement errors on the returns of the securities in the strategy. The IRC bias results from the fact that the hedge ratios in option strategies are calculated with prices that have measurement errors.

4.1 Simulation Procedure

We simulate true stock prices ($\tilde{S}$) and instantaneous variances ($\tilde{V}$) as well as a stock index and its variance using the model of Heston (1993). We provide full details on the simulated Heston model in the appendix. To keep the simulations numerically feasible, stocks are identical in terms of all the model parameters and differ only with respect to the starting values of the stock price and of the variance process and the amount of measurement error in the observed stock and option prices. The parameters of the simulated models are chosen to approximately match the empirical properties of the option portfolios we analyze. Specifically, the parameters we choose generate an 8% average stock return and approximately match the average stock volatility. Parameters also reflect the average pairwise correlations between

[28]See the appendix for proofs.

[29]The IMR bias results from error-induced correlations between the stock price and $\beta^f_{S,t-1}$. To see this, note that the proportional error $\epsilon_{S,t-1}$ in $\tilde{S}_{t-1}$ changes the observed $-\beta^f_{S,t-1}$ by approximately $-\frac{\partial \tilde{\beta}^f_{S,t-1}}{\partial \tilde{S}}\tilde{S}_{t-1}\epsilon_{S,t-1}$. The same error also affects the next-period observed stock return, changing it by approximately $-\epsilon_{S,t-1}$. The last term in Equation (6) is simply the expectation of the product of these two effects.

different stocks' returns, between different stocks' implied volatility changes, and between the same stock's returns and implied volatility changes. Finally, the parameters approximately match the volatility risk premium that we estimate in our adjusted FM regressions in Table 3.

We simulate 100 panels of stock and option prices, each with 500 stocks over 6,000 days. We assume 12 calls and 12 puts per stock-day, which yields 6,000 calls and 6,000 puts per day for a period of about 24 years (6,000/250), which approximately matches our actual sample. For each panel, we perform portfolio sorts and cross-sectional regressions. For computational efficiency, we assume that there is a continuum of expiration dates, and use options with one or two months in our portfolios and FM regressions. To be exact, the options have 29 and 59 calendar days at portfolio formation ($t-1$). For each maturity, there is a call and a put for six different randomly chosen strike prices K.[30] Strike prices are uniformly distributed such that they are within two standard deviations of the current stock price, where the standard deviation is the square root of the integral of expected $\tilde{V}_{i,t}$ over the lifetime of the option. We follow the common practice of using practitioner Black and Scholes hedge ratios to select samples, form portfolios, and compute delta-hedged returns.[31] In each of the tables described below, the statistics reported are averages over the 100 panels.

We assume that all stock and option prices are observed with errors. Specifically, we assume that we observe $S = \tilde{S}(1+\epsilon_S)$, $C = \tilde{C}(1+\epsilon_C)$, and $P = \tilde{P}(1+\epsilon_P)$. $\tilde{S}$, $\tilde{C}$, and $\tilde{P}$ are the true stock, call, and put prices, respectively, while S, C, and P are the observed prices. All measurement errors are drawn from symmetric triangular distributions, which are bounded distributions with probability density functions that are piecewise linear, increasing below the median and decreasing above it, reaching zero at either bound. This choice of density reflects the view that price errors are likely bounded by the size of the bid-ask spread. By choosing lower and upper bounds equal to $-1/2$ or $+1/2$ times the relative bid-ask spread, we ensure that the difference between observed prices and true prices is never larger than the spread, with differences closer to zero more likely than those further away.

[30]For index options, there are 100 different strikes. Random strike prices allow for a more realistic assessment of the CEIV bias associated with misclassification of options into portfolios.

[31]Hull and White (2017) define the practitioner Black and Scholes hedge ratios as the hedge ratios calculated with the Black and Scholes implied volatilities.

The simulation of option price errors further requires a model of relative option bid-ask spreads, which are chosen to match a number of patterns we observe in the data. Namely, we see that option spreads are higher when the underlying stock has a higher spread, when the option is further out of the money, and when the option has a shorter time until expiration. Our bid-ask spread model is motivated by the findings in De Fontnouvelle, Fishe, and Harris (2003). We provide full details on this model in the appendix.

We report simulation results for three scenarios. The first scenario uses simulated data without any measurement errors, resulting in infeasible estimates that represent the target of our bias adjustment method. The second uses simulated data with measurement errors and does not attempt to adjust for those errors, inducing the MR, RC, CEIV, and SS biases. The third scenario also uses data with errors, but it applies the methods we propose in Section 2 to reduce bias.

In addition to simulation results based on the three scenarios above, we also decompose the total bias by making different assumptions about which variables are impacted by measurement error. Specifically, we begin with a set of simulations in which only the direct component of the MR and RC biases is present (DMR and DRC). We do so by using the "true" simulated values of hedge ratios, sorting variables, and independent variables. Moreover, the sample is selected based on true prices. As a result, the indirect components of the RC and MR biases (IRC and IMR) as well as the SS and CEIV biases are not present in this first set of simulations.

We then allow for measurement errors in hedge ratios to see the incremental effect of the indirect MR and RC biases (IMR and IRC). That is, the β_S used to calculate each delta-hedged return is constructed from simulated prices containing errors. In this second set of simulations, we continue to use a sample whose selection is based on true simulated prices (avoiding the SS bias), and we use the "true" simulated values of sorting variables and independent variables (avoiding the CEIV bias).

Our third set of simulations allows us to gauge the size of the CEIV bias by introducing errors into the prices used to calculate the sorting and independent variables (moneyness and β_σ). Lastly, in the final set of simulations, which matches the scenario in which all biases are present, we make two additional changes that allow us to measure the SS bias.

First, we select the sample based on noisy prices rather than true prices. Second, we discard observations for which the hedge ratio or sorting variable is missing.[32] These two additional changes induce a covariance between sample selection and option returns that results in the SS selection bias.

4.2 Simulation Results

Table 7 presents average summary statistics for the simulations. In a number of aspects, the simulated data replicate the actual sample, as reported in Table 1. Means, medians, and standard deviations are similar, with a few exceptions. The extreme quantiles, however, are quite different in the simulations, with most variables displaying much thinner tails than we observe in the data. This difference is a natural consequence of our use of a single set of parameters to drive the price and variance processes of all stocks, a simplification that is necessary to make the simulation computationally feasible. Most importantly, the bid-ask spreads approximately match those in the actual data. For instance, for call (put) options written on individual stocks, the median relative bid-ask spread is about 9% (8%) in the simulated data, which closely match the medians that we observe in the actual data. As in the actual data, the mean and the median simulated bid-ask spreads for options written on the index are substantially smaller than those of options written on individual stocks.

We first analyze unhedged call returns, essentially replicating the empirical analysis reported in Table 2. Table 8 reports average coefficient estimates across the 100 simulated panels, the average t-statistics corresponding to those estimates, and the decomposition of the bias in the unadjusted estimates. The table shows that the returns on OTM calls calculated with prices without measurement errors (true prices) are significantly lower than returns on ITM calls. In fact, the average return of deep OTM calls is negative in the simulated model. As we explain in Section 2, this is a consequence of the negative volatility risk premium. This effect is present for both individual stocks and the stock index.

The bias decomposition reveals a complex interaction among the different biases. Panel A of Table 8 shows that the DMR, CEIV, and SS biases can be very large for calls written

[32] As in the real data, in some cases the implied volatility and hedge ratios calculated with the simulated prices with errors cannot be computed due to an apparent arbitrage violation resulting from mismeasurement.

on individual stocks, but the three biases affect each portfolio quite differently and hence it is difficult to pin down option characteristics that drive all the biases.[33]

The DMR bias is small for ITM options that trade with relatively small proportional spreads. On the other hand, the much larger percentage spreads of the deep OTM calls result in large MR biases. Deep OTM calls show the largest DMR bias, about 23 bps per day. In all cases, the MR bias is positive. Because the options are unhedged, there is no IMR bias.

The CEIV bias, on the other hand, varies in magnitude and sign from one portfolio to the next. To understand why this occurs, consider the case of an ITM call option whose implied volatility puts it close to the cutoff between the portfolio with low moneyness and Portfolio 2. If this option has a positive price error, its observed implied volatility increases, making it more likely to be assigned to Portfolio 2. (Recall that moneyness is inversely related to implied volatility.) The positive price error also means that the subsequent observed return is likely to be low. If, instead, the error were negative, then the observed moneyness would be reduced. The option would then be more likely to be placed in the low-moneyness portfolio, where its negative price error would lead to a high observed return. The patterns in the CEIV rows of the table reflect the effects of these misclassifications.

SS bias is generally negative, but the magnitude differs greatly across portfolios. For the low-moneyness (deep ITM calls) portfolio, sample selection bias is strongly negative due to the imposition of overly aggressive arbitrage filters. Standard practice is to exclude options for which the midpoint price is in violation of an arbitrage bound. The midpoints of ITM options are often below the immediate exercise value, preventing the calculation of an implied volatility based on the midpoint price. This does not actually imply arbitrage, which requires that the ask price be less than the exercise value. Discarding these observations leads to a negative sample selection bias because it tends to retain options for which pricing errors are positive. The SS bias is also quite economically significant for deep OTM options at about -18 bps per day. The low price and high bid-ask spread of OTM options results in option prices and bid-ask spreads that are often below our baseline minimum price filter and

[33]In untabulated results, we find that the biases in the FM coefficients decrease by excluding stock options with bid-ask spreads larger than 50% and zero open interest from the sample.

maximum bid-ask spreads. Consequently, price errors lead to a spurious covariance between sample selection criteria and returns for deep OTM options, which generates a large SS bias.

