

Mispricing of S&P 500 Index Options

George M. Constantinides

University of Chicago and NBER

Jens Carsten Jackwerth

University of Konstanz

Stylianos Perrakis

Concordia University

Widespread violations of stochastic dominance by 1-month S&P 500 index call options over 1986–2006 imply that a trader can improve expected utility by engaging in a zero-net-cost trade net of transaction costs and bid-ask spread. Although precrash option prices conform to the Black-Scholes-Merton model reasonably well, they are incorrectly priced if the distribution of the index return is estimated from time-series data. Substantial violations by postcrash OTM calls contradict the notion that the problem lies primarily with the left-hand tail of the index return distribution and that the smile is too steep. The decrease in violations over the postcrash period of 1988–1995 is followed by a substantial increase over 1997–2006, which may be due to the lower quality of the data but, in any case, does not provide evidence that the options market is becoming more rational over time.

(*JEL* G10, G13)

A robust prediction of the celebrated Black and Scholes (1973) and Merton (1973) (BSM) option pricing model is that the volatility implied by market prices of options is constant across strike prices. Rubinstein (1994) tested this prediction on the S&P 500 index options (SPX) traded on the Chicago Board Options Exchange, an exchange that comes close to the dynamically complete and perfect market assumptions underlying the BSM model. From the start of the exchange-based trading in April 1986 until the October 1987 stock market crash, the implied volatility is a moderately downward-sloping or U-shaped

We thank workshop participants at the German Finance Society Meetings 2004; the Bachelier 2004 and 2006 Congresses; the FMA 2005 European Conference; the EFA 2005 Meetings; the Frontiers of Finance Conference 2006; the Alberta/Calgary 2006 Conference; the Universities of Chicago, City, Concordia, Frankfurt, Iowa, Laval, Maryland, Princeton, Southern California, Svizzera Italiana, Texas-Austin, Torino, St. Gallen, Utah, and New York; London Business School; London School of Economics; and, in particular, Yacine Ait-Sahalia, David Bates, Duke Bristow, Larry Harris, Steve Heston, Jim Hodder, Mark Loewenstein, Matthew Richardson, Jeffrey Russell, Hersh Shefrin, Greg Willard, and the referees for their insightful comments and constructive criticism. We also thank Michal Czerwonko for excellent research assistance. We remain responsible for errors and omissions. Constantinides acknowledges financial support from the Center for Research in Security Prices of the University of Chicago, and Perrakis from the Social Sciences and Humanities Research Council of Canada. Send correspondence to George M. Constantinides, Graduate School of Business, University of Chicago, 5807 S. Woodlawn Avenue, Chicago, IL 60637, telephone: 773-702-7258. E-mail: gmc@ChicagoGSB.edu.

function of the strike price, a pattern referred to as the “volatility smile,” also observed in international markets and to a lesser extent in the prices of individual stock options. Following the crash, the volatility smile is typically more pronounced and downward sloping, often called “volatility skew.”¹

An equivalent statement of the above prediction of the BSM model, that the volatility implied by market prices of options is constant across strike prices, is that the *risk-neutral* stock price distribution is lognormal. Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (1998), Jackwerth (2000), and Ait-Sahalia and Duarte (2003) estimated the risk-neutral stock price distribution from the cross section of option prices.² Jackwerth and Rubinstein (1996) confirmed that, prior to the October 1987 crash, the risk-neutral stock price distribution is close to lognormal, consistent with a moderate implied volatility smile. Thereafter, the distribution is systematically skewed to the left, consistent with a more pronounced skew.

These findings raise important questions. Does the reasonable fit of the BSM model prior to the crash imply that options were rationally priced prior to the crash? Why does the BSM model typically fail after the crash? Were options priced rationally after the crash?

Whereas downward-sloping or U-shaped implied volatility is inconsistent with the BSM model, it is well understood that this pattern is not necessarily inconsistent with economic theory. Two fundamental assumptions of the BSM model are that the market is *dynamically complete* and *frictionless*. We empirically investigate whether the observed cross sections of 1-month S&P 500 index option prices over 1986–2006 are consistent with various economic models that explicitly allow for a dynamically incomplete market and also an imperfect market that recognizes trading costs and bid-ask spreads. To our knowledge, this is the first large-scale empirical study that addresses mispricing in the presence of transaction costs and intermediate trading.

