

Do Call Prices and the Underlying Stock Always Move in the Same Direction?

Gurdip Bakshi

University of Maryland

Charles Cao

Pennsylvania State University

Zhiwu Chen

Yale University

This article empirically analyzes some properties shared by all one-dimensional diffusion option models. Using S&P 500 options, we find that sampled intraday (or interday) call (put) prices often go down (up) even as the underlying price goes up, and call and put prices often increase, or decrease, together. Our results are valid after controlling for time decay and market microstructure effects. Therefore one-dimensional diffusion option models cannot be completely consistent with observed option price dynamics; options are not redundant securities, nor ideal hedging instruments—puts and the underlying asset prices may go down together.

Much of the extant knowledge about option pricing is based on the assumption that the underlying asset price follows a one-dimensional diffusion process. Examples of such option pricing models include the classic Black–Scholes (1973), Merton (1973), the Cox–Ross (1976) constant elasticity of variance, the ones studied in Derman and Kani (1994), Rubinstein (1994), Bergman, Grundy, and Wiener (1996), Bakshi, Cao, and Chen (1997, 2000), and Dumas, Fleming, and Whaley (1998). All models in the one-dimensional diffusion class share three basic properties. First, call prices are monotonically increasing and put prices are monotonically decreasing in the underlying asset price (*the monotonicity property*). Second, as the underlying asset price is the sole source of uncertainty for all of its options, option prices must be perfectly correlated with each other and with the underlying asset (*the perfect*

We would like to thank Kerry Back, Cliff Ball, David Bates, Steve Buser, Tarun Chordia, Alex David, Ian Domowitz, Phil Dybvig, Mark Fisher, Eric Ghysels, John Griffin, Frank Hatheway, Bill Kracaw, Craig Lewis, Dilip Madan, Jim Miles, Shashi Murthy, Ron Masulis, Louis Scott, Lemma Senbet, Hans Stoll, Rene Stulz, Guofu Zhou, and especially the anonymous referee, Bernard Dumas, and Ravi Jagannathan for their comments. This article has benefited from comments by seminar participants at the Federal Reserve Board, Pennsylvania State University, Seventh Annual Conference on Financial Economics and Accounting at SUNY-Buffalo, Vanderbilt University, and the 1998 Western Finance Association meetings. Any remaining errors are our responsibility alone. Address correspondence to Charles Cao, Smeal College of Business, Pennsylvania State University, University Park, PA 16802, or e-mail: charles@loki.smeal.psu.edu.

correlation property). Third, options can be replicated using the underlying and a risk-free asset, and are hence redundant securities (*the option redundancy property*). These model predictions have been the foundation of standard options textbooks. But are they consistent with observed option-price dynamics? From the perspective of pricing, hedging, and/or model internal consistency, many existing studies have examined the empirical performance of the Black–Scholes and other members in the one-dimensional diffusion class [see, e.g., Rubinstein (1985, 1994), Bakshi, Cao, and Chen (1997, 2000), Dumas, Fleming, and Whaley (1998), and Bates (2000)]. However, none has focused directly on the above model predictions. To be exact, in our inquiry, we address an important empirical question: Do call prices and the underlying stock always move in the same direction, and do put prices and the underlying stock always move in the opposite direction? If they do not, how often does it occur? In addition, if these predictions are violated, to what extent are they related to market microstructure factors and option time decay. Are additional state variables required to characterize empirical option pricing dynamics? This article serves to fill each one of these gaps.

Specifically, for our study we use bid-ask midpoint prices of S&P 500 index options, sampled at various intraday intervals (e.g., every half hour, 1 hour, 2 hours, and so on). The S&P 500 option market is one of the most active, and this index is also the basis for the most actively traded equity index futures contract. Furthermore, we focus on intraday sampling intervals, as they help minimize the impact of time decay in option premium on our results. This consideration is important because, unless one assumes a parameterized option pricing formula, it is not possible to decompose an option price change in a given period into a time-decay and a non-time-decay component. Overall, our findings can be summarized as follows.

First, depending on the intraday sampling interval, bid-ask midpoint call prices move in the opposite direction with the underlying asset between 7.2% and 16.3% of the time. We refer to such violations as type I violations. This is true whether the spot S&P 500 index or the lead-month S&P 500 futures are used as a proxy for the underlying asset. Thus this type of violation of the model predictions cannot be a consequence of stale S&P 500 component stock prices. When put option prices are used in place of call prices, similar violation rates are documented. Second, when the sampling interval changes from intraday to interday, the occurrence rate actually decreases, suggesting a role played by time decay in option premium. Third, the violation occurrence rate differs across options' maturity. Of a given moneyness, long-term calls are the most likely to move in the opposite direction, followed by medium-term and short-term calls. In general, there is no clear association between the moneyness of the option and its tendency to move in the opposite direction to the underlying stock. Finally, call and put prices with the same strike and the same expiration often move in the same direction, regardless of the changes in the underlying and irrespective of the intraday sampling interval. In fact,

when prices are sampled every 3 hours, call and put prices go up or down together as often as 16.9% of the time. Whether the underlying asset goes up or down, it is more likely for the call and put prices to go down together than up together. These observed option price movements are contrary to standard textbook predictions. Most of these occurrences cannot be treated as “outliers” since one cannot imagine throwing away as much as 17% of the observations. As these price quotes are usually binding at least up to 10 contracts, neither can they be treated as insignificantly misquoted prices.

Our empirical exercise also documents those occurrences in which the underlying asset price has changed during a given interval, but the option price quote has not (type II violations). Such occurrences are between 3.5% and 35.6%. Our analysis points to a class of violations in which the call/put prices have changed even though the spot index has not (type III violations), and the market makers overadjust option quotes in response to a change in the underlying stock price (type IV violations). The frequency of the latter occurrence is as much as 11.7% for calls and 13.7% for puts. But our study shows that type II and type IV violations are essentially due to market microstructure-related effects and minimum tick size restrictions, and type III violation occurrence frequency is relatively rare.

In our quest to comprehend observed option price dynamics, we focus our efforts on four candidates of interest: (1) market microstructure factors, (2) the violations of put-call parity, (3) the impact of time decay, and (4) a two-factor stochastic process for the underlying stock price. Maintaining a single-factor setting, our analysis reveals that market microstructure factors are important for some type of violations. For instance, the type II (type IV) violation occurrence rate is monotonically declining (increasing) in the dollar bid-ask spread. However, there is no association between type I violation rate and each of the market microstructure factors. When we adopt deviation from one round of transaction costs as a benchmark, about 3% of the intraday option sample violates put-call parity. Nonetheless, our results establish that type I violation frequency is robust to the inclusion or the exclusion of such observations. In evaluating the impact of time decay, we notice that the magnitude of time decay is comparatively larger for interday samples, but negligible for intraday samples. As a consequence, time decay explains more of type I violations for daily samples, but less for intraday samples (the samples stressed in our work). Each empirical result holds across calendar time, and is stable under alternative test designs.

In summary, it is evident from our study that some contradictory option-price movements are attributable to market microstructure factors and time decay. Still, it is difficult to explain why some types of violations occur in the first place. Regardless of the minimum tick size or bid-ask spread, market makers can at the minimum choose not to move call prices (bid or ask) up, for example, when the underlying price is going down. We are then led to reexamine model specifications beyond which much of the

extant option pricing knowledge is based. To search for an option pricing model that can explain the documented price dynamics (beyond what can be accounted for by market microstructure factors), one may need to incorporate, besides the underlying asset price, additional state variables. In this direction, we stipulate that if one is to introduce another state variable that affects option prices, this second stochastic process should not be perfectly correlated with the underlying. Furthermore, since call prices often go down (up) when the underlying goes up (down), this state variable should either affect call prices differently than the underlying, or be negatively correlated with the underlying price. Among the known ones, the stochastic volatility (SV) model of Heston (1993) possesses such features. For this reason we investigate the extent to which the SV model can explain the documented option price dynamics. Our simulation results show that about 11% of call prices generated from the SV model move in the opposite direction with the underlying price. On the other hand, the regression results indicate the SV model quantitatively falls short of completely explaining the observed option price movements. Using implied volatilities and SV model parameters, we find that about 47% of the type I violations become consistent with the predictions of the SV model. This analysis also uncovers the finding that type II and type IV violations are mostly outside the scope of the SV model. Thus a more plausible story of violation patterns should not leave out the role of market microstructure and time decay-related factors.

Our empirical results have important implications for investment management practice. They suggest that certain standard hedging strategies might not perform as well as one might expect. For instance, as students learn from textbooks, a typical hedge for a stock position involves shorting calls (or, buying puts), with the understanding that the call and the stock will move up or down in tandem. But as often as 7% of the time call prices go up even when the underlying goes down, a conventional hedge may actually double, rather than reduce, the hedger's loss. Another lesson from the textbooks is that when applying such models as the Black-Scholes to create a dynamic hedge (e.g., portfolio insurance), one should revise the hedge as often as the market condition changes. The reasoning is that, absent market frictions, hedging errors should converge to zero as the hedge revision interval shrinks to zero. Given our evidence that call (put) prices often move in the opposite (same) direction with the underlying, however, it is likely that beyond a certain point a higher frequency of hedge rebalancing will actually lead to higher hedging errors. Using the Black-Scholes formula as an example, we show that this is indeed the case: as revision takes place more frequently, the hedging errors decrease initially but increase after a certain point. This is true even without taking transaction costs into account.

The rest of the article is organized as follows. Section 1 develops the theoretical implications of one-dimensional diffusion option pricing models. In Section 2, we describe the S&P 500 option and futures data. Section 3

presents the main empirical results. Section 4 sheds light on a stochastic volatility option pricing model. In Section 5, we discuss the robustness of our findings. Concluding remarks are offered in Section 6.

1. Properties of Option Prices in a Diffusion Setting

In this section we discuss several properties possessed by all option pricing models in which the underlying asset price follows a one-dimensional diffusion (and hence Markov) process. We refer to each such model as a *one-dimensional diffusion (option pricing) model*.

1.1 Basic properties

The problem at hand is to determine the price of a European call with strike price K and τ years to expiration, written on some non-dividend-paying asset whose time t price is denoted by $S(t)$. To solve this problem, we need to specify (i) the process followed by $S(t)$ and (ii) the valuation rule of the economy, granted that $\{S(t) : t \geq 0\}$ is a well-defined stochastic process on some probability space. Assume that $S(t)$ follows a one-dimensional diffusion:

$$dS(t) = \mu[t, S]dt + \sigma[t, S]S(t)dW(t), \quad t \geq 0, \quad (1)$$

with $S(0) > 0$, where the drift, $\mu[t, S]$, and the volatility, $\sigma[t, S]$, are both functions of at most t and $S(t)$ and satisfy the usual regularity conditions, and $W(t)$ is a standard Brownian motion. This process contains many of those assumed in existing option pricing models. For example, in Black and Scholes (1973), $\sigma[t, S] = \sigma$, for some constant σ ; in the Cox and Ross (1976) constant elasticity of variance model, $\sigma[t, S] = \sigma S(t)^\alpha$, for some constants σ and α ; and in the models empirically investigated by Dumas, Fleming, and Whaley (1998), $\sigma[t, S]$ is the sum of polynomials of K and τ . The class of models covered by Equation (1) is also the focus of Bergman, Grundy, and Wiener (1996). For the valuation rule of the economy, assume, as is standard in the literature, that interest rates are constant over time and that the financial markets admit no free lunches.¹ Then there exists an equivalent martingale measure with respect to which any asset price today equals the expected risk-free discounted value of its future payoff. For example, letting $C(t, \tau, K)$ denote the time t price of the call option under consideration, we have

$$C(t, \tau, K) = E^*(e^{-r\tau} \max\{S(t + \tau) - K, 0\}) \quad (2)$$

where the expectation operator, $E^*(\cdot)$, is with respect to a given equivalent martingale measure, and r is the constant spot interest rate. Note that under

¹ The interest rates can be stochastic as in Bakshi, Cao, and Chen (1997) and others. But as found by Bakshi, Cao, and Chen (1997), stochastic interest rates may not be that important for pricing or hedging options. Given our focus on intraday option price changes, the impact of stochastic interest rates should also be negligible.

the equivalent martingale measure, the underlying asset price obeys the following stochastic differential equation:

$$\frac{dS(t)}{S(t)} = rdt + \sigma[t, S]dW(t), \quad (3)$$

that is, the expected rate of return on the asset is, under the martingale measure, the same as the risk-free rate r . By Ito's lemma, the call price dynamics are determined as follows:

$$dC = \left\{ C_t + \frac{1}{2} \sigma^2[t, S] S^2 C_{SS} \right\} dt + C_S dS, \quad (4)$$

where the subscripts on C stand for the respective partial derivatives. The result below is adopted from Bergman, Grundy, and Wiener (1996) (a proof is provided therein).