For call options on the index, the DMR bias is the most relevant for deep OTM calls, whereas the CEIV and SS biases are less relevant. These results are consistent with the fact that the bid-ask spreads for the simulated options on the index are in general much smaller than those of options on stocks.

Most importantly, the bias adjustment method that we propose addresses each of the biases and significantly decreases the total bias. For instance, both the average adjusted return and the true simulated return of deep OTM calls is about -38 bps per day.

The results in Table 8 indicate that a longer holding period is not necessary to address biases in the estimation of option expected returns. The empirical literature on stocks traditionally used monthly returns. The literature on options, on the other hand, tends to use high-frequency returns because these studies often rely on delta-hedged returns. The average returns in Table 8 are not based on hedged returns; hence, they could potentially be estimated with monthly holding periods, as in Ni (2008). However, the effectiveness of the bias adjustment method in Table 8 indicates that reliable average returns calculated at the daily frequency are perfectly feasible.

Table 9 presents average results for estimated FM coefficients. The table shows that biases in FM coefficients can be enormous. For put options on individual stocks, the strong negative relation between delta-hedged returns and β_σ that would be obtained without price noise is completely absent, on average, when estimated from observed prices without bias adjustment. For calls, the average slope coefficient is reduced by more than half. Biases in index options, which trade with much smaller spreads, are lower but qualitatively consistent. In all cases, price noise flattens the relationship between average returns and β_σ.

Our bias adjustment procedure performs well. In every case, the average intercept or slope coefficient is almost identical to the one obtained using true prices. Average t-statistics are also very close, except for moderate differences for a few intercepts that are close to zero on average. The table also shows that a number of different biases are present. The DRC bias, which is the one analyzed by Asparouhova, Bessembinder, and Kalcheva (2010, 2013), is sizable in most cases, but the CEV and SS biases can be larger. Furthermore, in some cases, the various biases reinforce one another, while in other cases they are partially offsetting.

One strong pattern shown in Table 9, as well as in the tables that follow, is the relatively small size of the IRC (later IMR) bias. As Equation 6 suggests, this bias arises due to measurement errors in the stock price, which are relatively small for liquid S&P 500 member stocks. As a result, the IRC bias can safely be ignored in the setting that we focus on. It is possible, of course, that options on a much less liquid underlying asset may be subject to a larger IRC bias, but untabulated results suggest that the decrease in liquidity would have to be very substantial. Moreover, we show in the appendix that the IMR bias is not small in the case of straddles. In straddles, the IMR bias depends on the variances of the measurement errors in calls and puts, which are much larger than those in stocks. As a result, the simple bias adjustment procedure that we propose does not work for straddles. Instead, a more elaborate scheme, which we describe in the appendix, is needed.

Our last observation regarding Table 9 is that the effect of bias adjustment in the simulation is quite similar to its effect in actual data, as reported in Table 3. In both cases, unadjusted intercept coefficients can be very negative, with some implausibly low t-statistics, and unadjusted slope coefficients are in many cases small, insignificant, or even positive. Bias adjustment raises intercepts and lowers slopes, by similar amounts across the two tables.

Table 10 reports the results of portfolio sorts of delta-hedged options on individual equity options. As in Table 8, the results in Table 10 show that DMR, CEIV, and SS biases can be very large, but the three biases affect each quintile quite differently. As with the FM coefficient estimates, the IMR bias is small enough to ignore.

As in the FM coefficients reported in Table 9, the total effect of these biases is a much weaker relationship between β_σ and average delta-hedged returns. The results of the portfolio sorts in Table 9 help us gain more clarity about the reasons for this weakness. Specifically, Low β_σ options have lower returns, mostly as the result of SS bias, while high β_σ options have higher returns due to a combination of DMR and CEIV biases. The effects are particularly strong for puts, for which the return on the high-minus-low portfolio is close to zero on average without bias adjustment, even though the true prices imply a large negative mean.

Bias adjustment is successful in returning average estimates that are very close to the infeasible estimates computed using true prices. Average t-statistics are also very close for most portfolios, though they are somewhat lower for low-β_σ portfolios. In general, the adjustments we propose are able to reduce the total bias by 90%, if not completely.

Table 11 repeats the same analysis but uses simulations of index options. Overall, the patterns are similar but subdued, as biases are generally small, only approaching economic significance for deep OTM options with high values of β_σ. As before, the bias adjustment procedure we propose is highly successful.

4.3 Model Dependence

An important question is how much the assumptions in our simulations affect our conclusion that the MR, RC, CEIV, and SS biases are economically significant. Naturally, our simulations are not meant to deliver exact estimates of the size of these biases. However, they do show that they can be economically significant and that they need to be addressed.

Three assumptions are crucial in our simulations. We next discuss how appropriate each of these assumptions is and how they affect our results with respect to the economic significance of the biases.

Our simulations use practitioner Black and Scholes hedge ratios, which are based on the Black and Scholes formula calculated with the volatility implied by the observed price of the same option. Practitioner Black and Scholes hedge ratios are by far the most commonly used hedge ratios in both industry (see Hull and White (2017)) and academia.[34] Consequently, these are the most appropriate hedge ratios to gauge the significance of the MR, RC, CEIV, and SS biases present in work done by practitioners and academics.

Second, to make the simulation computationally feasible, we use a single set of parameters to drive processes of all stocks in our simulations. However, our documented MR, RC, CEIV, and SS biases are independent of the simulated pricing model. To see this, note the following: First, the MR biases in delta-hedged and unhedged options depend only on the variances of the measurement errors ($E[\epsilon^2_{S,t-1}]$, $E[\epsilon^2_{f,t} - 1]$) and on $E[\beta^f_{S,t-1}]$ (see Equation 6). The variances of the measurement errors are independent of the simulated pricing models,

[34] Coval and Shumway (2001), Bakshi and Kapadia (2003a,b), Santa-Clara and Saretto (2009), Cao and Han (2013), Constantinides, Jackwerth, and Savov (2013), and Cao, Han, Tong, and Xintong (2017) are a few examples of academic papers that use practitioner Black and Scholes hedge ratios.

while $\beta^f_{S.t-1}$ are all calculated using the practitioner Black and Scholes approach, as is the common practice. Second, the CEIV and SS biases result from the correlation between the measurement errors in returns and the sorting variables (moneyness and β_σ) or selection variables (e.g., option prices). Again, the measurement errors are independent of the simulated pricing model, and the use of practitioner Black Scholes betas and moneyness reflects common practice rather than a modeling assumption. Applying the same type of argument, we also conclude that our RC bias results are independent of the model used for simulation.

Third, we assume that measurement errors have a triangular distribution with support within the bid-ask spread. Even though our simulation results are unaffected by the simulated pricing model, they are affected by the assumed distribution of measurement errors, because this distribution drives the variance of the measurement errors ($E[\epsilon^2_{S,t-1}]$, $E[\epsilon^2_{C,t-1}]$, and $E[\epsilon^2_{P,t-1}]$). As a result, one may question our estimates of the size of the simulated biases since they rely on an unobserved distribution. We note, however, that our assumption of triangular distribution is conservative. In fact, the biases are larger under the assumption that measurement errors are uniformly distributed. Moreover, the variances of the measurement errors are directly related to the support of the measurement error distribution, which is $\pm 1/2$ times the relative bid-ask spread. Given that the option spreads in the data are enormous, it seems unlikely that the variances of the simulated measurement errors are much larger than the actual ones. Hence, our simulation results and the enormous option relative bid-ask spreads in the data are strong indications that the MR, RC, CEIV, and SS biases are in many cases highly economically significant and too large for empirical studies to ignore.

5 Conclusion

We find that adjusting for microstructure biases in the estimation of expected returns is consequential for one of the most well-known stylized facts in the empirical option literature. Specifically, we show that the volatility risk premium on options written on stocks is not as different from that of options written on the S&P 500 Index as prior research suggests. Consistent with a negative volatility risk premium, we show that call option expected returns do not increase monotonically with moneyness. In addition, bias-adjusted FM regressions indicate that the volatility risk premium in individual equity options is economically large at about -5 bps per day. Finally, using a sample that includes a broad array of traded options, we find that the variance risk premium (VRP) is on average negative, at about -1.5 percentage points.

Our results indicate that the option literature needs to adopt procedures to deal with microstructure biases in the same way that the empirical equity literature has adopted approaches to deal with these biases at least since Blume and Stambaugh (1983). Our results, together with the large magnitude of the relative bid-ask spreads in options, constitute strong evidence that the microstructure biases in options do matter. Our findings also suggest that the current approach in the options literature of focusing only on close to ATM options can lead to conclusions that are not robust. Indeed, both the VRP estimated only with options close to ATM and the and bias-unadjusted FM regressions suggest that the volatility risk premium in equity options is zero. However, a zero volatility risk premium cannot explain the results in our entire sample. Overall, our findings therefore suggest that the option literature needs to move away from the approach of focusing only on a subset of the data to a systematic approach that addresses the microstructure biases in the entire sample. The bias-adjusted procedure that we propose is an initial contribution in this direction.

Our finding that volatility is priced in individual equity options is an important step towards understanding the price of volatility risk in equity options. Our results show that the price of volatility risk in equity options is negative and economically significant. There is, however, much to be done on understanding how the volatility risk premium varies across stocks and through time. These are topics for future research.