The absence of arbitrage in a frictionless market implies the existence of a risk-neutral probability measure, not necessarily unique, such that the price of any asset equals the expectation of its payoff under the risk-neutral measure, discounted at the risk-free rate. If a risk-neutral measure exists, the ratio of the risk-neutral probability density to the real probability density, discounted at the risk-free rate, is referred to as the *pricing kernel* or the *stochastic discount factor*. Thus, the absence of arbitrage implies the existence of a strictly positive pricing kernel.

¹ The shortcomings of the BSM model are addressed in the context of *no-arbitrage* models that generalize the stock price process by including stock price jumps and stochastic volatility and also generalize the processes for the risk premia. Many of these models are critically discussed in Jackwerth (2004), Hull (2006), McDonald (2006), and Singleton (2006).

² Jackwerth (2004) reviews the parametric and nonparametric methods for estimating the risk-neutral distribution. Ait-Sahalia and Duarte (2003) estimate the implied risk neutral distribution from a sample of simultaneously expiring European index option prices while constraining the option pricing function to be monotonic and convex. This approach may be extended to the estimation of the pricing kernel also.

Economic theory imposes restrictions on equilibrium models beyond merely ruling out arbitrage. In a frictionless representative agent economy with von Neumann-Morgenstern preferences, the pricing kernel equals the representative agent's intertemporal marginal rate of substitution over each trading period. If the representative agent has *state-independent* (derived) utility of wealth, then the concavity of the utility function implies that the pricing kernel is a decreasing function of wealth. Furthermore, if the representative agent's wealth at the end of each period is monotone increasing in the stock return over the period, then the pricing kernel is a decreasing function of the market return.

The monotonicity restriction on the pricing kernel does not critically depend on the existence of a representative agent. If there is not at least one pricing kernel that is a decreasing function of wealth over each trading period, then there is not even *one* economic agent with state-independent and concave (derived) utility of wealth and with wealth at the end of each period that is monotone increasing in the stock return over the period that is a marginal investor in the market. Hereafter, we employ the term *stochastic dominance violation* to connote the nonexistence of *even one* economic agent with the above attributes that is marginal in the market.³ This means that, if such an economic agent exists, then the return on the current portfolio is stochastically dominated (in the second degree) by the return of another feasible portfolio.

Under the two maintained hypotheses that the marginal investor's (derived) utility of wealth is state independent and wealth is monotone increasing in the market index level, the pricing kernel is a decreasing function of the market index level. Ait-Sahalia and Lo (2000), Jackwerth (2000), and Rosenberg and Engle (2002) estimated the pricing kernel implied by the observed cross section of prices of S&P 500 index options as a function of wealth, where wealth is proxied by the S&P 500 index level. Jackwerth reported that the pricing kernel is everywhere decreasing during the precrash period, 1986–1987, but widespread violations occur over the postcrash period, 1987–1995. Ait-Sahalia and Lo examined the year 1993 and reported violations; Rosenberg and Engle examined the 1991–1995 period and reported violations.⁴

Several extant models address the inconsistencies with the BSM model and the violations of monotonicity of the pricing kernel. While not all of these models explicitly address the monotonicity of the pricing kernel, they do address the problem of reconciling option prices with the time-series properties of the index returns. Essentially, these models introduce additional priced state variables and/or explore alternative specifications of preferences.⁵ These models

³ This line of research was initiated by Perrakis and Ryan (1984), Levy (1985), and Ritchken (1985). For more recent related contributions, see Perrakis (1986, 1993), Ritchken and Kuo (1988), Ryan (2000, 2003), and Oancea and Perrakis (2007).

⁴ Rosenberg and Engle (2002) found violations when they used an orthogonal polynomial pricing kernel but not when they used a power pricing kernel which, by construction, is decreasing in wealth.

⁵ These models are critically discussed in Singleton (2006). Bates (2006) introduced heterogeneous agents with utility functions that explicitly depend on the number of stock market crashes, over and above their dependence

are suggestive but stop short of *endogenously* generating the process of the risk premia associated with these state variables in the context of an *equilibrium* model of the macro economy and explaining on a month-by-month basis the cross section of the S&P 500 option prices.

In estimating the statistical distribution of the S&P 500 index returns, we refrain from adopting the BSM assumption that the index price is a Brownian motion and, therefore, that the arithmetic returns on the S&P 500 index are lognormal.