Proposition 1. *Let the underlying price $S(t)$ follow a one-dimensional diffusion as described in Equation (1). Then the option delta of any European call written on the asset must always be nonnegative and bounded from above by one:*

$$0 \leq C_S \leq 1. \quad (5)$$

The delta of any European put, denoted by P_S , must be nonpositive and bounded below by -1 :

$$-1 \leq P_S \leq 0. \quad (6)$$

That is, provided that everything else is fixed, the value of a European call (put) should be nondecreasing (nonincreasing) in the underlying asset price.

We refer to the above as the *monotonicity property*. Before discussing its implications further, we note that in a one-dimensional diffusion model the sole source of stochastic variations for all options is the underlying asset, and hence all option prices, regardless of moneyness and maturity, must covary perfectly with each other and with the underlying asset. This *perfect correlation property* imposes a simple, but potentially stringent, restriction on option price dynamics. It also implies another property of one-dimensional diffusion models, that is, *option contracts are redundant securities*, in the sense that they can be exactly replicated by dynamically mixing the underlying asset with a risk-free bond.

1.2 Testable Predictions

The monotonicity and the perfect correlation properties are both directly testable using properly sampled option data. As these properties are shared by all one-dimensional diffusion option pricing models, any rejection applies to the entire class. To formulate the testable predictions more directly, note that in our empirical exercises to follow we measure contemporaneous price changes in the underlying asset and its options by using sampling intervals ranging from 30 minutes to 1 day. That is, if we use Δt to denote the time length of the sampling interval, the highest value for Δt is a day and the lowest is 30 minutes. Under the assumption that $\{C_t + \frac{1}{2}\sigma^2[t, S]S^2C_{SS}\} \Delta t$ in Equation (4) is small and negligible for small Δt (this assumption will be empirically justified in Section 3.5), intraday changes in an option price are mostly in response to contemporaneous changes in the underlying market. This together with Proposition 1 leads to the following predictions:

1. Over any intraday interval, price changes in the underlying asset and in any call written on it should share the same sign: $\Delta S \Delta C \geq 0$, where ΔC denotes changes in the call price.
2. Over any intraday interval, price changes in the underlying asset and in any put on it should have opposite signs: $\Delta S \Delta P \leq 0$, where ΔP denotes changes in the put price.
3. Over any intraday interval, $\Delta C / \Delta S \leq 1$ and $\Delta P / \Delta S \geq -1$, provided $\Delta S \neq 0$.
4. Over any intraday interval, contemporaneous price changes in call and put options with the same strike price and the same maturity should be of opposite signs: $\Delta C \Delta P \leq 0$.

These predictions mainly examine the left-hand and right-hand inequalities in Equations (5) and (6). In order to derive exact relationships between changes in a call (or a put) and the underlying asset price, one would need to parameterize the underlying price process and the valuation framework in further detail, which is not the main purpose of the present article. Instead, our focus is on the empirical validity of the above model-independent predictions.

2. Intraday Index Option and Futures Data

The dataset employed in this study includes all intraday observations on (i) the S&P 500 spot index, (ii) lead-month S&P 500 futures prices, and (iii) bid-ask midpoint prices for S&P 500 index options. The use of bid-ask midpoint prices for options is to help eliminate the impact of bid-ask bounces. The sample period is from March 1, 1994, to August 31, 1994, with a total of 3.8 million observations on the index calls and puts. The intraday S&P 500 cash index and lead-month futures prices are obtained from the Chicago Mercantile Exchange. The source for the index options is the Berkeley Option Database. All intraday prices are time-stamped to the

second. We include intraday S&P 500 lead-month futures prices to control for the fact that index option prices usually reflect not only information contained in the spot index, but also innovations in the futures market.

Three filtering criteria are applied to the original data. First, as the New York Stock Exchange (NYSE) closes 15 minutes ahead of the options and futures markets, all prices time-stamped later than 3:00 P.M. Central Standard Time (CST) are eliminated. Second, index options with less than 6 days to expiration and with quoted prices lower than $\$ \frac{3}{8}$ are omitted to alleviate expiration-related and price discreteness-related biases. Third, option contracts with less than 10 quote revisions during any given day are dropped from that day's sample.

Option and spot price subsamples are collected using a sampling interval of 30 minutes, 1 hour, 2 hours, 3 hours, and 1 day (the last quote of each day prior to 3:00 P.M. CST). Other than these five call samples and five put samples, we tried 5-minute and 10-minute subsamples and found the results to be similar.

By convention, a call option is said to be *at the money* (ATM) if $S/K \in (0.97, 1.03)$, *out of the money* (OTM) if $S/K \leq 0.97$, and *in the money* (ITM) if $S/K \geq 1.03$. Similar terminology is defined for puts by replacing S/K with K/S . An option is said to be *short term* if it has less than 2 months to expiration, *medium term* if it has between 2 and 6 months to expiration, and *long term* otherwise.

To save space, we report summary statistics in Table 1 for the hourly sample for puts and calls, including (i) the average bid-ask midpoint price, (ii) the average bid-ask spread, (iii) the average percentage spread (the bid-ask spread divided by the bid-ask midpoint), and (iv) the total number of observations. Short-term and medium-term calls (puts) account for 43% and 37% (44% and 39%), respectively, of the entire hourly call (put) sample. As expected, OTM options have relatively wider percentage bid-ask spreads than their ATM or ITM counterparts.

Based on the daily option sample, Table 2 shows the number of quote revisions and trading volume across option moneyness and maturity categories. In each given moneyness category, short-term options are the most actively traded, followed by medium-term options; in each maturity category, OTM options have the highest trading volume, followed by ATM options. For example, in the short-term category there are on average 1080 OTM calls, 767 ATM calls, and 37 ITM calls traded per day. The only exception is short-term puts, for which ATM puts are the most actively traded. The quote revision picture is, however, quite different: in a given moneyness category, long-term option prices are the most frequently updated, whereas in a given maturity class ITM option prices are the most frequently updated (except for short-term options for which class ATM options are most often updated). For example, the quotes are revised on average every 0.4 minutes for long-term ITM calls and every 13.4 minutes for short-term OTM calls.

Table 2
Trading characteristics of S&P 500 calls and puts

Moneyness		Calls Term-to-Expiration			Puts Term-to-Expiration		
		Short	Medium	Long	Short	Medium	Long
OTM	No. of quote revisions	29	24	285	43	42	242
	No. of contracts traded	1080	833	164	784	261	139
ATM	No. of quote revisions	98	154	784	123	78	398
	No. of contracts traded	767	210	34	1000	198	109
ITM	No. of quote revisions	43	704	886	40	902	683
	No. of contracts traded	37	3	2	31	15	54

Reported below for each option moneyness/maturity category are (i) the average number of quote revisions per option contract per day, and (ii) the average number of contracts traded per option per day (trading volume). The sample period extends from March 1, 1994, through August 31, 1994. OTM, ATM, and ITM stand for out of the money, at the money, and in the money, respectively. Short-, medium-, and long-term refer to options with less than 60 days, with 60–180 days, and with more than 180 days, respectively, to expiration.

Overall the more sensitive an option’s value to the underlying price movements, the more frequently updated the option price.

To explain these observed patterns, it is worthwhile to briefly describe how the S&P 500 options market is structured. On the Chicago Board Options Exchange (CBOE), there are designated market makers who are responsible for the continual implementation of an “auto-quote” computer program. While other market makers can always offer more competitive quotes, the bid and ask quotes generated by the auto-quote program can at the minimum serve as the last source of liquidity and are binding up to 10 contracts. For each option contract, the designated market maker is responsible for providing a volatility input into the auto-quote program, where the volatility input may differ across option moneyness and maturity and may change with market conditions. Once the volatility value is given, the computer program automatically updates the quotes as the underlying index changes. For long-term ITM options, their quotes are more frequently revised even though they are rarely traded, because they have a delta very close to one. On the other hand, short-term OTM options have a delta close to zero and hence are relatively insensitive to underlying price changes. Given the minimum tick size of \$ $\frac{1}{16}$ or \$ $\frac{1}{8}$, their quotes are therefore rarely adjusted even though they tend to be actively traded.

3. Violations of Model Predictions

Based on the predictions given in Section 1.2, we define four distinct types of violation by the family of one-dimensional diffusion models:

- **Type I violation:** $\Delta S \Delta C < 0$, that is, either $\Delta S > 0$ but $\Delta C < 0$, or $\Delta S < 0$ but $\Delta C > 0$. Likewise, for puts, either $\Delta S > 0$ but $\Delta P > 0$, or $\Delta S < 0$ but $\Delta P < 0$.
- **Type II violation:** $\Delta S \neq 0$ but $\Delta C = 0$. For puts, $\Delta S \neq 0$ but $\Delta P = 0$.

- **Type III violation:** $\Delta S = 0$ but $\Delta C \neq 0$. For puts, $\Delta S = 0$ but $\Delta P \neq 0$.
- **Type IV violation:** $\Delta C/\Delta S > 1$, $\Delta S \neq 0$. For puts, $\Delta P/\Delta S < -1$, $\Delta S \neq 0$.

In the discussions to follow, we first present an overall picture regarding the four types of violations. Then we proceed to examine the violations from several perspectives, including the magnitude of violation, market microstructure effects, the potential role of time decay, violations of put-call parity, and regularity of occurrence.

3.1 Overall picture of empirical option price dynamics

Table 3 reports the respective occurrence frequencies of type I–type IV violations, each as a percentage of total observations in a given sample. Several systematic patterns emerge from this table. First, when the underlying index goes up (down), quite frequently call prices go down (up) and put prices go up (down), a phenomenon *fundamentally inconsistent* with the monotonicity and the perfect correlation properties of one-dimensional diffusion models. This is true regardless of sampling frequency and whether the cash index or the futures price is used as a surrogate for the underlying asset. For example, when price changes are sampled every hour, call prices and the underlying asset move in opposite directions (i.e., type I violations) 13.9% of the time when the cash index is used as the spot asset and 11.9% of the time when the futures price is used; put prices move in the same direction with the underlying asset 13.4% or 11.9% of the time, depending on whether the cash index or futures are used as the underlying asset. The occurrence of type I violations is remarkably persistent across all the sampling frequencies. For instance, based on the cash index, the occurrence of type I violations increases from 11.6% of the time (at the 30-minute sampling frequency) to 16.3% (at the 3-hour sampling frequency). When sampled daily, type I violations still account for 9.1% of the observations.