References

Asparouhova, Elena, Hendrick Bessembinder, and Ivalina Kalcheva, 2010, Liquidity biases in asset pricing tests, *Journal of Financial Economics* 96, 215–237.

———, 2013, Noisy prices and inference regarding returns, *Journal of Finance* 68, 665–714.

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

———, 2003b, Volatility risk premiums embedded in individual equity options: Some new insights, *Journal of Derivatives* 11, 45–54.

Battalio, Robert, Brian Hatch, and Robert Jennings, 2004, Toward a national market system for U.S. exchange–listed equity options, *Journal of Finance* 59, 933–962.

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

Blume, Marshall, and Robert Stambaugh, 1983, Biases in computed returns, *Journal of Financial Economics* 12, 387–404.

Bogousslavsky, Vincent, and Dmitriy Muravyev, 2021, Who trades at the close? Implications for price discovery and liquidity, *Working Paper*.

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.

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

———, Qing Tong, and Zhan Xintong, 2017, Option return predictability, *Working paper*.

Cao, Melanie, and Jason Wei, 2010, Option market liquidity: Commonality and other characteristics, *Journal of Financial Markets* 13, 20–48.

Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.

Christoffersen, Peter, Ruslan Goyenko, Kris Jacobs, and Mehdi Karoui, 2018, Illiquidity premia in the equity options market, *Review of Financial Studies* 31, 811–851.

Constantinides, George M., Jens Carsten Jackwerth, and Alexi Savov, 2013, The puzzle of index option returns, *Review of Asset Pricing Studies* 3, 229–257.

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

Cox, John C., Ross Stephen A., and Mark Rubinstein, 1979, Option pricing: A simplified approach, *Journal of Financial Economics* 7, 229–263.

De Fontnouvelle, Patrick, Raymond P. H. Fishe, and Jeffrey H. Harris, 2003, The behavior of bid-ask spreads and volume in options markets during the competition for listings in 1999, *Journal of Finance* 58, 2437–2463.

Dennis, Patrick, and Stewart Mayhew, 2009, Microstructural biases in empirical tests of option pricing models, *Review of Derivatives Research* 12, 169–191.

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

Duarte, Jefferson, and Christopher S. Jones, 2010, The price of market volatility risk, *Working paper.*

Eraker, Bjørn, and Daniela Osterrieder, 2018, Market maker inventory, bid-ask spreads, and the computation of option implied risk measures, *Working paper.*

Fama, Eugene F., 1984, The information in the term structure, *Journal of Financial Economics* 13, 509–528.

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

Fournier, Mathieu, Kris Jacobs, and Piotr Orlowski, 2021, Modeling conditional factor risk premia implied by index option returns, *Working paper.*

French, Kenneth R., 1983, A comparison of futures and forward prices, *Journal of Financial Economics* 12, 311–342.

Goyal, Amit, and Alessio Saretto, 2009, Cross-section of option returns and volatility, *Journal of Financial Economics* 94, 310–326.

Goyenko, Ruslan, and Chengyu Zhang, 2019, Option returns: Closing prices are not what you pay, *Working paper*.

Hentschel, Ludger, 2003, Errors in implied volatility estimation, *Journal of Financial and Quantitative Analysis* 38, 779–810.

Heston, Steven L, 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–43.

Hull, John, and Alan White, 2017, Optimal delta hedging for options, *Journal of Banking and Finance* 82, 180 – 190.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

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

——— , and Neil D Pearson, 2020, Options Trading Costs Are Lower than You Think, *Review of Financial Studies* 33, 4973–5014.

Ni, Sophie Xiaoyan, 2008, Stock option returns: A puzzle, *Working Paper*.

Rehman, Zahid, and Grigory Vilkov, 2012, Risk-neutral skewness: Return predictability and its sources, *Working Paper*.

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1127–1139.

Santa-Clara, Pedro, and Alessio Saretto, 2009, Option strategies: Good deals and margin calls, *Journal of Financial Markets* 12, 391 – 417.

Stambaugh, Robert F., 1988, The information in forward rates: Implications for models of the term structure, *Journal of Financial Economics* 21, 41–70.

Wolak, Frank A, 1989, Testing inequality constraints in linear econometric models, *Journal of Econometrics* 41, 205–235.

Zhang, Lan, Per A. Mykland, and Yacine Aït-Sahalia, 2005, A tale of two time scales: Determining integrated volatility with noisy high-frequency data, *Journal of the American Statistical Association* 100, 1394–1411.

Table 1: **Summary statistics.** This table presents the mean, standard deviation, median, first and third quartiles (Q1 and Q3), and first and 99^{th} percentiles (P1 and P99), as well as skewness of current daily returns, relative spreads, moneyness, implied volatilities, and the dollar value of open interest. The relative option spread is given by $2 \times (\text{ask} - \text{bid})/(\text{ask} + \text{bid})$, where ask and bid are the closing ask and bid option prices from IvyDB. The moneyness of an option (with maturity at time T) at time t is $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where S_t is the observed closing stock price on day t; K is the option strike price; r_t is the risk-free rate between t and the option expiration date T from IvyDB; and σ_t is the volatility implied by the closing option price from IvyDB. The dollar value of open interest ($\$OI$) is the closing mid-price of the option multiplied by its open interest from IvyDB. The relative bid-ask spread of a stock for a given day is the volume weighted average of the relative effective bid-ask spread of all transactions during the day. It is from WRDS Intraday Indicators. Return is displayed in basis points. Spreads and σ are displayed in percentages. Our sample period is from January 1996 to June 2019.

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
A: Call options on S&P 500 stocks (Avg. number of options per day: 5,503)								
Return	52.97	-129.45	3,268.91	-5,750.41	-1,333.37	991.91	9,582.72	35.94
Relative spread	11.92	8.00	10.96	0.83	4.08	15.89	48.65	1.53
Moneyness	-0.20	-0.17	0.93	-2.06	-0.93	0.52	1.71	-0.05
σ	35.71	30.73	19.95	13.27	23.31	41.69	113.31	2.90
$\$OI$	4,404.88	277.20	24,364.01	0.00	38.72	1,833.83	70,307.75	30.70
B: Put options on S&P 500 stocks (Avg. number of options per day: 5,427)								
Return	-62.56	-296.84	3,162.12	-5,531.25	-1,578.55	907.80	9,998.05	15.90
Relative spread	13.19	9.09	11.63	0.88	4.55	18.18	50.00	1.34
Moneyness	-0.06	-0.13	0.95	-1.87	-0.77	0.64	1.95	-0.19
σ	35.95	31.09	19.79	13.41	23.74	41.89	112.37	2.99
$\$OI$	2,823.42	178.20	27,144.64	0.00	26.33	1,105.50	43,907.85	35.65
C: Call options on S&P 500 Index (Avg. number of options per day: 190)								
Return	82.84	-80.86	2,876.85	-5,720.24	-1,255.82	950.44	10,487.26	2.28
Relative spread	7.65	3.39	9.97	0.50	1.76	8.70	46.15	2.24
Moneyness	-0.01	-0.11	1.11	-2.09	-0.92	0.87	2.16	0.16
σ	16.62	14.33	9.77	6.20	10.69	19.46	54.32	3.93
$\$OI$	4,442.42	182.00	13,338.47	1.00	21.00	2,417.00	64,648.00	6.82
D: Put options on S&P 500 Index (Avg. number of options per day: 274)								
Return	-221.38	-689.99	3,597.74	-5,581.77	-2,044.84	697.55	12,470.84	6.37
Relative spread	9.27	5.97	9.19	0.95	3.25	11.43	44.44	2.04
Moneyness	-0.84	-1.07	1.03	-2.24	-1.67	-0.20	1.90	0.81
σ	21.01	19.70	9.37	6.90	14.54	25.54	53.03	1.82
$\$OI$	4,821.79	253.00	15,113.04	1.00	35.00	2,436.00	75,745.00	6.78
E: Stocks (Avg. number of stocks per day: 483)								
Return	4.83	1.84	235.87	-636.33	-95.28	102.75	675.68	0.27
Effective spread	0.15	0.08	0.46	0.02	0.05	0.18	0.94	4.83

Table 2: **Mean returns of unhedged call options on S&P 500 stocks and on the S&P 500 Index.** Panels A and B present the average return of stock and index call options sorted by moneyness. The moneyness of an option at time t is $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where S_t is the observed closing stock price on day t; K is the option strike price; r_t is the risk-free rate between t and the option expiration date T from IvyDB; and σ_t is the volatility implied by the closing option price from IvyDB. Moneyness groups are defined as follows. Low: Moneyness < -1.25. Group 2: $-1.25 \leq$ Moneyness < -0.5. Group 3: $-0.5 \leq$ Moneyness < 0.5. Group 4: $0.5 \leq$ Moneyness < 1.25. High: Moneyness ≥ 1.25. The return of call option i is $(C_{i,t} - C_{i,t-1})/C_{i,t-1}$. The bias-unadjusted results are calculated with the *baseline* sample, which is constructed with selection criteria based on $t-1$ variables. Moreover, in the bias-unadjusted method, calls are sorted by their moneyness at the beginning of the holding period $(t-1)$. In contrast, the bias-adjusted results are calculated with a sample constructed with selection criteria based on $t-2$ variables, calls are sorted by their moneyness one day before the beginning of the holding period $(t-2)$, and returns are weighted by one-day-lagged call option gross returns $(C_{i,t-1}/C_{i,t-2})$. Section 1 describes the sample selection criteria. β_S and β_σ are the option sensitivities with respect to the underlying price and volatility. The median β_S, β_σ, relative spread, and moneyness are the mean values across time of the median $t-2$ values of these variables. Mean returns are displayed in basis points. Relative spreads are displayed in percentages. We also present results weighted by the dollar value of open interest ($\$OI_{i,t-2}$). The p-values are for a Wolak (1989) test of the null hypothesis that call option mean returns increase monotonically with moneyness. T-statistics are shown in parentheses.