We do not impose a parametric form on the distribution of the index returns and proceed in four different ways. In the first approach, we estimate the unconditional distribution as the (smoothed) histograms extracted from two different *historical* index data samples covering the 1928–1986 and 1972–1986 periods. In the second approach, we estimate the unconditional distribution as the histograms extracted from two different *forward-looking* samples: one that includes the October 1987 crash (1987–2006) and another that excludes it (1988–2006). In the third approach, we model the variance of the index returns as a generalized GARCH (1,1) process and scale the unconditional distribution for each month to have the above variance. Finally, in the fourth approach, we scale the unconditional distribution for each month to have standard deviation equal to the Black-Scholes implied volatility (IV) of the closest ATM option or, alternatively, equal to the Chicago Board Options Exchange Volatility Index (VIX) (1990–2006 only). Clearly, we have not exhausted all possible ways of estimating the statistical distribution of the S&P 500 index returns. One interpretation of our empirical results regarding mispricing is simply that the options market is priced with a different probability distribution than any of our estimated probability distributions.

We test the compliance of option prices to the predictions of a model that allows for market incompleteness, market imperfections, and intermediate trading over the life of the options. We consider a market with heterogeneous economic agents and investigate the restrictions on option prices imposed by a particular class of utility-maximizing agents that we simply refer to as traders. We assume that traders maximize state-independent increasing and concave

on the agent's terminal wealth. The calibrated economy exhibits the inconsistencies with the BSM model but fails to generate the nonmonotonicity of the pricing kernel. Brown and Jackwerth (2004) suggested that the reported violations of the monotonicity of the pricing kernel may be an artifact of the maintained hypothesis that the pricing kernel is state independent but concluded that volatility cannot be the sole omitted state variable in the pricing kernel. Garcia, Luger, and Renault (2003); Santa-Clara and Yan (2004); Brennan, Liu, and Xia (2006); and Christoffersen, Heston, and Jacobs (2006), among others, obtained plausible parameter estimates in models in which the pricing kernel is state dependent, using panel data on S&P 500 options. Others calibrated equilibrium models that generate a volatility smile pattern observed in option prices. David and Veronesi (2002) modeled the investors' learning about fundamentals, calibrated their model to earnings data, and provided a close fit to the panel of prices of S&P 500 options. Liu, Pan, and Wang (2005) investigated rare-event premia driven by uncertainty aversion in the context of a calibrated equilibrium model and demonstrated that the model generates a volatility smile pattern observed in option prices. Benzoni, Collin-Dufresne, and Goldstein (2007) extended the above approach to show that uncertainty aversion is not a necessary ingredient of the model. They also demonstrated that the model can generate the stark regime shift that occurred at the time of the 1987 crash. Alternative explanations include *buying pressure*, suggested by Bollen and Whaley (2004), and behavioral explanations based on *sentiment*, suggested by Shefrin (2005) and Han (2006).

utility functions and that each trader's wealth at the end of each period is monotone increasing in the stock return over the period. For example, an investor who holds 100 shares of stock and a net short position in 200 call options violates the monotonicity condition, while an investor who holds 200 shares of stock and a net short position in 200 call options satisfies the condition. Essentially, we assume that the traders have a sufficiently large investment in the stock, relative to their net short position in call options, such that the monotonicity condition is satisfied.

We do not make the restrictive assumption that all economic agents belong to the class of utility-maximizing traders. Thus, our results are robust and unaffected by the presence of agents with beliefs, endowments, preferences, trading restrictions, and transaction costs schedules that differ from those of the utility-maximizing traders modeled in this paper.

Whereas we assume that returns are *i.i.d.* and that traders have state-independent preferences, we also conduct tests that relax these assumptions and accommodate three implications associated with state dependence. First, each month we search for a pricing kernel to price the cross section of 1-month options without imposing restrictions on the time-series properties of the pricing kernel month by month. Second, we allow for intermediate trading; a trader's wealth on the expiration date of the options is generally a function not only of the price of the market index on that date but also of the entire path of the index level, thereby rendering the pricing kernel state dependent. Third, we allow the variance of the index return to be state dependent and employ the forecasted conditional variance.

The paper is organized as follows. In Section 1, we present a model for pricing options and state restrictions on the prices of options imposed by the absence of stochastic dominance violations. One form of these restrictions is a set of linear inequalities on the pricing kernel that can be tested by testing the feasibility of a linear program. The second form of these restrictions is a pair of upper and lower bounds on the prices of options. In Section 2, we test the compliance of bid and ask prices of 1-month index call options to these restrictions and discuss the results. In Section 3, we summarize the empirical findings and suggest directions for future research.