Type II violations also occur frequently, as shown in Table 3. In these cases call or put prices do not change, even after the underlying asset price has changed. But for this type of violation, the occurrence rate for calls decreases monotonically with the sampling interval, going from 35.6% of the time (at the 30-minute interval) to as low as 3.6% of the time (at the daily frequency). This suggests two possible explanations for type II violations: (i) The spot index and the futures price both change rapidly, but option prices change only slowly. Put differently, the options market is intrinsically slower than the spot and futures markets in adjusting to new information. (ii) Some changes in the underlying asset price are too small to warrant a change in call and put option prices, especially given a nontrivial minimum tick size or bid-ask spread. We will examine this issue later.

Type III violations are rare for both puts and calls and at each sampling frequency. This may not come as a surprise because the S&P 500 index

Table 3
Violation occurrences by type and by sampling interval

			Violations by calls				
Sampling interval	Number of observations	Spot asset used	Type I (%)	Type II (%)	Type III (%)	Type IV (%)	Total (%)
30 minutes	51363	Cash index	11.6	35.6	0.5	10.9	58.6
		Index futures	9.3	34.3	1.6	7.3	52.5
1 hour	25680	Cash index	13.9	23.0	0.4	11.0	48.3
		Index futures	11.9	22.2	1.6	8.9	44.6
2 hours	12840	Cash index	13.4	11.8	0.5	11.1	36.8
		Index futures	11.8	11.4	1.9	9.8	34.9
3 hours	4280	Cash index	16.3	8.2	0.0	11.7	36.2
		Index futures	15.8	7.7	2.5	7.2	33.2
1 day	3587	Cash index	9.1	3.6	0.0	11.5	24.2
		Index futures	7.2	3.5	0.0	7.7	18.4
Violations by Puts							
30 minutes	86088	Cash index	11.4	33.1	0.5	12.8	57.8
		Index futures	9.9	32.0	1.6	9.8	53.3
1 hours	43044	Cash index	13.4	20.9	0.4	13.1	47.8
		Index futures	11.9	20.0	1.4	10.5	43.8
2 hours	21522	Cash index	13.0	9.9	0.5	13.7	37.1
		Index futures	11.5	9.3	2.0	10.2	33.0
3 hours	7174	Cash index	15.7	7.7	0.0	13.4	36.8
		Index futures	13.5	7.2	1.8	8.8	31.3
1 day	6321	Cash index	5.4	2.8	0.0	13.2	21.4
		Index futures	6.5	2.7	0.0	9.6	18.8

Reported are, respectively, type I, type II, type III, and type IV violation occurrence rates, each as a percentage of total observations at a given sampling interval:

Type I:	$\Delta S \cdot \Delta C < 0$,	$\Delta S \neq 0$, $\Delta C \neq 0$	(or, $\Delta S \cdot \Delta P > 0$, $\Delta S \neq 0$,	$\Delta P \neq 0$, for puts)
Type II:	$\Delta S \cdot \Delta C = 0$,	$\Delta S \neq 0$, $\Delta C = 0$	(or, $\Delta S \cdot \Delta P = 0$, $\Delta S \neq 0$,	$\Delta P = 0$, for puts)
Type III:	$\Delta S \cdot \Delta C = 0$,	$\Delta S = 0$, $\Delta C \neq 0$	(or, $\Delta S \cdot \Delta P = 0$, $\Delta S = 0$,	$\Delta P \neq 0$, for puts)
Type IV:	$\frac{\Delta C}{\Delta S} > 1$,	$\Delta S \neq 0$	(or, $\frac{\Delta P}{\Delta S} < -1$,	$\Delta S \neq 0$, for puts)

The call (or put) option samples are separately obtained by sampling price changes once every (i) 30 minutes, (ii) 1 hour, (iii) 2 hours, (iv) 3 hours, and (v) 1 day. The rows under "Cash index" are obtained by using the S&P 500 cash index, while those under "Index futures" by using the lead-month SSP 500 futures, as a stand-in for the underlying asset.

and its futures price rarely stay unchanged during 30-minute or longer time intervals. Thus type III violations are not significant.

Type IV violations occur as frequently as 11.0% of the time for calls and 13.1% for puts at the hourly interval. That is, market makers tend to overadjust option quotes. However, our investigation suggests that type IV violations are closely related to the tick size restriction. To appreciate this point, take deep ITM call options (with delta close to 1) as an illustration. When the underlying index goes up by \$0.10, the implied increase in the call price is just \$0.10. Then if the minimum tick size is $\$ \frac{1}{8}$, market makers can either keep the price quotes unchanged (which results in a type II violation) or bump the bid or ask up by \$0.125 (which then results in a type IV violation). The market makers will face this dilemma whenever the implied increase or decrease in option value is between k and $(k+1)$ times the minimum tick size,

for any integer k . Therefore, like type II violations, type IV violations are mostly due to minimum tick size.

The results in Table 3 are virtually insensitive to whether the cash index or the lead-month futures price is used as a stand-in for the unobservable underlying asset. Thus the reported violations are *not* due to the fact that some S&P 500 component stock prices are stale, or to the fact that market makers for the index options often update price quotes based on innovations that first occur on the S&P futures market [e.g., Fleming, Ostdiek, and Whaley (1996)].

To determine the sensitivity of the above-documented violations to any potential breakdowns in the put-call parity, we conduct another experiment. In this test, all matched put-call pairs (in the maturity and strike dimension) violating the put-call parity relationship, $|\Delta C(t, \tau; K) - \Delta S - \Delta P(t, \tau; K)| < \eta$, are excluded. Letting η represent one round of transaction costs (i.e., the sum total of bid-ask spread for the put-call pair), this accounts for roughly 3% of the intraday samples. See the recent work by Kamara and Miller (1995), who also find parity violations to be small. Next, the occurrence frequencies are recomputed with this new sample.² Based on the cash index and hourly sample, the type I–type IV violation frequencies are 13.3%, 22.2%, 0.3%, and 11.9%, respectively. They are close to results reported in Table 3. With other sampling choices, the results are essentially the same. As a consequence, eliminating put-call parity violations will not overturn our empirical findings.

By the internal working of the put-call parity, note that a type IV violation for calls amounts to a type I violation for puts, and the reverse holds as well. To assess whether such connections will distort our inferences, we again construct a matched sample of puts and calls. Then we count observations for which the call (put) is a type I, and the put (call) a type IV, violation. At the representative hourly interval, this frequency is 2.8% (2.6%) for the cash index and 2.6% (2.1%) for index futures. Thus when violation interactions due to put-call parity are neutralized in Table 3, one is still left with a significantly large 11.1% of type I violations for calls and 10.8% for puts. Most documented type I violations cannot be due to tick size or bid-ask spread. The reason is that when the underlying asset goes down in value, regardless of bid-ask spread or tick size, the market makers can choose not to change the bid or ask price for, say, a call, instead of marking its bid or ask price upward (unless the true option value is driven by more than just the underlying asset price). We will address the impact of time decay and other market microstructure-related issues shortly.

² In generating the matched put-call sample, the sample size diminished significantly. To get a sense of this reduction, if the hourly call (put) sample has 25,680 (43,044) observations, the corresponding matched put-call sample only has 19,338 observations (compare the sample sizes in Tables 3 and 4). Persuaded by this constraint, the original call (or put) sample is retained in our later empirical work unless stated otherwise.

The overall picture can be summarized as follows. The occurrence of type I violations is robust, persistent, and relatively stable across different sampling intervals. Type II violations occur quite often, especially when sampling takes place as frequently as every half hour. But the fact that the type II occurrence rate decreases with the sampling interval indicates that they are highly sensitive to tick size or bid-ask spread. With types I–IV added together, the overall occurrence rate is as high as 48.3% of the time for calls and 47.8% for puts, based on the 1-hour sampling intervals. At the daily sampling frequency, the overall occurrence rates are reduced to 24.2% for calls and 21.4% for puts.

The frequent occurrence of type I violations represents strong evidence against the predictions of all one-dimensional diffusion option pricing models. Call and put prices do not change monotonically with the underlying price. Given the quality of the intraday S&P index option data, one cannot discard these type I observations from the sample (they are 16.3% of the total). One may point out that in order to reflect different market conditions, the designated market maker at the options exchange generally changes the volatility inputs to the auto-quote computer program over time. Depending on whether and how the volatility input is changed, the resulting option price quotes may not move in tandem with the underlying index as dictated by one-dimensional diffusion models, therefore the documented patterns are simply consequences of how the price quotes are generated. However, from the perspective of both an outside observer and an option pricing model developer, the way in which the prices are generated may not be as important. Perhaps the more important issue is whether the quoted prices are binding and valid. If they are, then an acceptable option pricing model's predictions must be consistent with the observed option price dynamics.³

3.2 How often do call and put prices go up or down together?

In this section we answer two related questions: (i) How often do call and put prices go up or down together? And (ii) when call and put prices move together, are they more likely to go down than to go up together? If call and put prices indeed change in the same direction, one of them must be changing in a way inconsistent with the predictions of Proposition 1. Thus we still refer to such occurrences as “violations” of the model predictions. Specifically, for a fixed τ and a given K , we distinguish among the four cases below:

- **Type A violation:** $\Delta S(t) > 0$, but $\Delta C(t, \tau, K) > 0$ and $\Delta P(t, \tau, K) > 0$.
- **Type B violation:** $\Delta S(t) > 0$, but $\Delta C(t, \tau, K) < 0$ and $\Delta P(t, \tau, K) < 0$.
- **Type C violation:** $\Delta S(t) < 0$, but $\Delta C(t, \tau, K) > 0$ and $\Delta P(t, \tau, K) > 0$.
- **Type D violation:** $\Delta S(t) < 0$, but $\Delta C(t, \tau, K) < 0$ and $\Delta P(t, \tau, K) < 0$.

³ It is possible that in a given type I violation, the number of contracts at which the changed bid or ask price is binding may be small (e.g., 10 contracts). Consequently the ability to trade and profit from such violations may be limited. While the economic significance of trading on type I and other violations is an interesting topic, the Berkeley Options Database does not include information on bid or ask sizes. Hence we cannot address this and other related issues in detail.

To ensure that any of the violations is not due to a violation of the put-call parity, we exclude all pairs of call and put price changes that violate the put-call parity (about 3%).

For each sampling frequency, Table 4 reports [Total (%)], the percentage of observed put-call pairs that represent type A, B, C, or D violations. Again, we begin with the results based on the cash index. The occurrence rate for any type of joint violation lies between 0.7% and 6.6%, and it increases with the intraday sampling interval. For example, the type A violation frequency is 2.0% at the 30-minute, 2.7% at the 1-hour, 2.9% at the 2-hour, and 3.0% at the 3-hour sampling interval. As a result, the total percentage of observed put-call pairs that represent a joint violation (type A through type D altogether) also increases with the intraday sampling interval.