	Low	2	3	4	High	H-L	
A: Stocks							
Median β_σ	0.15	0.80	3.24	8.31	13.17	13.03	
Median β_S	4.97	8.50	13.70	19.59	22.01	17.04	
Median relative spread	4.02	5.78	8.22	18.89	31.35	27.33	
Median moneyness	-1.58	-0.86	0.00	0.81	1.43	3.01	
	Average daily call option returns						Wolak p-value
Bias-unadjusted							
Equal weighted	1.39	0.03	51.59	83.50	-14.70	-15.93	0.00
	(0.21)	(0.00)	(3.09)	(3.75)	(-0.55)	(-0.70)	
Weighted by $\$OI_{i,t-2}$	8.50	10.33	25.87	-8.26	-161.40	-169.75	0.00
	(1.17)	(0.91)	(1.53)	(-0.36)	(-5.52)	(-6.76)	
Bias-adjusted							
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}}$	12.96	22.05	47.99	73.64	38.77	25.85	0.03
	(1.97)	(2.07)	(2.92)	(3.37)	(1.42)	(1.12)	
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}} \times \$OI_{i,t-2}$	15.99	19.56	32.99	24.84	-73.48	-89.39	0.00
	(2.22)	(1.75)	(1.98)	(1.09)	(-2.46)	(-3.46)	
B: Index							
	Low	2	3	4	High	H-L	
Median β_σ	0.32	1.48	6.47	20.65	39.36	39.04	
Median β_S	8.69	14.51	26.49	45.83	63.99	55.31	
Median relative spread	1.66	2.71	5.56	13.04	28.45	26.79	
Median moneyness	-1.54	-0.86	-0.03	0.85	1.57	3.11	
	Average daily call option returns						Wolak p-value
Bias-unadjusted							
Equal weighted	20.81	29.76	41.86	61.49	-59.29	-103.69	0.00
	(2.02)	(1.77)	(1.43)	(1.29)	(-1.05)	(-2.16)	
Weighted by $\$OI_{i,t-2}$	21.60	27.67	39.01	58.61	-12.74	-59.60	0.00
	(2.07)	(1.60)	(1.37)	(1.27)	(-0.22)	(-1.21)	
Bias-adjusted							
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}}$	20.07	28.96	40.47	29.93	-52.11	-94.88	0.00
	(1.96)	(1.73)	(1.38)	(0.63)	(-0.95)	(-1.98)	
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}} \times \$OI_{i,t-2}$	19.09	26.29	38.17	27.70	-41.03	-86.53	0.00
	(1.86)	(1.52)	(1.34)	(0.61)	(-0.74)	(-1.80)	

Table 3: **FM regressions of returns of delta-hedged options on S&P 500 stocks and on the S&P 500 Index.** This table displays the results of FM regressions: $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t} = \lambda_0 + \lambda_\sigma \times \beta^f_{\sigma,i,t-1} + \eta_{i,t}$. The dependent variable is delta-hedged option excess return, where $R_{f,i,t}$ is the excess return of call or put i between $t-1$ and t; $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$ is the β of the call or put with respect to the underlying asset; and $R_{S,t}$ is the excess return of the underlying stock or index between $t-1$ and t. In the bias-unadjusted method, the independent variable is the β of the option with respect to the volatility of the underlying at time $t-1$ ($\beta^f_{\sigma,i,t-1}$). Moreover, the sample used in the bias-unadjusted method is constructed with selection criteria based on $t-1$ variables, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, in the bias-adjusted method, the independent variable is $\beta^f_{\sigma,i,t-2}$, and the regressions are estimated with WLS using gross returns as weights (either $C_{i,t-1}/C_{i,t-2}$ or $P_{i,t-1}/P_{i,t-2}$). Section 1 describes the sample selection criteria. We also present results of regressions estimated with WLS using the dollar value of open interest ($\$OI_{i,t-2}$) as weights. The options are written on the S&P 500 Index and individual stocks in the S&P 500 Index. T-statistics are shown in parentheses.

	A: Bias-unadjusted			
	Stocks		Index	
	λ_0	λ_σ	λ_0	λ_σ
Calls				
OLS	-17.83	1.74	17.37	-6.70
	(-11.34)	(1.24)	(3.92)	(-4.73)
WLS, weighted by $\$OI_{i,t-2}$	-8.81	-3.25	4.79	-5.25
	(-7.43)	(-2.00)	(1.56)	(-3.77)
Puts				
OLS	-24.26	1.58	32.24	-12.89
	(-13.22)	(1.07)	(2.35)	(-5.34)
WLS, weighted by $\$OI_{i,t-2}$	-24.56	2.55	32.88	-11.80
	(-13.66)	(1.50)	(4.10)	(-6.13)
	B: Bias-adjusted			
	Stocks		Index	
	λ_0	λ_σ	λ_0	λ_σ
Calls				
WLS, weighted by $\frac{C_{i,t-1}}{C_{i,t-2}}$	-3.57	-0.32	8.06	-6.36
	(-2.43)	(-0.24)	(1.85)	(-4.26)
WLS, weighted by $\frac{C_{i,t-1}}{C_{i,t-2}} \times \$OI_{i,t-2}$	-0.67	-5.40	-2.00	-5.79
	(-0.58)	(-3.57)	(-0.69)	(-3.98)
Puts				
WLS, weighted by $\frac{P_{i,t-1}}{P_{i,t-2}}$	-2.39	-5.04	32.24	-14.67
	(-1.41)	(-3.76)	(1.72)	(-6.77)
WLS, weighted by $\frac{P_{i,t-1}}{P_{i,t-2}} \times \$OI_{i,t-2}$	-6.98	-4.72	35.19	-13.10
	(-5.16)	(-3.01)	(3.41)	(-7.00)

Table 4: **Mean excess returns of delta-hedged calls and puts on individual stocks in the S&P 500 Index.** This table presents average delta-hedged returns of calls and puts sorted by their β with respect to the volatility of the underlying asset (β_σ). Delta-hedged option excess returns are $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$, where f_{t-1} is the price of the call or put i at time $t-1$; $R_{f,i,t}$ is its excess return between $t-1$ and t; $R_{S,t}$ is the excess return of the underlying stock; and $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$. The sample used in the bias-unadjusted method is built with selection criteria based on $t-1$ variables, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. The sorting variable is $\beta^f_{\sigma,i,t-1}$ in the bias-unadjusted method. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, the bias-adjusted method uses $\beta^f_{\sigma,i,t-2}$ as the sorting variable, and average returns are calculated with weighted averages using gross returns as weights. Section 1 describes the sample selection criteria. The median β_σ, relative spread, and moneyness are the mean values across time of the median $t-2$ values of these variables. Mean returns are displayed in basis points. Relative spreads are displayed in percentages. We also present averages weighted by the dollar value of open interest ($\$OI_{i,t-2}$). T-statistics are shown in parentheses.

	A: Calls						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.19	0.86	2.24	4.83	9.97	9.78	2.24
Median relative spread	4.25	5.54	6.92	10.79	22.00	17.75	7.60
Median moneyness	-1.50	-0.84	-0.26	0.32	0.93	2.44	-0.26
	Average daily delta-hedged returns						
Bias-unadjusted							
Equal weighted	-30.96	-18.10	-1.65	-2.74	-0.10	30.86	-10.70
	(-31.73)	(-8.74)	(-0.40)	(-0.35)	(-0.01)	(2.21)	(-1.89)
Weighted by $\$OI_{i,t-2}$	-15.97	-10.56	-10.72	-22.95	-40.92	-24.95	-19.09
	(-20.02)	(-4.93)	(-2.47)	(-2.94)	(-2.96)	(-1.86)	(-5.33)
Bias-adjusted							
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}}$	-2.30	-1.79	-1.68	-7.07	-10.65	-8.36	-3.73
	(-3.36)	(-0.90)	(-0.41)	(-0.92)	(-0.75)	(-0.61)	(-0.68)
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}} \times \$OI_{i,t-2}$	-1.57	-3.95	-12.03	-25.08	-53.04	-51.46	-11.88
	(-1.84)	(-1.88)	(-2.83)	(-3.38)	(-4.06)	(-4.05)	(-3.54)
	B: Puts						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.33	1.48	3.32	5.90	9.90	9.56	3.32
Median relative spread	4.80	6.04	8.19	13.24	23.75	18.96	8.80
Median moneyness	1.31	0.56	-0.05	-0.58	-1.11	-2.42	-0.04
	Average daily delta-hedged returns						
Bias-unadjusted							
Equal weighted	-34.51	-15.53	-12.51	-28.39	-20.58	13.93	-22.30
	(-36.96)	(-7.49)	(-2.84)	(-3.74)	(-1.70)	(1.17)	(-4.36)
Weighted by $\$OI_{i,t-2}$	-23.99	-17.10	-19.46	-22.82	2.09	26.08	-22.86
	(-28.76)	(-6.84)	(-3.94)	(-2.85)	(0.14)	(1.81)	(-6.96)
Bias-adjusted							
Weighted by $\frac{P_{i,t-1}}{P_{i,t-2}}$	-3.94	-5.85	-18.19	-44.93	-65.84	-61.90	-25.35
	(-5.07)	(-2.72)	(-4.07)	(-5.97)	(-5.69)	(-5.54)	(-4.98)
Weighted by $\frac{P_{i,t-1}}{P_{i,t-2}} \times \$OI_{i,t-2}$	-4.92	-12.95	-22.87	-38.53	-54.01	-49.09	-18.41
	(-5.89)	(-5.10)	(-4.65)	(-4.87)	(-4.24)	(-3.96)	(-5.78)