1. Restrictions on Option Prices Imposed by Stochastic Dominance

1.1 The model and assumptions

Trading occurs on a finite number of dates, $t = 0, 1, \dots, T, \dots, T'$. The utility-maximizing traders are allowed to hold and trade only two primary securities in the market, a stock and a bond. The stock has the natural interpretation as the market index. The bond is risk free and pays constant interest $R - 1$ each period. The traders may buy and sell the bond without incurring transaction costs. On date t , the *cum dividend* stock price is $(1 + \delta_t)S_t$, the cash dividend

is $\delta_t S_t$, and the *ex dividend* stock price is S_t , where δ_t is the dividend yield. We assume that the rate of return on the stock, $(1 + \delta_{t+1})S_{t+1}/S_t$, is identically and independently distributed over time.

Stock trades incur proportional transaction costs charged to the bond account as follows. On each date t , the trader pays $(1 + k)S_t$ out of the bond account to purchase one *ex dividend* share of stock and is credited $(1 - k)S_t$ in the bond account to sell (or sell short) one *ex dividend* share of stock. We assume that the transaction costs rate satisfies the restriction $0 \leq k < 1$.

On date 0, the utility-maximizing traders are also allowed to buy or sell J European call options that mature on date T .⁶ Also, on this date, a trader can buy the j th option at price $P_j + k_j$ and sell it at price $P_j - k_j$, net of transaction costs. Thus $2k_j$ is the bid-ask spread plus the round-trip transaction costs that the trader incurs in trading the j th option.

On each date, a trader chooses the investment in the bond, stock, and call options to maximize the expected utility of net worth at the terminal date T' . We make the plausible assumption that utility is state independent and is increasing and concave in net worth. Later on, we relax the assumption of state independence.

One may formulate this problem as a dynamic program. As shown in Constantinides (1979), the value function is monotone increasing and concave in the dollar values in the bond and stock accounts, properties that it inherits from the monotonicity and concavity of the utility function. This implies that, on any date, the *marginal utility of wealth out of the bond account* is strictly positive and decreasing in the dollar value of the bond account, and the *marginal utility of wealth out of the stock account* is strictly positive and decreasing in the dollar value of the stock account. Furthermore, as shown in Fama (1970), the joint assumptions that (1) the rate of return on the stock is identically and independently distributed over time and (2) utility is state independent ensure that the state space on each date is defined solely by the stock return realizations without additional state variables.

Finally, we assume that each trader's wealth at the end of each period is monotone increasing in the stock return over the period. Essentially, we assume that the traders do not write naked calls: they have sufficiently large investment in the stock, relative to their net short position in call options such that the monotonicity condition is satisfied. The implication of the monotonicity condition is that a trader's marginal utility of wealth out of the stock account is strictly positive and decreasing in the *stock return*.

Our model assumptions are weaker than the assumptions made in the derivation of the capital asset pricing model. Thus, our model implies that the pricing kernel is monotone decreasing in the index return but, unlike the capital asset

⁶ In this paper, we empirically investigate only 1-month call options. We first investigate the case where trading is allowed only once per month by setting $T = 1$ and considering the time between trading dates to be of calendar length 1 month. Later on, we investigate the case where trading is allowed N times per month by setting $T = N$ and considering the time between trading dates to be of calendar length $1/N$ months.

pricing model, does not necessarily imply that the pricing kernel is *linearly* decreasing in the index return.

We search for marginal utilities with the above properties that support the prices of the bond, stock, and derivatives at a given point in time. If we fail to find such a set of marginal utilities, then any trader (as defined in this paper) can increase expected utility by trading in the options, the index, and the risk-free rate—hence equilibrium does not exist. These strategies are termed *stochastically dominant* for the purposes of this paper, insofar as they would be adopted by all traders, in the same way that all risk-averse investors would choose a dominant portfolio over a dominated one in conventional second-degree stochastic dominance comparisons. Stochastic dominance then implies that at least one agent, *but not necessarily all agents*, increases expected utility by trading. In our empirical investigation, we report the percentage of months for which the problem is feasible. These are months for which stochastic dominance violations are ruled out.