Note that at a given sampling interval, type B and type D violations are more likely to occur than type A and type C violations. That is, *whether the underlying asset price is up or down, it is always more likely for both prices of a put-call pair to go down than to go up together*. For instance, at the 3-hour sampling interval, we have the occurrence rate at 3.0% for type A, 5.2% for type B, 4.1% for type C, and 4.6% for type D violations. There are two possible contributing factors for this phenomenon. First, option premiums are subject to inevitable time decay as the options come closer to expiration, which affects option prices negatively. Thus, whether the underlying is up or down during a given time interval, both call and put prices

Table 4
Joint violation occurrences by type and by sampling interval

Sampling interval	Number of observations	Spot asset used	Type A (%)	Type B (%)	Type C (%)	Type D (%)	Total (%)
30 minutes	39,074	Cash index	2.0	2.4	1.6	2.1	8.1
		Index futures	2.0	2.3	1.4	2.2	8.0
1 hour	19,338	Cash index	2.7	3.4	2.4	3.3	11.8
		Index futures	3.0	3.2	2.1	3.4	11.8
2 hours	9570	Cash index	2.9	4.4	3.1	3.9	14.3
		Index futures	3.1	4.0	2.8	3.9	13.8
3 hours	3226	Cash index	3.0	5.2	4.1	4.6	16.9
		Index futures	4.2	5.8	2.6	3.4	16.0
1 day	2482	Cash index	0.9	6.6	0.7	2.6	10.8
		Index futures	0.7	4.9	0.9	4.2	10.7

For matched pairs of call and put options (i.e., the call and the put with the same strike price and maturity), reported below are type A, type B, type C, and type D violation rates, each as a percentage of total observations at a given sampling interval:

Type A: $\Delta S > 0$, $\Delta C > 0$, $\Delta P > 0$

Type B: $\Delta S > 0$, $\Delta C < 0$, $\Delta P < 0$

Type C: $\Delta S < 0$, $\Delta C > 0$, $\Delta P > 0$

Type D: $\Delta S < 0$, $\Delta C < 0$, $\Delta P < 0$

The samples are separately obtained by sampling price changes once every (i) 30 minutes, (ii) 1 hour, (iii) 2 hours, (iv) 3 hours, and (v) 1 day. The rows under "Cash index" are obtained by using the S&P 500 cash index, while those under "Index futures" by using the lead-month S&P 500 futures, as a stand-in for the underlying asset. Observations that violate the put-call parity are dropped from the sample.

have a slightly stronger tendency to go down together than to go up. Second, suppose that we go beyond the one-dimensional diffusion models and allow volatility to be stochastic over time [say, as in Heston (1993)]. Then volatility is generally believed to be negatively correlated with stock returns: When the stock price goes up, volatility is likely to go down, and vice versa. As volatility affects option premiums positively, a stock price increase-induced decline in volatility will exert a negative impact on both put and call premiums. If volatility effect dominates stock price effect, both put and call prices will decrease together. This reasoning, however, only helps explain why type B violations are more likely to occur than type A violations,⁴ but not the finding that when the underlying goes down, type D violations are more likely than type C violations. In actuality, the documented patterns in Table 4 should be mostly due to the joint working of both the time decay and the negatively correlated volatility factor: When the underlying price goes up, both the lower volatility and the time decay factor should make call and put premiums lower. This joint working hypothesis is also consistent with the fact that in Table 4, the type B occurrence frequency is, at each given sampling interval, higher than type D, that is, it is more likely for call and put prices to go down together in a rising than in a declining stock market. This is especially true at the daily sampling interval, in which case the likelihood for a pair of put-call prices to go down together is 6.6% in a rising day and 2.6% in a declining day.⁵

To continue the above theme, observe that the rates at which put-call pairs move up or down together differ significantly between interday and intraday intervals. Specifically, as noted earlier, when the *intraday sampling interval* increases successively from 30 minutes to 3 hours, the occurrence rates for type A, type C, and type D each increase monotonically. But changing from the 3-hour sampling interval to the daily interval, we see the occurrence rates for all three types going down (e.g., from 4.1% to 0.7% for type C). It is particularly striking that during intraday intervals, put and call prices go up together, ranging from 3.6% to 7.1% of the time (type A and type C combined), but from day to day, it is relatively rare for put and call prices to go up together (about 1.6%). This is puzzling. Why do option prices behave differently intraday than interday? Finally, as seen in Table 4, the results

⁴ Implicit in this discussion is the assumption that when volatility and underlying price changes are negatively correlated, the probability for volatility to decrease is, conditional on an underlying price increase, higher than for volatility to increase. As the simulations will show in a later section, this implicit assumption holds at least under the stochastic volatility model framework. From another standpoint, a vast GARCH literature has pinpointed that index volatility increases substantially more when the level of the index unexpectedly goes down than when it unexpectedly goes up [see Glosten, Jagannathan, and Runkle (1993), Amin and Ng (1997), and Kroner and Ng (1998)].

⁵ It is worthwhile to acknowledge that the documented type A–D joint violations are related to type I (type IV) violations for calls, and type IV (type I) violations for puts, provided the put-call parity is strictly satisfied. To see this point, note that type B and type C violations imply that the call is a type I and the put a type IV violation. By the same logic, type A and type D violations imply that the call is a type IV and the put a type I violation.

hold even if we replace the cash index with the lead-month futures price as a stand-in for the unobservable underlying asset.

3.3 Do violation occurrences differ across moneyness and maturity?

We now study the structure of violation occurrences across moneyness and maturity. For this we focus on the hourly call option sample, as the results are similar for puts and for other sampling intervals.

As a first check we plot, in Figure 1, the hourly changes in call prices against those in the S&P 500 index during the same interval, for nine different moneyness-maturity categories. For almost every category of calls, there is a significant number of $(\Delta S, \Delta C)$ pairs that lie in the second and the fourth quadrants of each plot, which should not occur if the predictions by one-dimensional diffusion models were consistent with option price dynamics. Visually, OTM and ATM calls of each term to expiration seem to have more $(\Delta S, \Delta C)$ observations in the second and the fourth quadrants. When the lead-month futures price is used in place of the cash index, the resulting plots are similar to those in Figure 1, but with slightly more scattered points (see Figure 2).

More precisely, in Table 5 we report six numbers for each moneyness maturity category: (i) occurrence rate (as a percentage of total observations in the moneyness maturity category) of type I, type II, type III, and type IV violations combined; (ii) occurrence rate of type I violations; (iii) occurrence rate of type I violations in which the underlying price goes down but the call price goes up; (iv) occurrence rate of type I violations in which the underlying asset goes up but the call price goes down; (v) type II violation rate; and (vi) type IV violation rate. We again use both the cash index and the lead-month futures price to determine changes in the underlying asset.

We start with the cash index-based results in Table 5. Of a given term to expiration, OTM calls violate the model predictions most frequently, followed by ATM and then ITM calls (except for short-term options, in which case the ITM calls come second and the ATM calls are last). The fact that the combined violation rate is much higher for the relatively cheaper calls (i.e., OTM calls and short-term calls) suggests that market microstructure factors must be playing a role. This is especially evident from the patterns of type II violations. Of a given maturity, the OTM calls in Table 5 have the most frequent type II violations, followed by ATM and ITM calls (except for short-term ones). This is the case because, as discussed before, the OTM calls have the lowest option deltas and hence are least sensitive to underlying price changes. Thus, unless the market moves drastically, the implied change in value for an OTM call may not be large enough to overcome the minimum tick size of $\$ \frac{1}{16}$ or $\$ \frac{1}{8}$. For type IV violations, ITM options have the highest violation rate and OTM options the least. Note that type IV violations are as high as 23.6% for short-term ITM options and only 0.4% for short-term




Figure 1
For all calls of a given moneyness/maturity group, their price changes (denoted by ΔC) are plotted against the underlying S&P 500 cash index changes (denoted by ΔS). All price changes are sampled every hour for the period from March 1, 1994, to August 31, 1994. OTM, ATM, and ITM stand for out of the money, at the money, and in the money options, respectively.




Figure 2 For all calls of a given moneyness/maturity group, their price changes (denoted by ΔC) are plotted against the lead-month S&P 500 futures price changes (denoted by ΔS). All price changes sampled every hour for the period from March 1, 1994, to August 31, 1994, OTM, ATM, and ITM stand for out of the money, at the money, and in the money options, respectively.

Table 5

Violation rates across moneyness and maturity

$\frac{S}{K}$	Type of violation	Cash index term to expiration			Index futures term to expiration		
		Short	Medium	Long	Short	Medium	Long
OTM	Type I, II, III and IV	63.8	64.3	58.3	65.5	65.0	53.3
	Type I only	14.6	14.9	15.0	15.5	17.2	11.3
	Type I, and $\Delta S < 0$, $\Delta C > 0$	6.1	5.9	7.1	6.6	6.8	4.6
	Type I, and $\Delta S > 0$, $\Delta C < 0$	8.5	9.0	7.9	8.9	10.5	6.7
	Type II only	48.6	46.1	38.3	47.9	45.2	36.5
	Type IV only	0.4	3.1	4.7	1.4	2.6	4.8
ATM	Type I, II, III and IV	40.2	56.1	43.5	40.9	55.4	35.4
	Type I only	13.0	15.1	15.9	13.3	15.3	9.2
	Type I, and $\Delta S < 0$, $\Delta C > 0$	6.0	6.8	7.8	6.0	6.8	3.8
	Type I, and $\Delta S > 0$, $\Delta C < 0$	7.0	8.3	8.1	7.3	8.5	5.4
	Type II only	18.2	32.4	17.1	17.7	31.4	15.8
	Type IV only	8.5	8.3	10.2	8.0	7.6	9.1
ITM	Type I, II, III and IV	62.1	38.0	41.7	60.6	25.4	33.9
	Type I only	11.5	13.7	15.5	12.3	6.2	8.9
	Type I, and $\Delta S < 0$, $\Delta C > 0$	5.0	6.0	7.6	5.5	2.8	3.6
	Type I, and $\Delta S > 0$, $\Delta C < 0$	6.5	7.6	7.9	6.7	3.4	5.3
	Type II only	26.2	10.6	12.2	26.2	9.8	11.1
	Type IV only	23.6	13.4	13.8	20.2	7.8	12.2

For each option moneyness/maturity category, the reported numbers are violation occurrence rates (%) of (i) type I, type II, type III, and type IV combined, (ii) type I only, (iii) type I with $\Delta S < 0$ and $\Delta C > 0$, (iv) type I with $\Delta S > 0$ and $\Delta C < 0$, (v) type II only, and (vi) type IV only, each as a percentage of total observations. The results are based on call prices *sampled every hour*. The columns under "Cash index" are obtained by using the S&P 500 cash index, while those under "Index futures" by using the lead-month S&P 500 futures, as a stand-in for the underlying asset. OTM, ATM, and ITM stand for out of the money, at the money, and in the money, respectively. Short-, medium-, and long-term refer to options with less than 60 days, with 60–180 days, and with more than 180 days, respectively, to expiration.

OTM options. Between ITM and OTM options, the former have deltas close to one, and their prices are more likely to be adjusted. Consequently, ITM options are more likely to be affected by the tick size restriction.

The type I violation rates, reported under "Type I only" in Table 5, display a somewhat different structure than the one presented for type II and type IV. First, within each moneyness class, the type I violation rate monotonically increases with the term to expiration: the longer an option's remaining life, the more likely its price goes in the opposite direction with the underlying asset. Second, for short-term options, the more in the money, the less likely the call price will change in the opposite direction with the underlying. But for medium- and long-term options, the ATM calls are most likely to change in the opposite direction. Therefore, while there is some association between type I violation rate and the option's time to expiration, there is no clear relationship between option moneyness and type I rate.

While the two occurrence rates for "Type I: $\Delta S < 0$ but $\Delta C > 0$ " and "Type I: $\Delta S > 0$ but $\Delta C < 0$ " exhibit the same patterns across moneyness and maturity as does the "Type I only" occurrence rate (Table 5), they also suggest another regularity: of type I violations, there are more cases with " $\Delta S > 0$ but $\Delta C < 0$ " than with " $\Delta S < 0$ but $\Delta C > 0$," as the former always has a higher occurrence rate regardless of moneyness and maturity.