Table 5: **Mean excess returns of delta-hedged calls and puts written on the S&P 500 Index.** This table presents average delta-hedged returns of calls and puts sorted by their β with respect to the volatility of the underlying asset (β_σ). Delta-hedged option excess returns are $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$, where f_{t-1} is the price of the call or put i at time $t-1$; $R_{f,i,t}$ is its excess return between $t-1$ and t; $R_{S,t}$ is the excess return of the underlying stock; and $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$. The sample used in the bias-unadjusted method is built with selection criteria based on $t-1$ variables, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. The sorting variable is $\beta^f_{\sigma,i,t-1}$ in the bias-unadjusted method. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, the bias-adjusted method uses $\beta^f_{\sigma,i,t-2}$ as the sorting variable, and average returns are calculated with weighted averages using gross returns as weights. Section 1 describes the sample selection criteria. The median β_σ, relative spread, and moneyness are the mean values across time of the median $t-2$ values of these variables. Mean returns are displayed in basis points. Relative spreads are displayed in percentages. We also present averages weighted by the dollar value of open interest ($\$OI_{i,t-2}$). T-statistics are shown in parentheses.

	A: Calls						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.52	1.94	5.31	13.77	32.46	31.94	5.53
Median relative spread	1.82	2.99	4.93	9.07	21.84	20.02	5.08
Median moneyness	-1.40	-0.77	-0.20	0.45	1.29	2.69	-0.17
	Average daily delta-hedged returns						
Bias-unadjusted							
Equal weighted	11.82	5.82	-4.65	-27.72	-178.26	-189.37	-40.29
	(5.34)	(1.31)	(-0.52)	(-1.45)	(-4.53)	(-4.94)	(-2.83)
Weighted by $\$OI_{i,t-2}$	13.33	4.90	-7.35	-26.09	-115.18	-127.81	-9.70
	(6.13)	(1.09)	(-0.86)	(-1.46)	(-3.41)	(-3.90)	(-1.30)
Bias-adjusted							
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}}$	-2.59	-7.98	-17.42	-48.71	-197.85	-195.27	-53.53
	(-1.21)	(-1.64)	(-1.62)	(-2.27)	(-4.76)	(-4.80)	(-3.64)
Weighted by $\frac{C_{i,t-1}}{C_{i,t-2}} \times \$OI_{i,t-2}$	-3.01	-9.11	-21.65	-47.04	-159.22	-156.22	-17.67
	(-1.41)	(-1.91)	(-2.10)	(-2.37)	(-4.47)	(-4.50)	(-2.47)
	B: Puts						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	3.13	7.96	12.36	16.21	19.76	16.62	12.58
Median relative spread	3.73	5.76	9.46	15.43	24.64	20.91	9.09
Median moneyness	0.72	-0.13	-0.76	-1.27	-1.73	-2.45	-0.77
	Average daily delta-hedged returns						
Bias-unadjusted							
Equal weighted	3.12	-30.95	-76.66	-153.86	-216.37	-219.82	-97.33
	(0.61)	(-2.84)	(-3.95)	(-6.01)	(-6.17)	(-6.94)	(-5.21)
Weighted by $\$OI_{i,t-2}$	4.83	-30.73	-75.84	-143.46	-182.22	-187.38	-55.27
	(0.90)	(-2.83)	(-4.37)	(-5.86)	(-5.24)	(-6.02)	(-4.88)
Bias-adjusted							
Weighted by $\frac{P_{i,t-1}}{P_{i,t-2}}$	-3.87	-33.07	-79.85	-149.52	-228.48	-224.94	-99.01
	(-0.76)	(-3.05)	(-4.62)	(-6.48)	(-6.71)	(-7.32)	(-5.62)
Weighted by $\frac{P_{i,t-1}}{P_{i,t-2}} \times \$OI_{i,t-2}$	-5.44	-33.17	-78.09	-148.21	-229.13	-224.02	-60.54
	(-1.00)	(-3.02)	(-4.55)	(-6.14)	(-6.60)	(-7.20)	(-5.38)

Table 6: **Averages of realized variance (RV) and model-free implied variance ($MFIV$) for stocks in the S&P 500 and for the S&P 500 Index.** RV is calculated from daily returns over a one-month window, and MFIV is calculated from a cross-section of one-month options. Panel A presents the results for stocks in the S&P 500. The first column of Panel A presents results using the method in Driessen, Maenhout, and Vilkov (2009) to calculate $MFIV$. This method uses OTM calls and puts with $\Delta_C > 0.15$ and $\Delta_P < -0.05$. The second column presents results using our sample, which includes OTM options with $\Delta_C > 0.01$ and $\Delta_P < -0.01$. The third column presents results using the $MFIV$ in Rehman and Vilkov (2012). For options on stocks, $\overline{MFIV}$ ($\overline{RV}$) is the time series average of the equally weighted cross-sectional averages of the $MFIV$ (RV) across all stocks. Panel B presents the results for the S&P 500 Index. The first column of Panel B presents results using the sample selection criteria in Driessen, Maenhout, and Vilkov (2009). The second column presents results using our sample selection criteria. The third column presents results using the VIX as an estimate for the $MFIV$ of the S&P 500 Index. For options on the index, $\overline{MFIV}$ ($\overline{RV}$) is time series average of the index $MFIV$ (RV). T-statistics are shown in parentheses.

		A: Stocks	
	OTM Only $\Delta_C > 0.15$ $\Delta_P < -0.05$	OTM Extended $\Delta_C > 0.01$ $\Delta_P < -0.01$	Rehman & Vilkov (2012)
Average number of stocks per day	253.9	351.9	468.6
Average number of options per stock	5.3	7.3	
$\overline{MFIV}$	0.4136^2	0.4157^2	0.3999^2
$\overline{RV}$	0.4176^2	0.3958^2	0.3816^2
$\overline{RV} - \overline{MFIV}$	0.0033	-0.0161	-0.0143
	(0.47)	(-2.52)	(-2.41)
	Tests of $H_0 : \overline{RV} - \overline{MFIV} = 0$ for each stock		
	Number of stocks		
$t - statistic > 1.96$	1	0	3
$t - statistic < -1.96$	67	334	379
$\lvert t - statistic \rvert < 1.96$	501	358	479
		B: Index	
	OTM Only $\Delta_C > 0.15$ $\Delta_P < -0.05$	OTM Extended $\Delta_C > 0.01$ $\Delta_P < -0.01$	VIX
Average number of options per day	36.1	74.6	
$\overline{MFIV}$	0.2110^2	0.2421^2	0.2160^2
$\overline{RV}$	0.1921^2	0.1921^2	0.1921^2
$\overline{RV}$ - $\overline{MFIV}$	-0.0076	-0.0217	-0.0098
	(-2.66)	(-6.36)	(-3.39)

Table 7: **Summary statistics for simulated option data.** This table presents mean, median, standard deviation, first and third quartiles (Q1 and Q3), and first and 99^{th} percentiles (P1 and P99), as well as skewness of the simulated daily returns, relative spreads, moneyness, and implied volatilities. The moneyness of an option (with maturity at time T) at time t is $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where S_t is the observed closing stock price on day t; K is the option strike price; and r_t is the risk-free rate, which we assume is zero in the simulation. The stock and index returns as well as volatilities are simulated with the Heston (1993) stochastic volatility model. The displayed statistics on returns, moneyness, and stochastic volatilities are based on the simulated values without measurement errors. The simulated bid-ask spreads are based on an empirical model of the bid-spreads motivated by the findings in De Fontnouvelle, Fishe, and Harris (2003). Section 4.1 explains the simulation procedure. Return is displayed in basis points. Relative spread and σ are displayed in percentages.