1.2 Restrictions in the single-period model

We specialize the general model by setting $T = 1$. We do not rule out trading after the options expire; we just rule out trading over the 1-month life of the options. In Section 1.3, we consider the more realistic case in which traders are allowed to trade the bond and stock at one intermediate date over the life of the options.

As stated earlier, the joint assumptions that the rate of return on the stock is identically and independently distributed over time and utility is state independent ensure that the state space at the options' maturity is defined solely by the stock return realizations and not by additional state variables.⁷ Furthermore, the joint assumptions that utility is concave and wealth is increasing in the stock return (the monotonicity condition) ensure that the marginal utility of wealth out of the stock account is strictly positive and decreasing in the *stock return*.

We specialize the general notation as follows. The stock market index has price S_0 at the beginning of the period; *ex dividend* price S_{1i} with probability π_i in state i , $i = 1, \dots, I$, at the end of the period; and *cum dividend* price $(1 + \delta)S_{1i}$ at the end of the period. We order the states such that S_{1i} is increasing in i .

We define $M^B(0)$ as the marginal utility of wealth out of the bond account at the beginning of the period; $M^S(0)$ as the marginal utility of wealth out of the stock account at the beginning of the period; $M_i^B(1)$ as the marginal utility of wealth out of the bond account at the end of the period; and $M_i^S(1)$ as the marginal utility of wealth out of the stock account at the end of the period. The marginal utility of wealth out of the bond and stock accounts at the beginning

⁷ One may replace the assumption of *i.i.d.* returns with the assumption that the investment horizon ends on date one, $T = T' = 1$.

of the period is strictly positive:

$$M^B(0) > 0 \quad (1)$$

and

$$M^S(0) > 0. \quad (2)$$

The marginal utility of wealth out of the bond account at the end of the period is strictly positive:⁸

$$M_i^B(1) > 0, \quad i = 1, \dots, I. \quad (3)$$

The marginal utility of wealth out of the stock account at the end of the period is strictly positive and decreasing in the stock return:

$$M_1^S(1) \geq M_2^S(1) \geq \dots \geq M_I^S(1) > 0. \quad (4)$$

On each date, the trader may transfer funds between the bond and stock accounts and incur transaction costs. Therefore, the marginal rate of substitution between the bond and stock accounts differs from unity by, at most, the transaction cost rate:

$$(1 - k) M^B(0) \leq M^S(0) \leq (1 + k) M^B(0) \quad (5)$$

and

$$(1 - k) M_i^B(1) \leq M_i^S(1) \leq (1 + k) M_i^B(1), \quad i = 1, \dots, I. \quad (6)$$

Marginal analysis on the bond holdings leads to the following condition on the marginal rate of substitution between the bond holdings at the beginning and end of the period:

$$M^B(0) = R \sum_{i=1}^I \pi_i M_i^B(1), \quad (7)$$

where R is one plus the risk-free rate. Marginal analysis on the stock holdings leads to the following condition on the marginal rate of substitution between the stock holdings at the beginning of the period and the bond and stock holdings at the end of the period:

$$M^S(0) = \sum_{i=1}^I \pi_i \left[\frac{S_{1i}}{S_0} M_i^S(1) + \frac{\delta S_{1i}}{S_0} M_i^B(1) \right]. \quad (8)$$

⁸ Since the value of the bond account at the end of the period is independent of the state i , we cannot impose the condition that the marginal utility of wealth out of the bond account is decreasing in the dollar value of the bond account.

Marginal analysis on the option holdings leads to the following condition on the marginal rate of substitution between the bond holdings at the beginning of the period and the bond holdings at the end of the period:

$$(P_j - k_j)M^B(0) \leq \sum_{i=1}^I \pi_i M_i^B(1) X_{ij} \leq (P_j + k_j)M^B(0), \quad j = 1, \dots, J. \quad (9)$$

Each month in our empirical analysis, we check for the feasibility of conditions (1)–(9) by using the linear programming features of the optimization toolbox of MATLAB 7.0. We report the percentage of months in which the conditions are feasible and, therefore, stochastic dominance is ruled out.

1.3 Restrictions in the two-period model

We relax the assumption of the single-period model that, over the 1-month life of the options, markets for trading are open only at the beginning and end of the period; we allow for a third trading date in the middle of the month. We define the marginal utility of wealth out of the bond account and out of the stock account at each one of the three trading dates and set up the linear program as a direct extension of the program (1)–(9) in Section 1.2. The explicit program is given in Appendix A. In our empirical analysis, we report the percentage of months in which the conditions are feasible and, therefore, stochastic dominance is ruled out.