Both such occurrences are, of course, inconsistent with the one-dimensional diffusion models. But they are consistent with the fact that in Table 4 type B violations are more likely than type C. Thus the negatively correlated stochastic volatility and the time decay explanation discussed earlier may apply here as well.

When the lead-month S&P futures price is used as surrogate for the underlying asset, some of the aforementioned patterns no longer apply. The most evident difference from the cash index case is how the type I violation rate relates to option maturity: for OTM and ATM calls, the relationship is now strongly hump-shaped, with the medium-term calls having the most frequent type I violations, while for ITM calls the relationship is U-shaped, with the medium-term options showing the least violations. Why do medium-term OTM and ATM call prices move in the opposite direction with the underlying more often than their short-term and long-term counterparts? Why do medium-term ITM calls have the lowest type I violation rate among all ITM calls? These patterns are quite puzzling.

3.4 Are violation occurrences related to market microstructure factors?

Next we study violation frequencies by separately sorting options according to (i) time of day, (ii) dollar bid-ask spread, (iii) number of quote revisions, and (iv) daily trading volume. As before, we focus on the hourly call-option sample. First, according to each of the four criteria we divide the overall sample into six groups of, whenever possible, equal size. Next, for each given group, we obtain the ratio between the number of option observations representing a type I (type II or type IV) violation and the total number of observations in the group. We report this ratio as the group's type I (type II or type IV) violation occurrence rate in Table 6. The purpose is to see whether and how the violation rates are related to market microstructure factors.

We begin with the "time of day" patterns in Table 6. It is clear that the occurrence rates for type I and type II are the highest (17.4% and 30.3%, respectively) from 11:00 to 12:00 A.M. CST. The type IV violations appear to be the highest from 12:00 to 1:00 P.M. CST. Generally these violations occur more frequently during the middle hours of the day than during the initial and the final hours. This intraday occurrence pattern is somewhat consistent with the findings in the existing literature that both intraday trading volume and volatility in the stock market are persistently U-shaped from the morning to market close, indicating relatively active trading at both ends of each trading day [e.g., Wood, McNish, and Ord (1985) and Chan, Christie, and Schulz (1995)]. Still, it is hard to explain why such an intraday pattern exists.

Note that the occurrence rates for type II (type IV) violations are monotonically decreasing (increasing) in the dollar bid-ask spread size: the narrower (wider) the bid-ask spread, the more (less) often violations occur. Both types of violations are strongly related to tick size and bid-ask spread. These results

Table 6
The impact of market microstructure factors

Time of day	Violation rate (%)	Dollar bid-ask spread	Violation rate (%)	No. of quote revisions	Violation rate (%)	Trading volume	Violation rate (%)
9:00–10:00 A.M.	13.0 (20.5) {10.4}	$< \frac{3}{16}$	14.3 (35.4) {1.8}	< 16	12.7 (46.5) {11.3}	0	13.6 (20.2) {13.6}
10:00–11:00 A.M.	11.7 (21.3) {9.5}	$\frac{3}{16} - \frac{1}{4}$	15.3 (26.0) {4.3}	16–30	13.7 (34.0) {11.7}	0–14	15.0 (27.3) {15.1}
11:00 A.M.–12:00 P.M.	17.4 (30.3) {11.3}	$\frac{1}{4} - \frac{3}{8}$	15.2 (24.8) {7.2}	30–67	13.9 (25.5) {11.0}	14–115	13.5 (30.1) {10.1}
12:00–1:00 P.M.	15.1 (22.9) {12.5}	$\frac{3}{8} - \frac{1}{2}$	14.4 (24.8) {10.9}	67–310	13.9 (16.5) {9.5}	115–350	14.5 (25.6) {7.9}
1:00–2:00 P.M.	13.9 (23.1) {12.2}	$\frac{1}{2} - \frac{3}{4}$	13.1 (18.3) {14.9}	310–755	15.2 (10.8) {9.9}	350–935	13.7 (25.8) {6.3}
2:00–3:00 P.M.	12.4 (21.6) {10.2}	$\geq \frac{3}{4}$	12.7 (19.7) {17.1}	≥ 755	14.1 (5.6) {12.7}	≥ 935	14.3 (21.8) {3.7}

Reported below are type I, type II (in parentheses), and type IV (in curly brackets) violation frequencies partitioned according to (i) time of the day (Central Standard Time), (ii) dollar bid-ask spread, (iii) daily number of quote revisions, and (iv) daily trading volume (number of contracts traded). Each violation rate is expressed as a percentage of option observations in a given group that represents a type I (type II or type IV) violation. All calculations use the S&P 500 cash index to infer changes in the underlying asset. The results are based on the hourly call option sample.

are consistent with those of Table 5 in that OTM (ITM) options lead to more frequent type II (type IV) violations than ITM (OTM) options, as the former tend to have lower bid-ask spreads and option deltas. Type I violations are, by and large, more uniformly distributed across option groups with different bid-ask spreads than its type II and type IV counterparts. For example, type I violation frequencies are 14.3% and 12.7% for options with bid-ask spreads less than $\$ \frac{3}{16}$ and greater than $\$ \frac{3}{4}$, respectively. These results reinforce our earlier deduction that the majority of type I violations for calls (puts) are unrelated to type IV violations for puts (calls), and to market-microstructure factors.

The occurrence of type I and type II violations may be related to how often price quotes are revised. If the occurrence rates for both types of violation are highly negatively related to the frequency of quote updating, our evidence may not be interpreted as being against the model predictions. Instead, it may be treated only as a consequence of infrequent and perhaps improper quote updating, especially if the most frequently revised options would have no violations. In Table 6, it is shown that the type I occurrence does not vary significantly, whereas the frequency of type II violations decreases dramatically (from 46.5% to 5.6%) with the number of quote revisions per day. This

striking contrast reaffirms our assertions that type I violations reflect something fundamental in the actual option valuation process but missing from the models, and that type II violations are much more sensitive to market microstructure factors.

Finally, the type I occurrence rate shows no clear association with an option contract's trading volume. Thus actively traded calls do not necessarily have a lesser chance to move in the opposite direction with the underlying asset. Observe that the type IV violation rate decreases monotonically with trading volume. Because they are the most sensitive to changes in the underlying price, prices of inactively traded options (e.g., the ITM options) are more likely to be overadjusted.

3.5 To what extent are type I violations related to time decay?

Having presented a large body of empirical evidence supporting violation of model predictions, we now devote our attention to understanding the role of time decay. That is, to what extent can it help reconcile type I violations for which $\Delta C < 0$ and $\Delta S > 0$ (it is of no appeal in characterizing $\{\Delta C > 0, \Delta S < 0\}$ violation pairs). Since it is impossible to divide option price changes into their time-decay and non-time-decay constituents without specifying a model, we examine this issue in the Black–Scholes (henceforth, BS) setting. First, we assess the likelihood that a fraction of $\{\Delta C < 0, \Delta S > 0\}$ pairs are indeed a consequence of time decay. In pursuing such a goal, a two-step procedure is followed. In the initial step, we pick a representative day, say, March 1, 1994. Correspondingly, we set $r = 5\%$, $\sigma = 15\%$, $S(t) = 460$, and then compute the BS option price with maturity τ and strike K . Fixing the observation interval (say, 1 hour or 1 day), we recompute the option price. Provided the increase in the spot price over the observation interval is no larger than some critical level α , the change in the call will be negative (just due to time decay). Using an algorithm to determine α subject to the constraints $\Delta C < 0$ and $\Delta S > 0$, we calculate the conditional probability (for expected stock return μ):

$$\Pr(0 < \Delta S < \alpha \mid \Delta S > 0) = \frac{\mathcal{N}(d_2) - \mathcal{N}(d_1)}{1 - \mathcal{N}(d_1)}, \quad (7)$$

where $d_1 \equiv -1/\sigma (\mu - (1/2)\sigma^2) \sqrt{\tau}$, $d_2 \equiv d_1 + (1/\sigma \sqrt{\tau}) \ln[1 + \alpha S(t)^{-1}]$, and $\mathcal{N}(\cdot)$ represents the cumulative normal distribution function. This probability captures the percentage of $\{\Delta C, \Delta S > 0\}$ pairs that are caused by time decay. Keeping $\mu = 15\%$, panel A of Table 7 tabulates α and the resulting probabilities for calls with 30 days, 60 days, and 180 days to expiration. At the 1-hour interval and for an OTM call, there is only a 2.1% chance that the $\{\Delta C, \Delta S > 0\}$ pair will have $\Delta C < 0$. The picture is different at the daily interval: for an OTM call with 30 days to expiration, the time-decay component is significant with a 11.3% type I violation probability. For options with longer maturities, the corresponding probabilities are smaller.

Table 7
The impact of time decay on the option premium

Panel A: Conditional probability of type I violations with $\Delta S > 0$ and $\Delta C < 0$

Δt		30-day call		60-day call		180-day call	
		Cutoff α	Prob. (%)	Cutoff α	Prob. (%)	Cutoff α	Prob. (%)
1 hour	OTM	0.02	2.1	0.00	0.0	0.00	0.0
	ATM	0.00	0.0	0.00	0.0	0.00	0.0
	ITM	0.00	0.0	0.00	0.0	0.00	0.0
2 hours	OTM	0.04	3.0	0.02	1.5	0.00	0.0
	ATM	0.02	1.5	0.00	0.0	0.00	0.0
	ITM	0.00	0.0	0.00	0.0	0.00	0.0
3 hours	OTM	0.06	2.4	0.04	2.4	0.02	1.2
	ATM	0.02	1.2	0.02	1.2	0.00	0.0
	ITM	0.00	0.0	0.00	0.0	0.00	0.0
1 day	OTM	0.52	11.3	0.32	6.9	0.16	3.5
	ATM	0.28	6.1	0.20	4.3	0.12	2.6
	ITM	0.12	2.6	0.12	2.6	0.10	2.1

Panel B: Changes in options prices during Δt

Term to expiration		Call price	Magnitude of time decay during Δt for $\Delta t =$			
			1 hour	2 hours	3 hours	1 day
$\tau = 30$ days	OTM	\$2.08	\$0.00	\$0.01	\$0.01	\$0.10
	ATM	8.30	0.00	0.01	0.01	0.16
	ITM	23.08	0.00	0.01	0.01	0.12
$\tau = 60$ days	OTM	5.15	0.01	0.01	0.02	0.10
	ATM	13.17	0.01	0.01	0.02	0.13
	ITM	26.56	0.00	0.01	0.01	0.11
$\tau = 180$ days	OTM	15.75	0.00	0.01	0.01	0.08
	ATM	25.33	0.00	0.01	0.01	0.09
	ITM	38.01	0.01	0.01	0.01	0.08

Reported in panel A is the probability ($0 < \Delta S < \alpha \mid \Delta S > 0$) with the cutoff α so determined that $\Delta S > 0$ but $\Delta C < 0$, in a Black–Scholes setting. In computing this probability (over interval Δt), the inputs are taken from March 1, 1994. Specifically, $S(t) = \$460$, $\sigma = 15\%$ (the implied volatility), $r = 5\%$. In this illustration, the compounded (annualized) S&P 500 index return is fixed at 15%. In panel B, we report the change in the option price during Δt ($= 1$ hour, 2 hours, 3 hours, or 1 day).

Continuing to use March 1, 1994, inputs as the basis, we can also infer call price changes over a short interval Δt . Panel B of Table 7 displays a few benchmark calculations. For example, the time decay (as surrogated by the instantaneous change in the call price) is large at the interday frequency; it diminishes substantially, however, as the sampling interval shrinks. In particular, it justifies the construction of intraday samples to minimize the impact of time decay.