A: Call options on simulated stocks

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
Return	0.09	-127.86	2,389.05	-5,346.40	-1,316.94	1,081.60	7,627.67	1.13
Relative spread	12.51	8.74	10.99	0.77	4.31	17.35	46.68	1.34
Moneyness	-0.05	-0.08	1.03	-1.96	-0.92	0.79	1.90	0.05
σ	39.31	38.73	7.39	24.71	34.21	43.73	59.68	0.69

B: Put options on simulated stocks

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
Return	-38.74	-102.42	2,056.97	-5,008.89	-1,097.60	872.97	6,453.00	0.95
Relative spread	11.43	7.69	10.63	0.51	3.69	15.68	46.08	1.47
Moneyness	0.33	0.40	1.04	-1.71	-0.52	1.21	2.17	-0.17
σ	39.10	38.57	7.58	23.77	33.92	43.63	59.75	0.68

C: Call options on simulated index

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
Return	1.85	-65.85	2,220.96	-5,386.64	-991.76	812.73	7,433.94	1.53
Relative spread	8.13	3.54	10.42	0.08	1.17	10.91	45.21	1.86
Moneyness	-0.22	-0.33	1.28	-2.24	-1.34	0.82	2.27	0.22
σ	16.48	16.22	3.55	9.51	13.93	18.74	25.76	0.45

D: Put options on simulated index

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
Return	-99.72	-159.65	2,468.13	-5,861.15	-1,348.29	933.39	7,866.86	1.14
Relative spread	6.92	2.06	10.27	0.00	0.01	9.78	44.29	1.91
Moneyness	-0.05	-0.12	1.32	-2.24	-1.19	1.08	2.30	0.11
σ	16.20	15.96	3.43	9.40	13.73	18.41	25.06	0.39

E: Simulated stocks

	Mean	Median	Std Dev	P1	Q1	Q3	P99	Skewness
Return	2.18	0.16	198.87	-474.37	-125.09	127.05	498.87	0.08
Relative spread	0.11	0.09	0.07	0.02	0.06	0.13	0.34	2.10

Table 8: **Mean returns of simulated call options sorted by moneyness.** Panels A and B present the average of simulated returns of stock and index call options sorted by moneyness. The moneyness of an option (with maturity at time T) at time t is $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where S_t is the stock price; K is the option strike price; r_t is the risk-free, which is set equal to zero; and σ_t is the volatility implied by the simulated option price. Moneyness groups are defined as follows. Low: Moneyness < -1.25. Group 2: $-1.25 \leq$ Moneyness < -0.5. Group 3: $-0.5 \leq$ Moneyness < 0.5. Group 4: $0.5 \leq$ Moneyness < 1.25. High: Moneyness ≥ 1.25. β_S and β_σ are the options betas with respect to the underlying price and the volatility. The return of call option i is $(C_{i,t} - C_{i,t-1})/C_{i,t-1}$. "True prices" refer to results calculated with simulated prices without measurement errors. The bias-unadjusted results are calculated with a simulated sample constructed with selection criteria based on $t-1$ variables that contain measurement errors. Moreover, in the bias-unadjusted method, calls are sorted by their moneyness at the beginning of the holding period $(t-1)$. In contrast, the bias-adjusted results are calculated with a sample constructed with selection criteria based on $t-2$ variables that contain measurement errors, calls are sorted by their moneyness one day before the beginning of the holding period $(t-2)$, and returns are weighted by the one-day-lagged call option gross return $(C_{i,t-1}/C_{i,t-2})$. The statistics on β_S, β_σ, relative spreads, and moneyness are the mean values across time of the median $t-2$ values of these variables. The total bias is decomposed into its different parts: the direct mean return bias (DMR), indirect mean return bias (IMR), CEIV bias, and sample-selection bias (SS). Section 4.1 describes the simulation procedure. The p-values shown are average values, across simulation trials, of p-values of a Wolak (1989) test with the null hypothesis that call option mean returns increase monotonically with moneyness. Relative spreads are displayed in percentages. Returns and biases are in basis-points per day. Average t-statistics are shown in parentheses.

A: Stocks							
	Low	2	3	4	High	H-L	
Median β_σ	0.20	0.69	2.60	6.90	12.28	12.09	
Median β_S	4.90	6.55	10.05	14.56	17.87	12.98	
Median relative spread	3.97	6.03	10.02	15.40	17.29	13.33	
Median moneyness	-1.51	-0.89	-0.01	0.88	1.56	3.07	
	Average daily call option returns						Wolak p-value
True prices	9.06	9.73	5.05	-11.99	-37.95	-47.01	0.00
	(1.49)	(1.15)	(0.40)	(-0.66)	(-1.67)	(-2.73)	
Bias-unadjusted	20.97	-12.58	5.05	4.12	-17.59	-38.55	0.00
	(3.39)	(-1.51)	(0.40)	(0.22)	(-0.77)	(-2.25)	
Bias-adjusted	9.10	9.63	5.08	-11.55	-37.72	-46.82	0.00
	(1.48)	(1.16)	(0.40)	(-0.64)	(-1.66)	(-2.75)	
Biases							
DMR	2.97	4.95	11.65	20.03	22.88	19.92	
IMR	0.00	0.00	0.00	0.00	0.00	0.00	
CEIV	39.75	-20.59	-11.65	-2.81	15.65	-24.09	
SS	-30.80	-6.66	-0.01	-1.11	-18.16	12.64	
Total	11.91	-22.30	0.00	16.11	20.37	8.46	
B: Index							
	Low	2	3	4	High	H-L	
Median β_σ	0.26	1.66	6.41	17.88	38.39	38.13	
Median β_S	9.26	15.10	24.33	37.58	53.26	44.01	
Median relative spread	0.81	2.11	5.35	12.55	22.69	21.88	
Median moneyness	-1.79	-0.90	-0.02	0.87	1.74	3.53	
	Average daily call option returns						Wolak p-value
True prices	18.53	23.71	18.24	-13.76	-89.54	-108.07	0.00
	(2.03)	(1.64)	(0.80)	(-0.41)	(-1.91)	(-2.73)	
Bias-unadjusted	18.82	22.51	18.55	-6.24	-59.18	-78.00	0.02
	(2.07)	(1.56)	(0.81)	(-0.19)	(-1.26)	(-1.96)	
Bias-adjusted	18.37	23.70	18.42	-13.44	-86.82	-105.19	0.00
	(2.03)	(1.64)	(0.81)	(-0.40)	(-1.85)	(-2.64)	
Biases							
DMR	0.11	0.47	3.22	12.67	29.66	29.55	
IMR	0.00	0.00	0.00	0.00	0.00	0.00	
CEIV	1.88	-1.66	-2.90	-5.15	8.85	6.96	
SS	-1.70	-0.01	0.00	0.00	-8.14	-6.44	
Total	0.29	-1.20	0.31	7.52	30.37	30.07	

Table 9: **FM regressions of simulated delta-hedged option returns on their β with respect to volatility (β_σ).** This table displays the results of FM regressions: $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t} = \lambda_0 + \lambda_\sigma \times \beta^f_{\sigma,i,t-1} + \eta_{i,t}$. The dependent variable is delta-hedged option excess return ($R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$), where $R_{f,i,t}$ is the excess return of call or put i between $t-1$ and t; $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$ is the β of the call or put with respect to the underlying asset; and $R_{S,t}$ is the excess return of the underlying stock or index between $t-1$ and t. "True prices" refer to results calculated with simulated prices without measurement errors. The bias unadjusted and adjusted results are based on simulated prices with added measurement errors. In the bias-unadjusted method, the independent variable is the β of the option with respect to the volatility of the underlying at time $t-1$ ($\beta^f_{\sigma,i,t-1}$). Moreover, the sample used in the bias-unadjusted method is constructed with selection criteria based on $t-1$, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, in the bias-adjusted method, the independent variable is $\beta^f_{\sigma,i,t-2}$, and the regressions are estimated with WLS using gross returns as weights (either $C_{i,t-1}/C_{i,t-2}$ or $P_{i,t-1}/P_{i,t-2}$). The total bias is decomposed into its different parts: the direct and indirect regression coefficient (DRC and IRC, respectively), CEIV bias, and sample-selection bias (SS). Section 4.1 describes the simulation procedure. The regression coefficients and t-statistics (in parentheses) are averages across simulation trials.

	A: Calls			
	Stocks		Index	
	λ_0	λ_σ	λ_0	λ_σ
True prices	-1.00	-5.96	0.93	-5.43
	(-4.07)	(-9.10)	(1.04)	(-7.17)
Bias-unadjusted	-15.70	-1.71	-5.39	-4.32
	(-52.89)	(-2.60)	(-5.21)	(-5.67)
Bias-adjusted	-0.95	-5.87	0.53	-5.38
	(-3.10)	(-9.05)	(0.46)	(-6.90)
	Biases in unadjusted estimates of λ_0 and λ_σ			
DRC	5.38	1.61	-0.68	0.79
IRC	0.35	-0.04	-0.18	0.01
CEIV	-9.30	2.37	-5.62	0.55
SS	-11.13	0.30	0.15	-0.23
Total	-14.70	4.25	-6.32	1.11
	B: Puts			
	Stocks		Index	
	λ_0	λ_σ	λ_0	λ_σ
True prices	-1.14	-5.51	-1.35	-4.62
	(-7.18)	(-9.23)	(-1.67)	(-7.02)
Bias-unadjusted	-25.10	1.35	-8.93	-3.53
	(-104.36)	(2.25)	(-9.60)	(-5.31)
Bias-adjusted	-1.13	-5.44	-1.53	-4.65
	(-5.01)	(-9.17)	(-1.52)	(-6.85)
	Biases in unadjusted estimates of λ_0 and λ_σ			
DRC	4.33	2.05	-3.00	0.83
IRC	-0.73	0.11	0.01	0.00
CEIV	-7.15	2.55	-6.75	0.60
SS	-20.41	2.15	2.17	-0.34
Total	-23.96	6.86	-7.58	1.09

Table 10: **Mean excess returns of simulated delta-hedged calls and puts on individual stocks.** This table presents averages of simulated delta-hedged returns of calls and puts sorted by their βs with respect to the volatility of the underlying asset (β_σ). Delta-hedged option excess returns are calculated as $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$, where f_{t-1} is the price of the call or put i at time $t-1$; $R_{f,i,t}$ is its excess return between $t-1$ and t; $R_{S,t}$ is the excess return of the underlying stock; and $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$. "True prices" refer to results calculated with simulated prices without measurement errors. The bias unadjusted and adjusted results are based on simulated prices with added measurement errors. The sample used in the bias-unadjusted method is constructed with selection criteria based on $t-1$ variables, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. The sorting variable is $\beta^f_{\sigma,i,t-1}$ in the bias-unadjusted method. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, the bias-adjusted method uses $\beta^f_{\sigma,i,t-2}$ as the sorting variable, and average returns are calculated with weighted averages using gross returns as weights. The statistics on β_σ, relative spreads, and moneyness are the mean values across time of the median $t-2$ values of these variables. The total bias is decomposed into its different parts: the direct and indirect mean return biases (DMR and IMR, respectively), CEIV bias, and sample-selection bias (SS). Section 4.1 describes the simulation procedure. The reported means and t-statistics (in parentheses) are the averages across simulation trials. Relative spreads are displayed in percentages. Returns and biases are in basis points per day.