1.4 Restrictions in the multiperiod model

In principle, we may allow for more than one intermediate trading date over the 1-month life of the options. However, the numerical implementation becomes tedious as both the number of constraints and variables in the linear program increase exponentially with the number of intermediate trading dates.

Constantinides and Perrakis (2002) derived testable implications of the absence of stochastic dominance that are invariant to the allowed frequency of trading the bond and stock over the life of the options. This generality is achieved under the assumption that the trader's universe of assets consists of the bond, stock, and a 1-month call option with a certain strike price. Specifically, Constantinides and Perrakis derived an upper and a lower bound to the price of a call option of given strike and maturity. The bounds have the following interpretation. If one can buy the option for less than the lower bound, then there is a stochastic dominance violation between the bond, stock, and the given option. Likewise, if one can write the option for more than the upper bound, then again there is a stochastic dominance violation between the bond, stock, and the given option.⁹

⁹ These bounds on call prices (and the corresponding bounds on put prices) may not be the tightest possible bounds for any given frequency of trading. However, they are presented here because of their universality in

Below, we state without proof the bounds on call options. At any time t prior to expiration, the following is a partition-independent upper bound on the price of a call:

$$\bar{c}(S_t, t) = \frac{(1+k)}{(1-k)R_S^{T-t}} E [[S_T - K]^+ | S_t], \quad (10)$$

where R_S is the expected return on the stock per period.

A lower bound for a call option can also be found, but only if it is additionally assumed that there exists at least one trader for whom the investment horizon coincides with the option expiration. In such a case, transaction costs become irrelevant in the put-call parity, and the following is a lower bound:¹⁰

$$\underline{c}(S_t, t) = (1+\delta)^{t-T} S_t - K/R^{T-t} + E[(K - S_T)^+ | S_t] / R_S^{T-t}, \quad (11)$$

where R is one plus the risk-free rate per period. We present the upper and lower bounds in Figures 1–4 and discuss their violations in Section 2.6.

2. Empirical Results

2.1 Data and estimation

We use the historical daily record of the S&P 500 index and its daily dividend record over the 1928–2006 period, as explained in Appendix B. The monthly index return is based on 30 calendar day (21 trading day) returns. In order to avoid difficulties with the estimated historical mean of the returns, we demean all our samples and reintroduce a mean 4% annualized *premium* over the risk-free rate. Our results remain practically unchanged if we do not make this adjustment because the prices of 1-month options are insensitive to the expected return on the stock.

We estimate both the *unconditional* and the *conditional* distribution of the index. The *unconditional* distribution is extracted from four alternative samples of 30-day index returns: two *historical* returns samples over the 1928–1986 and 1972–1986 periods; a *forward-looking* returns sample over the 1987–2006 period that includes the 1987 stock market crash; and a *forward-looking* returns sample over the 1988–2006 period that excludes the stock market crash. The

that they do not depend on the frequency of trading over the life of the option. For a comprehensive discussion and derivation of these and other, possibly tighter, bounds that are specific to the allowed frequency of trading, see Constantinides and Perrakis (2002). Constantinides and Perrakis (2007) provided bounds for American-style options and futures options. These bounds were tested with data on S&P 500 futures options by Constantinides et al. (2007), who identified options violating the bounds and derived strategies exploiting these mispricings. For alternative approaches to option bounds under transaction costs, see also Leland (1985), Bensaid et al. (1992), and Constantinides and Zariphopoulou (1999, 2001).

¹⁰ In the special case of zero transaction costs, the assumption $T = T'$ is redundant because the put-call parity holds.

annualized volatilities are 21.0% (1928–1986), 15.8% (1972–1986), 15.2% (1987–2006), and 14.8% (1988–2006).

For each sample, we use a discrete state space of 61 values from $e^{-0.60}$ to $e^{0.60}$, spaced 0.02 apart in log spacing. This span covers all observed returns in any of our samples. We use the standard Gaussian kernel of Silverman (1986, pp. 15, 43, and 45). The resulting probabilities for different states can vary greatly in scale and cause numerical problems in solving the resulting LPs. We thus resort to eliminating states with probabilities smaller than 0.00001 and rescaling the remaining probabilities (typically 99.998%) to sum to 1.