In the same spirit, we also utilize the BS model to get a sense of type I violations that can possibly be attributable to time decay. For this purpose, set $\sigma[t, S] \equiv \sigma$ in Equation (4) and analytically compute $\{C_t + (1/2)\sigma^2 S^2 C_{SS}\}$ (using BS-implied volatility). Then, for each given Δt ranging from an hour to a day, we adjust the observed call price changes by this time decay. We

repeat this procedure for each option contract and at each sampling frequency. The average option time decay is $-\$0.001$, $-\$0.003$, and $-\$0.031$ for the 1 hour, 3-hour, and 1-day samples, respectively. Further, as the magnitude of the time decay is not large (both in absolute or relative terms), the type I violation frequencies are rendered stable. We still obtain type I violations of 13.9% (1-hour interval), 16.3% (3-hour interval), and 8.9% (daily interval). These statistics are very close to the ones reported in Table 3. Collectively these exercises corroborate our assertions that time decay can explain some type I violations, and it is more important in daily sampling schemes, but not so much for options sampled intraday. When we investigate type A through D violation frequencies in the matched sample, our results are consistent with the ones recorded in Table 4. To the extent that the magnitude of time decay is comparable across alternative option pricing models, our conclusions are relatively robust.

3.6 How large are price changes in a typical violation?

To gauge the relative magnitude of the documented violations, we report in Table 8 the average changes in the underlying asset, denoted by $\overline{\Delta S}$, and the corresponding average changes in call prices, denoted by $\overline{\Delta C}$, where the averaging is based only on those $(\Delta S, \Delta C)$ pairs in a given moneyness maturity category that each represent a type I (or type II in panel B) violation. To avoid cancellation of price changes of the opposite sign, we separate those observations in which $\Delta S > 0$ but $\Delta C < 0$ from those in which $\Delta S < 0$ but $\Delta C > 0$. If both $\overline{\Delta S}$ and $\overline{\Delta C}$ are small relative to the minimum tick size, then the violations are on average economically insignificant.

Since the results based on the futures price are similar, we focus on those based on the cash index and shown in panel A of Table 8. Take the type I violations in which $\Delta S < 0$ but $\Delta C > 0$. Of the OTM calls, the average value of ΔS in a violation is $-\$0.71$, $-\$0.65$, and $-\$0.40$, whereas the corresponding average value of ΔC in a violation is $\$0.15$, $\$0.17$, and $\$0.35$ for short-, medium-, and long-term options, respectively. The order of these magnitudes are all larger than the minimum tick sizes. Similar conclusions can be drawn based on the average values of ΔS and ΔC for type I violations by other moneyness maturity categories. The documented type I violations are thus economically significant.

Panel B of Table 8 displays the average value of ΔS in a type II violation, together with the average bid-ask spread, for each option moneyness maturity category. Again, we separate those observations with $\Delta S > 0$ from those with $\Delta S < 0$. Then, except for OTM options, type II violations occur, on average, when the underlying price changes by an amount close to a typical bid-ask spread for options of the corresponding moneyness and maturity. For OTM options, the required price adjustment (after taking into account the appropriate option deltas) may also be close to the respective bid-ask spreads.

Thus this evidence also supports our claim that type II violations are significantly related to the size of bid-ask spread. Since the magnitudes of type IV violations are not so informative, they are omitted from our discussion.

3.7 Implications for hedging

The persistent occurrence of type I, II, and IV violations indicates that call and put prices are not perfectly correlated with the underlying price. Hence these contracts are not as ideal a hedging instrument for equity portfolios as textbooks suggest. Another lesson learned from standard textbooks is that in applying an option pricing model to implement a dynamic options-based hedge, one should rebalance the hedge as frequently as possible, ideally continuously. For example, managers who adopted portfolio insurance were all told to do so prior to and during the October 1987 market crash. Is it empirically true that the more often one revises an option hedge, the smaller the hedging errors? Given the type I, II, and IV violations, our conjecture is that it may not. The reason is that if one applies a one-dimensional diffusion model to design a delta-neutral hedge, then increasing the frequency of hedge revision will more severely compound the errors caused by the imperfect correlation between the options and the underlying asset.

Table 8
Average magnitude of type I and type II violations in the hourly sample

Moneyness		$\Delta S < 0$ and $\Delta C > 0$			$\Delta S > 0$ and $\Delta C < 0$		
		Term to expiration			Term to expiration		
		Short	Medium	Long	Short	Medium	Long
OTM	$\overline{\Delta S}$	-\$0.71	-\$0.65	-\$0.40	\$0.64	\$0.52	\$0.54
	$\overline{\Delta C}$	\$0.15	\$0.17	\$0.35	-\$0.14	-\$0.22	-\$0.46
ATM	$\overline{\Delta S}$	-\$0.29	-\$0.34	-\$0.31	\$0.33	\$0.43	\$0.32
	$\overline{\Delta C}$	\$0.19	\$0.36	\$0.35	-\$0.24	-\$0.38	-\$0.35
ITM	$\overline{\Delta S}$	-\$0.21	-\$0.20	-\$0.19	\$0.36	\$0.21	\$0.22
	$\overline{\Delta C}$	\$0.40	\$0.20	\$0.33	-\$0.45	-\$0.23	-\$0.33

		$\Delta S < 0$ and $\Delta C = 0$			$\Delta S > 0$ and $\Delta C = 0$		
		Short	Medium	Long	Short	Medium	Long
OTM	$\overline{\Delta S}$	-\$0.80	-\$0.62	-\$0.79	\$0.64	\$0.67	\$0.69
	Bid-ask spread	\$0.10	\$0.24	\$0.51	\$0.10	\$0.24	\$0.51
ATM	$\overline{\Delta S}$	-\$0.42	-\$0.55	-\$0.71	\$0.47	\$0.55	\$0.59
	Bid-ask spread	\$0.36	\$0.50	\$0.74	\$0.36	\$0.50	\$0.74
ITM	$\overline{\Delta S}$	-\$0.37	-\$0.37	-\$0.59	\$0.35	\$0.37	\$0.57
	Bid-ask spread	\$0.84	\$0.70	\$0.80	\$0.84	\$0.70	\$0.80

Reported are (i) the average change in the S&P 500 index, denoted by $\overline{\Delta S}$, and (ii) the average corresponding change in the call option price, denoted by $\overline{\Delta C}$, where the averaging is based on all option observations, in a given moneyness maturity category, that represent a type I violation (for panel A) or a type II violation (for panel B). For panel A, we distinguish between the " $\Delta S > 0$ but $\Delta C < 0$ " case and the " $\Delta S < 0$ but $\Delta C > 0$ " case, so as to avoid cancellation of changes of opposite signs. All calculations use the S&P 500 cash index to infer to changes in the underlying asset. The results are based on the hourly call-option sample. For ease of comparison, panel B also reports the average bid-ask spread for options in a given moneyness maturity category.

Table 9
Rebalancing frequency and hedging errors

Moneyness	Rebalancing frequency					
	30 minutes	1 hour	2 hours	3 hours	1 day	2 days
OTM	\$0.92	\$0.90	\$0.88	\$0.64	\$0.86	\$1.49
ATM	1.43	1.14	1.01	1.37	1.77	3.38
ITM	1.83	1.72	1.82	2.31	0.88	1.86

Constructed using the Black–Scholes model, all delta-neutral hedges of calls use the underlying asset as the instrument. Each target call has about 2 months to expiration. To implement the Black–Scholes hedges, we first obtain the implied spot volatility by minimizing the sum of squared pricing errors across all options collected at the time of hedge revision. Next, the implied volatility is used as input to calculate the Black–Scholes delta, which determines both the number of shares long of the underlying stock and the required risk-free bond position in the replicating portfolio of a given target call. During the 2-month period, if there is at any time a shortage (surplus) in the value of the replicating portfolio relative to that of the target call, we borrow (lend) the amount at the risk-free rate (which allows us to compute hedging errors only at the end of hedging period). Finally, at the end of the hedging period we calculate the difference between the target option price and the liquidation value of the replicating portfolio, which gives the hedging error for the option. These steps are repeated for each 2-month option and separately for each rebalancing interval. Reported below are the average absolute hedging errors for each moneyness category.

To illustrate this point, we choose the standard BS model again as an example, even though we realize that the BS model is surely misspecified for both pricing and hedging. But our purpose is to see how documented violations affect hedging effectiveness by examining the relationship between a delta-neutral strategy's hedging error and the hedge-revision frequency. Specifically, each BS-based delta-neutral strategy is to hedge an S&P 500 call option with about 2 months to expiration initially.⁶ Table 9 reports the average absolute hedging errors for OTM, ATM, and ITM calls. Hedge-revision intervals considered range from 30 minutes to 1 hour, 2 hours, and so on, with 2 days being the longest.

For each given moneyness category, the average absolute hedging errors show a pronounced U-shaped pattern in relation to the revision interval. For instance, the absolute hedging error for ATM calls is, on average, \$1.43, \$1.14, \$1.01, \$1.37, \$1.77, and \$3.38 if the hedge is revised every 30 minutes, every trading hour, every 2 hours, every 3 hours, every day, and every other day, respectively. Thus the more often a hedge is adjusted within a day, the larger the hedging errors. That is, as the hedge-revision frequency increases within a day, it, on the one hand, makes the replication portfolio mimic the option target more and more closely and, on the other hand, allows type I, II, and IV violations to exert a stronger impact on the hedging errors. Beyond a certain point, the negative impact induced by type I, II, and IV violations can dominate the benefit of the additional rebalancing efforts. Therefore, in light of our documented option price patterns, rebalancing a hedge as often as possible is not necessarily optimal.

⁶ The focus is on these options as they are actively traded (see Table 2). We also use options with other terms to expiration as hedging targets and find the results are similar to the ones reported.

4. Beyond One-Dimensional Diffusion Models

Having clarified the role of market microstructure, put-call parity violations, and option time decay as partial descriptors of violation patterns, this part of our empirical analysis examines the explanatory power of a model in the two-dimensional (Markov) diffusion class. On a more profound level, we check whether this model is capable of overturning type I violations under the one-dimensional paradigm. This emphasis is guided by the intuition that if type II and type IV violations are market microstructure-dependent, then even a more elaborate forcing process will have difficulty fitting such observations.

Among the choices of second state variable, return volatility is clearly a natural candidate as there is abundant evidence (both academic and casual empirical) showing that stock volatility changes stochastically over time [see Amin and Ng (1993, 1997), Glosten, Jagannathan, and Runkle (1993), and Kroner and Ng (1998)]. To address the above challenges, we use Heston's (1993) stochastic volatility model (the *SV model*) as an example (partly because it offers a closed-form formula). In Heston's model, stochastic volatility affects call and put prices positively, and it is not necessarily perfectly correlated with the underlying price. Specifically, let the underlying price and its return volatility, denoted by $V(t)$, follow (under the equivalent martingale measure) the respective processes below:

$$\frac{dS(t)}{S(t)} = r dt + \sqrt{V(t)} dW_s(t), \quad (8)$$

$$dV(t) = [\theta_v - \kappa_v V(t)] dt + \sigma_v \sqrt{V(t)} dW_v(t), \quad (9)$$

where the structural parameter, κ_v , reflects the speed at which $V(t)$ approaches its long-run mean θ_v/κ_v , σ_v is its variation coefficient, and W_s and W_v are two standard Brownian motions with a correlation of ρ . From Heston (1993), the call pricing formula is

$$C(t, \tau) = S(t) \Pi_1(t, \tau) - K e^{-r\tau} \Pi_2(t, \tau), \quad (10)$$

where the probabilities, $\Pi_j(t, \tau) = 1/2 + 1/\pi \int_0^\infty \text{Re} \left[\frac{K^{-i\phi}}{i\phi} f_j(t, \tau; \phi) \right] d\phi$, for $j = 1, 2$, with the characteristic functions, f_1 and f_2 , are given in the appendix. The call dynamics are determined by

$$dC(t, \tau, K) = \mu_c(t, \tau, K) dt + C_S dS(t) + C_V dV(t), \quad (11)$$

where the two partial derivatives, C_S and C_V , are given in Equations (A3) and (A4) of the appendix, and $\mu_c(t, \tau, K)$ is determined according to Ito's lemma.