	A: Calls						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.24	0.88	2.40	5.38	10.81	10.56	2.40
Median relative spread	4.33	6.34	9.61	14.12	16.98	12.65	9.08
Median moneyness	-1.41	-0.74	-0.07	0.62	1.37	2.77	-0.07
	Average daily delta-hedged returns						
True prices	-1.29	-5.03	-14.41	-33.63	-67.74	-66.44	-24.42
	(-9.66)	(-9.61)	(-9.57)	(-9.56)	(-9.22)	(-9.20)	(-9.23)
Bias-unadjusted	-16.98	-21.74	-16.80	-20.25	-37.77	-20.80	-22.71
	(-80.35)	(-36.01)	(-10.47)	(-5.57)	(-5.09)	(-2.86)	(-9.01)
Bias-adjusted	-1.45	-5.49	-15.12	-34.31	-67.73	-66.29	-24.69
	(-7.56)	(-9.16)	(-9.38)	(-9.46)	(-9.17)	(-9.15)	(-9.11)
	Biases in unadjusted means						
DMR	3.04	5.15	10.30	17.66	22.31	19.27	11.69
IMR	0.89	0.07	0.02	0.03	0.03	-0.86	0.21
CEIV	20.86	-23.85	-12.92	-3.90	19.82	-1.04	0.00
SS	-40.47	1.93	0.20	-0.41	-12.20	28.27	-10.19
Total	-15.68	-16.70	-2.40	13.38	29.96	45.65	1.71
	B: Puts						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.16	0.54	1.57	3.83	8.42	8.26	1.57
Median relative spread	3.55	5.54	7.91	12.48	16.76	13.20	8.01
Median moneyness	1.65	1.05	0.39	-0.33	-1.16	-2.81	0.39
	Average daily delta-hedged returns						
True prices	-0.73	-2.85	-8.78	-22.16	-48.43	-47.70	-16.59
	(-9.39)	(-9.35)	(-9.40)	(-9.47)	(-9.30)	(-9.29)	(-10.40)
Bias-unadjusted	-19.14	-36.91	-18.44	-15.09	-16.28	2.86	-21.17
	(-113.84)	(-80.74)	(-17.10)	(-6.00)	(-3.05)	(0.54)	(-9.82)
Bias-adjusted	-1.05	-3.47	-9.85	-23.53	-48.95	-47.91	-17.31
	(-6.32)	(-8.66)	(-9.12)	(-9.36)	(-9.24)	(-9.23)	(-9.96)
	Biases in unadjusted means						
DMR	2.76	4.04	7.40	14.63	21.61	18.85	10.09
IMR	-1.91	-0.33	0.06	0.06	0.06	1.97	-0.41
CEIV	51.07	-42.33	-18.77	-7.87	17.90	-33.17	0.00
SS	-70.32	4.56	1.64	0.25	-7.41	62.90	-14.25
Total	-18.40	-34.07	-9.66	7.07	32.15	50.55	-4.58

Table 11: **Mean excess returns of simulated delta-hedged calls and puts on the index.** This table presents averages of simulated delta-hedged returns of calls and puts sorted by their β with respect to the volatility of the underlying index (β_σ). Delta-hedged option excess returns are calculated as $R_{f,i,t} - \beta^f_{S,i,t-1} R_{S,t}$, where f_{t-1} is the price of the call or put i at time $t-1$; $R_{f,i,t}$ is its excess return between $t-1$ and t; $R_{S,t}$ is the excess return of the underlying index; and $\beta^f_{S,i,t-1} = S_{t-1}/f_{t-1} \times \Delta^f_{S,i,t-1}$. "True prices" refer to results calculated with simulated prices without measurement errors. The bias unadjusted and adjusted results are based on simulated prices with added measurement errors. The sample used in the bias-unadjusted method is constructed with selection criteria based on $t-1$ variables, including the requirement that the option's delta ($\Delta^f_{S,i,t-1}$) is non-missing. The sorting variable is $\beta^f_{\sigma,i,t-1}$ in the bias-unadjusted method. In contrast, the sample used in the bias-adjusted method is constructed with selection criteria based on $t-2$ variables, with missing deltas replaced by their lagged values ($\Delta^f_{S,i,t-2}$) when calculating the delta-hedged return. In addition, the bias-adjusted method uses $\beta^f_{\sigma,i,t-2}$ as the sorting variable, and average returns are calculated with weighted averages using gross returns as weights. The statistics on β_σ, relative spreads, and moneyness are the mean values across time of the median $t-2$ values of these variables. The total bias is decomposed into its different parts: the direct and indirect mean return biases (DMR and IMR, respectively), CEIV bias, and sample-selection bias (SS). The simulation parameters are calibrated to match moments of the S&P 500 returns and bid-ask spreads. Section 4.1 describes the simulation procedure. The reported means and t-statistics (in parentheses) are the averages across simulation trials. Relative spreads are displayed in percentages. Returns and biases are in basis points per day.

	A: Calls						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.20	1.04	4.19	13.23	34.59	34.39	4.19
Median relative spread	0.68	1.61	3.77	9.33	20.99	20.30	3.82
Median moneyness	-1.88	-1.14	-0.32	0.57	1.59	3.47	-0.32
	Average daily delta-hedged returns						
True prices	-0.93	-5.14	-21.39	-68.78	-184.13	-183.20	-56.07
	(-6.81)	(-6.73)	(-6.82)	(-6.85)	(-6.91)	(-6.90)	(-6.73)
Bias-unadjusted	-3.52	-7.43	-23.35	-68.93	-153.97	-150.46	-51.43
	(-19.99)	(-9.19)	(-7.25)	(-6.77)	(-5.73)	(-5.62)	(-6.67)
Bias-adjusted	-0.91	-5.25	-21.74	-69.91	-183.55	-182.63	-54.95
	(-5.04)	(-6.36)	(-6.62)	(-6.74)	(-6.74)	(-6.73)	(-6.84)
	Biases in unadjusted means						
DMR	0.69	0.28	1.47	7.66	26.90	26.20	7.40
IMR	-0.55	-0.02	0.01	0.01	0.01	0.56	-0.11
CEIV	2.49	-1.95	-3.12	-7.35	9.93	7.44	0.00
SS	-5.22	-0.60	-0.33	-0.47	-6.68	-1.46	-2.66
Total	-2.58	-2.29	-1.97	-0.15	30.15	32.74	4.63
	B: Puts						
	Low β_σ	2	3	4	High β_σ	H-L	All
Median β_σ	0.32	2.24	7.72	17.18	30.52	30.20	7.72
Median relative spread	0.00	0.14	2.16	8.30	19.98	19.98	2.31
Median Moneyness	1.83	0.82	-0.12	-0.98	-1.79	-3.62	-0.12
	Average daily delta-hedged returns						
True prices	-1.89	-11.47	-36.00	-76.54	-135.70	-133.82	-52.32
	(-6.58)	(-6.56)	(-6.68)	(-6.84)	(-6.89)	(-6.88)	(-6.64)
Bias-unadjusted	-1.86	-11.87	-36.83	-81.53	-106.11	-104.25	-47.64
	(-6.10)	(-6.75)	(-6.83)	(-7.27)	(-5.33)	(-5.30)	(-6.62)
Bias-adjusted	-1.87	-11.60	-36.50	-77.46	-135.44	-133.57	-51.28
	(-5.73)	(-6.40)	(-6.61)	(-6.75)	(-6.70)	(-6.69)	(-6.74)
	Biases in unadjusted means						
DMR	0.01	0.20	0.70	6.15	24.65	24.64	6.34
IMR	0.02	0.01	0.00	0.00	-0.01	-0.02	0.00
CEIV	0.30	-0.48	-1.60	-11.21	12.99	12.69	0.00
SS	-0.30	-0.13	0.07	0.07	-8.05	-7.74	-1.67
Total	0.02	-0.40	-0.83	-5.00	29.59	29.57	4.68

Figure 1: **Cumulative distribution of dollar open interest.** Dollar open interest ($\$OI$) is the closing open interest of an option multiplied by its closing quote midpoint. For every day in our sample, we sort options according to their $\$OI$ percentile and compute the $\$OI$ accumulated up to that percentile divided by the total $\$OI$ on that day. The graph shows the averages of these cumulative distributions across all days in the sample. The figure shows that the options on the top quintile of $\$OI$ account for approximately 90% of the total $\$OI$.