We estimate the *conditional* distribution of the index each month over the 1972–2006 period in three different ways: by GARCH (1,1), as the implied volatility (IV), and as the *revised VIX* index.¹¹ In the first way, we apply the semi-parametric GARCH (1,1) method of Engle and Gonzalez-Rivera (1991), a method that does not impose the restriction that conditional returns are normally distributed, as explained in Appendix C.¹² In the second way, we estimate the conditional volatility as the Black-Scholes IV of the closest ATM 1-month option and scale the unconditional distribution every month to match the conditional volatility. In the third way, we estimate the conditional volatility as 0.01 times the *revised VIX* index and scale the unconditional distribution every month over the 1990–2006 subperiod (when the VIX is available) to match the conditional volatility.¹³

For the S&P 500 index options, we use two data sources. For the 1986–1995 period, we use the tick-by-tick Berkeley Options Database of all quotes and trades. We focus on the most liquid call options that have a K/S ratio (moneyness) in the 0.90–1.05 range. For 107 months we retain only the call option quotes for the day corresponding to options with 30 days to expiration.¹⁴ For each day retained in the sample, we aggregate the quotes to the minute and pick the minute between 9:00 and 11:00 AM with the most quotes as our cross section for the month. We present these quotes in terms of their bid and

¹¹ An alternative Wall Street approach of obtaining a conditional distribution is to bootstrap from 22-day overlapping returns with a rolling window of several months where each day is the beginning of another 22-day return. 100 days is a common choice. We are indebted to an anonymous referee of this journal for pointing out this alternative approach.

¹² The index return sample and the option price sample do not align. We use the conditional volatility of the 30-day return period, which starts *before* the option sample and covers it partly at the beginning. We recalculated the results by using the conditional volatility of the 30-day return period, which starts *during* the option sample and covers it partly at the end and then continues beyond the option sample. The two sets of results are practically indistinguishable, and thus we do not report the latter results here.

¹³ The scaling of the distribution by the conditional volatility does not change the skewness and kurtosis. However, our numerical implementation does cause these moments to vary slightly due to slightly different discretizations. For the 30-day conditional index return distribution, 1972–2006, based on GARCH (1,1), the average skewness is -0.415765 and kurtosis is 1.693435 ; based on implied volatilities, the average skewness is -0.393942 and kurtosis is 1.650050 .

¹⁴ We lose 9 months for which we do not have sufficient data, i.e., months with less than five different strike prices, months after the crash of October 1987 until June 1988, and months before the introduction of S&P 500 index options in April 1986.

ask implied volatilities. Details on this database are provided in Appendix B, Jackwerth and Rubinstein (1996), and Jackwerth (2000).

We do not have options data for 1996. For the 1997–2006 period, we obtain call option bid and ask prices from the Option Metrics Database, described in Appendix B. We calculate a hypothetical noon option cross section from the closing cross section and the index observed at noon and the close. Here we assume that the implied volatilities do not change between noon and the close. We start out with 109 raw cross sections and are left with 108 final cross sections. The time to expiration is 29 days.

Since the Berkeley Options Database provides less noisy data than the Option Metrics Database, we expect a higher incidence of stochastic dominance violations over the 1997–2006 period than over the 1986–1995 period. Thus we are cautious in comparing results across these two periods.

2.2 Assumptions on bid-ask spreads and trading fees

There is no presumption that all agents in the economy face the same bid-ask spreads and transaction costs as the traders do. We assume that the traders are subject to the following bid-ask spreads and trading fees. For the index, we model the combined one-half bid-ask spread and one-way trading fee as a one-way proportional transaction costs rate equal to 50 bps of the index price.

For the options, we model the combined one-half bid-ask spread and one-way trading fee either as *fixed* or as *proportional* transaction costs. Under the fixed-costs regime, we set the fixed transaction costs equal to 5, 10, or 20 bps of the index price. This corresponds to about 19, 38, or 75 cents as one-way fee per call, respectively. Fixed transaction costs probably overstate the actual transaction costs on OTM calls and understate them on ITM calls.

Under the proportional-costs regime, the proportional transaction costs for an ATM call are set equal to the transaction costs under the fixed-costs regime. However, for an OTM (or ITM) call with a price equal to a fraction (or multiple) x of the price of the ATM call, the proportional transaction costs are equal to a fraction (or multiple) x of the transaction costs of the ATM call. Proportional transaction costs probably understate the actual transaction costs on OTM calls and overstate them on ITM calls. In the tables, we present results under both fixed-cost and proportional-cost regimes.