Table 10
Simulated occurrence of violations

Sampling interval	Type I for calls (%)	Type I for puts (%)	Type A (%)	Type B (%)	Type C (%)	Type D (%)
30 minutes	11.5	11.1	5.8	6.0	5.5	5.3
1 hour	11.4	11.1	5.8	6.0	5.4	5.3
2 hours	11.4	11.1	5.8	6.1	5.3	5.3
3 hours	11.3	11.1	5.8	6.1	5.2	5.3
1 day	11.3	11.2	5.7	6.5	4.8	5.5

In the simulation experiment, each $(S(t + \Delta t), V(t + \Delta t))$ pair is generated according to the discretized version of the stochastic volatility model of Heston (1993):

$$S(t + \Delta t) - S(t) = rS(t)\Delta t + \sqrt{V(t)}S(t)\epsilon_s(t)\sqrt{\Delta t},$$

$$V(t + \Delta t) - V(t) = [\theta_v - \kappa_v V(t)]\Delta t + \sigma_v \sqrt{V(t)}\epsilon_v(t)\sqrt{\Delta t},$$

with $\rho \equiv \text{cov}(\epsilon_s(t), \epsilon_v(t))$. We set $S(t) = \$460$, $\sqrt{V(t)} = 15\%$, $r = 5\%$, $\theta_v = 0.80$, $\kappa_v = 2.18$, $\rho = -0.70$, and $\sigma_v = 0.53$. In each simulation trial, we compute the model-implied changes in the calls and puts (using the closed-form stochastic volatility option model) and hence the occurrence of (i) type I violations for calls; (ii) type I violations for puts; and (iii) type A through type D (joint) violations. The reported numbers are the fraction of the violation occurrence across 10,000 simulation trials.

4.1 Simulation experiment

Option prices generated according to the SV model can indeed exhibit patterns that are similar to the documented type I violations. To see this, we again take the data from March 1, 1994, to back out the SV model's parameters (see the appendix for a description of the estimation method), which leads to $\theta_v = 0.08$, $\kappa_v = 2.18$, $\sigma_v = 0.53$, $\rho = -0.7$, and $\sqrt{V(t)} = 0.15$. Recall on that day, $S(t) = 460$ and $r = 5\%$. Substituting these values into the discretized version of Equations and (8) and (9), we simulate changes in $S(t)$ and $C(t)$ [or $P(t)$] over a given intraday interval 10,000 times, where the call and put prices are determined according to Equation (10) and the put-call parity. Based on simulated pairs $(\Delta S, \Delta C)$, we calculate the occurrence frequencies of type I and type A through type D violations and report the results in Table 10 for different intraday intervals. Note that the simulated option-price changes exhibit a type I violation between 11.1% and 11.5% of the time, for both calls and puts and for intraday intervals ranging from 30 minutes to 1 day. These simulated type I occurrence rates mimic the corresponding ones in Table 3. The SV model-generated type A through type D violation occurrence rates are higher in general than those reported in Table 4. Clearly, simulated type B violations are more likely to occur than simulated type A violations, a pattern consistent with Table 4. However, based on the simulated SV model prices, the relative ordering between type C and type D occurrence rates is not as clear-cut, suggesting that even the SV model cannot be perfectly consistent with the documented option price patterns.

4.2 Marginal explanatory power of volatility

The focus so far has been on violations of the left- and right-hand inequalities in Equation (5). In so doing, we have not taken into account the fact that option prices should change not only in the right direction but also by the right amount (as dictated by a specific model). Therefore examining the quantitative fit by each model should help shed light on how much a two-state variable option pricing model can add beyond the one-dimensional diffusion class. For this purpose, we choose the familiar BS formula and the previously discussed SV model as the respective representatives of the two model classes and compare their quantitative fit of option price changes.

Before we examine its promise quantitatively, it makes sense to describe what we mean by a violation under the SV model. If the observed dC and the SV model predicted dC do not share the same sign, then the observation is deemed inconsistent with its predictions. This is closest in principle to a type I violation under the one-dimensional class. In transitioning from the predictions of the one-dimensional diffusion models, four possibilities need to be contrasted: (1) the call is a type I violation under one-dimensional models, but not a violation under the SV model; (2) the call is not a type I violation under one-dimensional models, but is a violation under the SV model; (3) the call remains a violation under both models; and (4) the call is not a violation under either model. By assessing the respective frequencies, we can get some idea about the empirical potential of the volatility state variable. Applying the parameters and $V(t)$ values implied by *all* the S&P 500 call prices at the beginning of each sampling interval, we calculated the time t value of each option delta. From the implied-volatility time series, the changes, ΔV , are determined accordingly and applied to every call option. These frequencies at the hourly sampling interval are 6.5%, 5.3%, 7.4%, and 80.8%.

It is noted that 6.5% of the observations are classified as type I violations, but are nonviolations under the SV. Stated differently, 47% of the type I violations (i.e., 6.5/13.9) are in compliance with the SV model predictions. Nevertheless, we discover that 5.3% of the call observations are *not* a type I violation, but do violate SV model prediction. As most of these observations (84%) are type IV violations originally, we can conclude that market microstructure factors are hindering the SV model's ability to characterize empirical option pricing dynamics as well. In the hourly call sample, 7.4% of the observations violate both one-dimensional and SV model predictions. Among observations *not* a type I, II, III, or IV violation, only a minute fraction (i.e., 0.8% of the total sample) are inconsistent with the SV model. In addition, the SV model is incapable of explaining any type II violations. Thus the SV model explains about half of type I violations, but it cannot be expected to be completely consistent with the option data.

Keeping the hourly call sample as the focal point of discussion, now return to the quantitative fit of the BS and SV models. Adopting the dis-

cretized dynamics of Equation (11) and setting $C_V \equiv 0$, we run the following regression:

$$\Delta C(t, \tau_n, K_n) = \beta_0 + \beta_1 [C_S^{bs} \Delta S] + \epsilon(t, \tau_n, K_n), \quad (12)$$

where n stands for the n th call option in the sample, coefficient β_0 reflects each option's deterministic component, and $C_S^{bs} \Delta S$ is the contract-specific call price change in response to ΔS and as predicted by the BS model. The BS delta for each option is calculated using the volatility implied by *all call options* at the beginning of a given hourly interval and obtained by minimizing the sum of squared pricing errors for all calls. Under the BS hypothesis, it should hold that $\beta_1 = 1$ (and $\beta_0 \approx 0$), with the R^2 equal to one. Several features can be noted from Table 11. First, the β_1 estimate is always positive, with its magnitude ranging from 0.25 (for short-term OTM calls) to 0.87 (for short-term ITM calls). Based on the reported magnitudes and the associated standard errors, the null hypothesis $\beta_1 \equiv 1$ is overwhelmingly rejected. Of each given maturity, the BS fits ITM calls better than both OTM and ATM options. Still, the fact that each β_1 estimate is statistically far from one indicates that the BS model is overall inconsistent with the data. Second, the estimates for β_0 are all close to zero. This further substantiates the claim that during an intraday time interval, the option time decay is small. Finally, the adjusted R^2 for the regressions is between 13% (for OTM options) and 73% (for ITM options). For the full sample, the BS model can explain 51% of the observed call price changes from hour to hour and the estimated β_1 is 0.73.

To examine the SV model quantitatively, we adopt the following discretized version:

$$\Delta C(t, \tau_n, K_n) = \beta_0 + \beta_1 [C_S \Delta S] + \beta_2 [C_V \Delta V] + \epsilon(t, \tau_n, K_n), \quad (13)$$

where $[C_S \Delta S]$ and $[C_V \Delta V]$ reflect the respective option price changes induced by ΔS and ΔV , as determined by the SV model. Under the assumptions of the SV model, it should hold that $\beta_0 \approx 0$ and $\beta_1 = \beta_2 = 1$, with the regression $R^2 = 1$. As seen in Table 11, several patterns for the SV are similar to those for the BS model. For example, the β_0 estimates are insignificantly different from zero. Of a given maturity, the calls lead to a higher β_1 estimate, and a higher R^2 value, as the moneyness goes from OTM, to ATM, and to ITM, indicating a better fit for ITM calls by the SV model. Compared to the R^2 values under the BS model, the corresponding ones under the SV are generally higher, showing a better fit by the latter model. In the full-sample case, the R^2 increases from the BS model's 51% to the SV's 59%, a 16% increase in relative terms. Improvements of a smaller magnitude by the

Table 11
Quantitative fit of hourly call price changes

		BS model			SV model				F-test
		β_0	β_1	R^2	β_0	β_1	β_2	R^2	(P-value)
Full sample		0.00 (0.00)	0.73 (0.00)	0.51	0.00 (0.00)	0.80 (0.00)	0.41 (0.01)	0.59	(< 0.01)
OTM	Short	-0.01 (0.01)	0.25 (0.02)	0.16	-0.01 (0.01)	0.20 (0.01)	0.15 (0.05)	0.20	(< 0.01)
	Medium	-0.01 (0.01)	0.31 (0.02)	0.13	-0.01 (0.01)	0.29 (0.03)	0.10 (0.04)	0.15	(0.01)
	Long	0.00 (0.01)	0.42 (0.03)	0.19	-0.01 (0.01)	0.41 (0.03)	0.01 (0.05)	0.20	(0.83)
ATM	Short	0.00 (0.00)	0.79 (0.01)	0.59	0.01 (0.01)	0.86 (0.01)	0.48 (0.02)	0.62	(< 0.01)
	Medium	-0.01 (0.01)	0.57 (0.01)	0.35	-0.01 (0.01)	0.56 (0.01)	0.11 (0.03)	0.36	(< 0.01)
	Long	0.01 (0.01)	0.66 (0.01)	0.47	0.01 (0.01)	0.68 (0.02)	0.23 (0.04)	0.48	(< 0.01)
ITM	Short	0.01 (0.01)	0.87 (0.01)	0.64	0.01 (0.01)	1.00 (0.01)	1.65 (0.18)	0.68	(< 0.01)
	Medium	0.01 (0.01)	0.82 (0.01)	0.73	0.01 (0.01)	0.93 (0.01)	0.68 (0.03)	0.76	(< 0.01)
	Long	0.01 (0.01)	0.75 (0.02)	0.58	0.01 (0.01)	0.87 (0.02)	0.74 (0.04)	0.62	(< 0.01)

Respectively for the Black–Scholes (denoted BS) and the stochastic volatility (denoted SV) models, the reported results are based on the regression specifications below:

$$\Delta C(t, \tau_n, K_n) = \beta_0 + \beta_1 \left[C_S^{bs} \Delta S \right] + \epsilon(t, \tau_n, K_n),$$
$$\Delta C(t, \tau_N, K_n) = \beta_0 + \beta_1 \left[C_S \Delta S \right] + \beta_2 [C_V \Delta V] + \epsilon(t, \tau_n, K_n),$$

where $\Delta C(t, \tau_n, K_n)$, $\Delta S(t)$, and $\Delta V(t)$ are a stand-in for the changes in the n th call price, the cash index, and the spot volatility. The sampling frequency is set to 1 hour. The standard errors (reported in parentheses) are based on White's heteroscedasticity-consistent estimator. OTM, ATM, and ITM stand for out of the money, at the money, and in the money, respectively. Short-, medium-, and long-term refer to options with less than 60 days, with 60–180 days, and with more than 180 days, respectively, to expiration. The regression results are based on the entire sample and not on the violation sample alone. The reported R^2 is the adjusted R^2 . The null hypothesis of the F -test in $H_0 : \beta_2 = 0$, which is based on conditional sum of squares for comparison of full and reduced models.