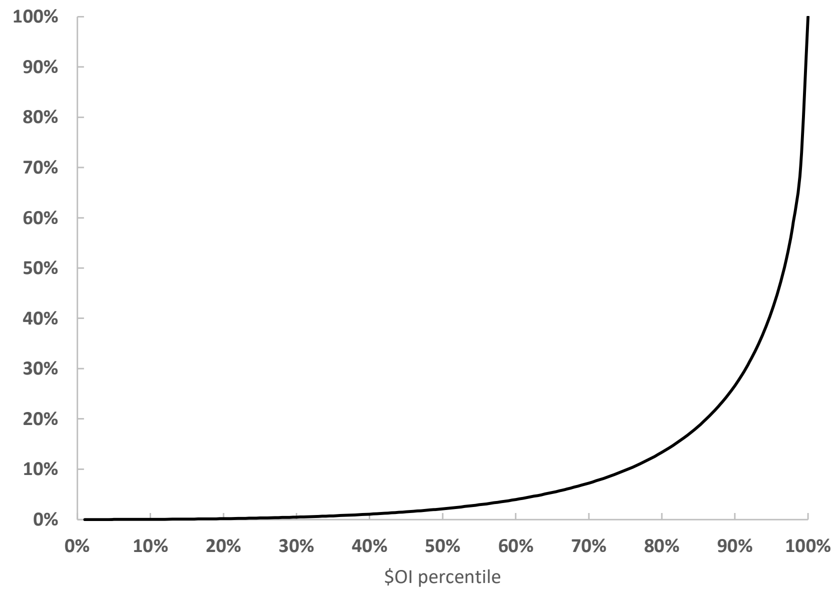

Figure 2: **Average closing bid-ask spreads of options.** Panels A and B show the average across options of the closing relative bid-ask spreads of options written on stocks in the S&P 500. Panels C and D show the average across options of the closing relative bid-ask spreads of options written on the S&P 500 Index. Panels A and C show the results for the baseline sample while panels B and D show the results for ATM options with 21 to 23 trading days to maturity. We classify all options with moneyness between -0.5 and 0.5 as ATM options. Both equal-weighted and dollar open-interest ($\$OI$) weighted averages are presented. Option spreads are in percentage terms relative to quote midpoints.

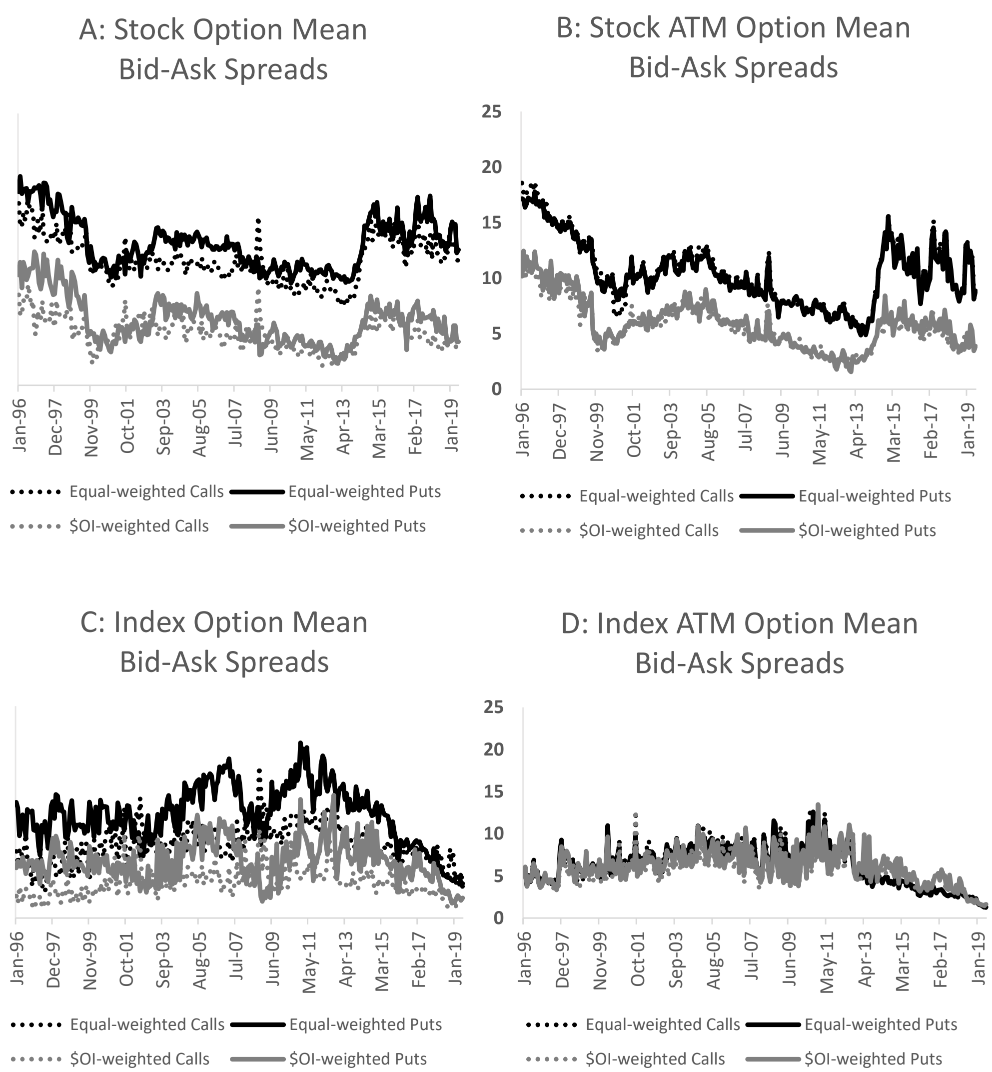

Figure 3: **Option β_S and β_σ as a function of moneyness.** The graphs below display the averages in our sample of the option βs with respect to the underlying (β_S) and with respect to volatility (β_σ) as a function of moneyness. The moneyness of an option (with maturity at time T) at time t is $\ln(e^{-r_t(T-t)}K/S_t)/(\sigma_t\sqrt{T-t})$, where S_t is the stock price; K is the option strike price; r_t is the risk-free rate; and σ_t is the volatility implied by the simulated option price. The option βs are calculated with the option Δs and νs in IvyDB. The β of an option with respect to the underlying (β_S^f) is $\Delta_t^f \times S_t/f_t$, where Δ_t^f and f_t are the delta of the option and its price, respectively, both given by IvyDB. The β of an option with respect to the volatility of the underlying (β_σ^f) is defined as ν_t^f/f_t, where ν_t^f is the vega of the option, also given by IvyDB.

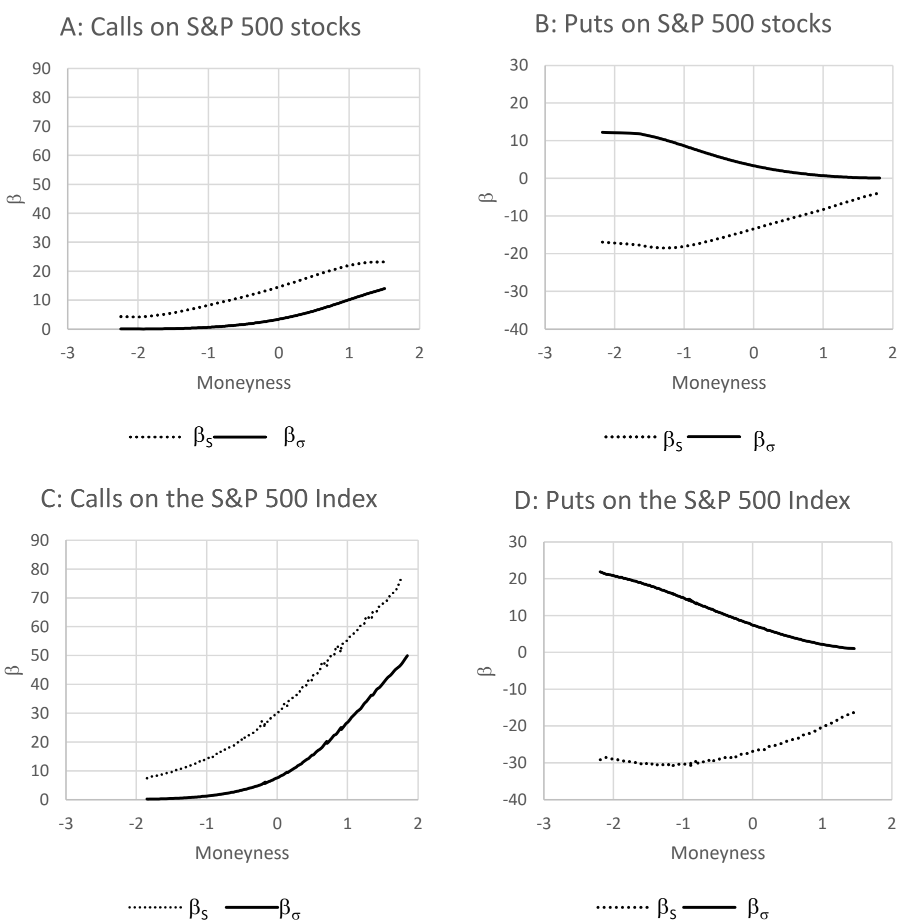

Figure 4: **Average volatility smile used to estimate the model-free implied variance ($MFIV$).** The graph shows the averages in our sample of the implied volatilities of out-of-the-money options used to estimate the $MFIV$s in Table 6. The circles (squares) are average implied volatilities of out-of-the-money puts (calls) from IvyDB. The lines are the average interpolated implied volatilities used to numerically calculate the $MFIV$ by an integral of a function of option prices on strike prices. The continuous line represents the volatilities used in Driessen, Maenhout, and Vilkov (2009). This line keeps implied volatility constant for OTM put options with $\Delta^P > -0.05$ and OTM calls with $\Delta^C < 0.15$. The dotted line represents the implied volatilities used to estimate the $MFIV$ in the Extended OTM method in Table 6. The implied volatilities are plotted as a function of the ratio of strike-to-spot price (K/S).