SV model are observed as well when the various moneyness maturity-based subsamples are employed in the regressions. Based on the p -values from conditional sum of squares F -tests, we strongly reject the null hypothesis that the coefficient on volatility is zero for the full sample and for eight (of the nine) moneyness maturity subsamples.

The marginal contribution by the volatility variable can also be seen from two other perspectives. First, the β_1 estimates under the SV are in general closer to unity than under BS. This is particularly true for ITM calls. Second, the β_2 estimates are persistently positive and significant for each subsample (except for long-term OTM calls), indicating that volatility is a significant

factor in explaining changes in call prices. These estimates are closer to the hypothesized value of unity for ITM calls, but less so for OTM calls.⁷

Innovations in the underlying asset price are, therefore, by far the most important and they are responsible for about 51% of option price changes, option price innovations induced by volatility fluctuations are also significant, and they can explain an additional 8% of option price changes. Still, about 40% of intraday option price variations remain unaccounted for by the SV model.

One possible explanation for this unsatisfactory performance by the SV model is that volatility does not change sufficiently intraday. To see this point, we divide the hourly call price changes that each constitute a type I violation into two groups: (i) $\Delta S < 0$ but $\Delta C > 0$, and (ii) $\Delta S > 0$ but $\Delta C < 0$. For each observation in group (i), we substitute the observed ΔS , ΔC , the SV model's C_S and C_V into Equation (13) and set $\beta_0 = 0$ (ignoring time decay), $\beta_1 = \beta_2 = 1$ and $\epsilon(t) = 0$, to solve for the required volatility change, ΔV . Then for the SV model to explain all the type I violations in group (i), the average level of required volatility change is about 1%. On the other hand, the actual implied volatility (based on the SV model) only increases, on average, by 0.25% from hour to hour and for all the observations in group (i). Conducting the same calculations for group (ii), we find the average level of required volatility change to be about -1%, but in actuality the implied volatility only has an average hourly decline of 0.25% for the observations in group (ii). Therefore in order for the SV model to explain all type I violations, volatility on average needs to vary more than the implied volatility does.

5. Issues of Robustness

The option price patterns documented in preceding sections are robust even under different test designs. To examine the robustness issue, we have partitioned the call option sample into two subperiods: March 1–May 31, 1994, and June 1–August 31, 1994. For both subperiods, we have found the violation patterns to be similar to those displayed in Tables 3 and 4. At the hourly sampling interval, for instance, type I and type II violations occur at 13.1% and 21.60% of the time in the first subperiod, and 14.70% and 24.5% over the later subperiod. Joint put-call violations are also consistent between the subperiods and the full sample: (i) type A through type D violations combined account for 12.0% of the sample in the first subperiod, and 11.7%

⁷ As the BS and the SV models do not have the power to explain any type II violations, we repeated our regression analysis by omitting such observations. The premise is that violations driven by market microstructure are model independent and will not require any state variable (the spot price or otherwise). Confirming this insight, we observed that the regression R^2 for the BS model and the SV model increased by approximately the same amount (i.e., 10%). Barring this departure from the full-sample counterpart, all of our earlier conclusions remained invariant. For this reason and to save on space, the results based on this sample are not reported.

in the second; and (ii) it is still more likely for put-call pairs to go down together than to go up. Again, these conclusions hold whether the S&P 500 futures price or the spot index is used.

When the empirical analysis in Table 8 is repeated for the two subperiods, the respective magnitudes of changes in the underlying and option prices are found to be similar. For example, for those medium-term ATM call observations with $\Delta S < 0$ and $\Delta C > 0$, we find the average call-price change to be \$0.38 in the first subperiod and \$0.34 in the second subperiod, both changes are quite close (these results are available upon request). Further, we have examined the quantitative fit of option price changes by the BS and SV models. The regression R^2 and coefficient estimates based on the two subperiods are close to the ones reported in Table 11. Our empirical findings are therefore robust to the sample period choice.

6. Concluding Remarks

Until recently, much of the focus in option pricing has been on models within the one-dimensional diffusion class. In any such option pricing model, however, volatility can at most change, and hence be perfectly correlated, with the underlying price. As a result, most properties of the Black–Scholes model still hold under a general one-dimensional diffusion setting. In particular, call prices should be monotonically increasing and put prices monotonically decreasing in the underlying asset price. Also, option prices are perfectly correlated with each other and with the underlying asset, making option contracts redundant securities. The monotonicity property and perfect correlation property thus impose a stringent constraint on how option prices can change with the underlying price. Our empirical investigation indeed demonstrates that the predictions by these models are often violated.

Our analysis points to a particular role to be played by market microstructure factors and the second state variable in the modeled underlying price process. Microstructure effects (e.g., bid-ask spread, tick size restriction, etc.) can explain why option prices do not change while the underlying price changes, and why the adjustment in option prices can be larger than that in the underlying. On the other hand, an additional state variable can potentially explain why call prices move in the opposite direction with the underlying asset price. A specific example in the family of two-factor models is the Heston's (1993) stochastic volatility option pricing model, in which volatility is not necessarily perfectly correlated with the underlying price. With this imperfectly correlated state variable affecting option prices, option contracts are no longer redundant, hence option prices can move independently of the underlying price. Indeed, our simulation exercise demonstrates that the SV model-based option price changes exhibit patterns qualitatively similar to the documented type I and type A through type D violations. Further, using the

SV model parameters, we find that 47% of type I violations become consistent with the prediction of the SV model in terms of the sign of option price changes. Thus the volatility variable offers some needed flexibility to better explain option price dynamics.

Appendix

The characteristic functions for the Heston (1993) SV model are as presented below:

$$f_1 = \exp \left\{ -\frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v\tau})}{2\xi_v} \right) \right] \right. \\ \left. - \frac{\theta_v}{\sigma_v^2} [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v] \tau + i\phi r\tau + i\phi \ln[S(t)] \right. \\ \left. + \frac{i\phi(i\phi + 1)(1 - e^{-\xi_v\tau}) V(t)}{2\xi_v - [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v\tau})} \right\}, \quad (A1)$$

$$f_2 = \exp \left\{ -\frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^*\tau})}{2\xi_v^*} \right) \right] \right. \\ \left. - \frac{\theta_v}{\sigma_v^2} [\xi_v^* - \kappa_v + i\phi\rho\sigma_v] \tau + i\phi r\tau + i\phi \ln[S(t)] \right. \\ \left. + \frac{i\phi(i\phi - 1)(1 - e^{-\xi_v^*\tau}) V(t)}{2\xi_v^* - [\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^*\tau})} \right\}. \quad (A2)$$

Under the Heston model, the partial derivatives (or deltas) of the call price with respect to $S(t)$ and $V(t)$ are as follows:

$$C_S = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi} f_1(t, \tau; \phi)}{i\phi} \right] d\phi \quad (A3)$$

$$C_V = S(t) \left\{ \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi}}{i\phi} \frac{\partial f_1}{\partial V} \right] d\phi \right\} \\ - K e^{-r\tau} \left\{ \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi}}{i\phi} \frac{\partial f_2}{\partial V} \right] d\phi \right\}. \quad (A4)$$

Following a standard practice [e.g., Bakshi, Cao, and Chen (1997)], we back out the SV model parameters and the spot volatility from observed option prices. The estimation procedure is as follows. Let $\Phi \equiv \{\kappa_v, \theta_v, \sigma_v, \rho\}$. Then at each time t (the beginning of each hourly sampling interval) we collect N (> 5) S&P 500 option prices and choose values for Φ and $V(t)$ to minimize the sum of squared pricing errors:

$$\min_{V(t), \Phi} \sum_{n=1}^N |\hat{C}(t, \tau_n, K_n) - C(t, \tau_n, K_n)|^2, \quad (A5)$$

where $\hat{C}(t, \tau_n, K_n)$ is the observed price for the n th call, $C(t, \tau_n, K_n)$ is its model price as determined in Equation (10), and $S(t)$ is the dividend-adjusted spot index. The estimated values for Φ and $V(t)$ are next applied to calculate the time t partial derivatives for every call option included in the regression in Equation (A1) and only for the time interval starting at t . These steps are repeated for each hourly interval of the sample period.

The implied volatility for the Black–Scholes is similarly estimated using all call options, except that in this case only the spot volatility is backed out of the option prices.

References

- Amin, K., and V. Ng, 1993, "Option Valuation with Systematic Stochastic Volatility," *Journal of Finance*, 48, 881–910.
- Amin, K., and V. Ng, 1997, "Inferring Future Volatility from the Information in Implied Volatility in Eurodollar Options: A New Approach," *Review of Financial Studies*, 10, 333–368.
- Bakshi, G., C. Cao, and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, 2003–2049.
- Bakshi, G., C. Cao, and Z. Chen, 2000, "Pricing and Hedging Long-Term Options," *Journal of Econometrics*, 94, 277–318.
- Bates, D., 2000, "Post-87 Crash Fears in S&P 500 Futures Options," *Journal of Econometrics*, 94, 145–180.
- Bergman, Y., B. Grundy, and Z. Wiener, 1996, "General Properties of Option Prices," *Journal of Finance*, 51, 1573–1610.
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.
- Chan, K. C., W. Christie, and P. Schulz, 1995, "Market Structure and the Intraday Pattern of Bid-Ask Spreads for NASDAQ Securities," *Journal of Business*, 68, 35–60.
- Cox, J., and S. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145–166.
- Derman, E., and I. Kani, 1994, "Riding on a Smile," *Risk*, 7(2), 32–39.
- Dumas, B., J. Fleming, and R. Whaley, 1998, "Implied Volatility Smiles: Empirical Tests," *Journal of Finance*, 53, 2059–2106.
- Fleming, J., B. Ost diek, and R. Whaley, 1996, "Trading Costs and the Relative Rates of Price Discovery in the Stock, Futures, and Options Markets," *Journal of Futures Markets*, 16, 353–387.
- Glosten, L., R. Jagannathan, and D. Runkle, 1993, "On the Relation Between the Expected Value and Volatility of Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779–1802.
- Heston, S., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–343.
- Kamara, A., and T. Miller, 1995, "Daily and Intradaily Tests of European Put-Call Parity," *Journal of Financial and Quantitative Analysis*, 30, 519–539.
- Kroner, K., and V. Ng, 1998, "Modeling Asymmetric Comovements of Asset Returns," *Review of Financial Studies*, 11, 817–844.
- Merton, R., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics*, 4, 141–183.
- Rubinstein, M., 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Options Classes from August 23, 1976 Through August 31, 1978," *Journal of Finance*, 40, 455–480.
- Rubinstein, M., 1994, "Implied Binomial Trees," *Journal of Finance*, 49, 771–818.
- Wood, R., T. McInish, and K. Ord, 1985, "An Investigation of Transactions Data for NYSE Stocks," *Journal of Finance*, 40, 723–739.