2.3 Stochastic dominance violations in the single-period case

Each month we check for the feasibility of conditions (1)–(9). Infeasibility of these conditions implies stochastic dominance: *any* trader can improve utility by trading in these assets without incurring any out-of-pocket costs. If we rule out bid-ask spreads and trading fees, we find that these conditions are violated in all months. Thus, we introduce bid-ask spreads and trading fees as described in Section 2.2.

The time series of option prices is divided into seven periods, and stochastic dominance violations in each period are reported in different columns,

Table 1A
Percentage of months without stochastic dominance violations with (A) proportional transaction costs and (B) fixed transaction costs

	A: 8605– 8710	B: 8807– 9103	C: 9104– 9308	D: 9309– 9512	E: 9702– 9912	F: 0002– 0305	G: 0306– 0605
(A) Proportional transaction costs ^a							
Number of months	15	29	28	26	35	37	36
Unconditional index return distribution, 1928–1986	67 73 (5)	24 76 (10)	21 50 (11)	0 4 (5)	6 26 (9)	11 27 (6)	6 17 (4)
	80	90	82	35	60	49	39
Unconditional index return distribution, 1972–1986	27 53	28 48	32 54	0 15	3 9	5 8	3 11
	67	76	93	81	29	19	47
Unconditional index return distribution, 1987–2006	13 20	38 59	43 68	0 35	0 9	5 11	3 14
	33	76	96	88	14	16	47
Unconditional index return distribution, 1988–2006	7 20	38 55	39 71	0 27	0 0	5 8	3 11
	47	66	93	81	20	16	61
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	13 27	34 59	61 82	42 69	0 11	8 27	6 11
	33	76	86	96	29	43	47
Conditional index return distribution, 1972–2006, based on implied vol.	53 87	55 83	71 96	42 73	6 29	14 38	3 19
	100	93	86	96	71	84	50
Conditional index return distribution, 1972–2006, based on VIX	N/A	N/A	61 86	23 54	6 17	8 24	3 19
			89	92	51	86	58

^aThe panel displays the percentage of months in which stochastic dominance violations are absent in the cross section of option prices. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are proportional. In each row, the one-way transaction costs rate on the ATM calls is 5 bps of the index price (top entries), 10 bps (bold, middle entries), or 20 bps (bottom entries). The numbers inside parentheses in the first row are bootstrap standard deviations of the first-row entries, based on 200 runs.

labeled as panels A–G in Tables 1–4. The first period extends from May 1986 to October 16, 1987, just prior to the stock market crash. The other six periods are all postcrash and span July 1988 to March 1991, April 1991 to August 1993, September 1993 to December 1995, February 1997 to December 1999, February 2000 to May 2003, and June 2003 to May 2006. Note that we do not have options data for 1996 from either data source. The average annualized implied volatility is 0.1641, and the panel averages are 0.1753 (A), 0.179 (B), 0.1307 (C), 0.1089 (D), 0.2 (E), 0.2173 (F), and 0.1228 (G).

In Table 1A, the one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are proportional. The entries display the percentage of months without stochastic dominance violations. In each row, the one-way transaction costs rate on the ATM calls is 5 bps of the index price (top entries), 10 bps (bold middle entries), or 20 bps (bottom entries). The number of calls in each (filtered) monthly cross section fluctuates between 5 and 23 with a median of 10. The bracketed numbers in the first row are bootstrap standard deviations of the first-row middle entries (10 bps transaction costs), based

Table 1B
A percentage of months without stochastic dominance violations with (A) proportional transaction costs
and (B) fixed transaction costs

	A: 8605– 8710	B: 8807– 9103	C: 9104– 9308	D: 9309– 9512	E: 9702– 9912	F: 0002– 0305	G: 0306– 0605
(B) Fixed transaction costs ^b							
Number of months	15	29	28	26	35	37	36
Unconditional index return distribution, 1928–1986	47 67 87	7 41 86	0 7 46	0 0 0	9 31 74	3 35 51	0 11 25
Unconditional index return distribution, 1972–1986	33 53 60	21 52 72	4 32 68	4 4 15	3 11 29	8 16 24	3 14 53
Unconditional index return distribution, 1987–2006	13 40 47	38 55 66	21 54 82	0 8 38	6 20 31	11 11 16	8 17 58
Unconditional index return distribution, 1988–2006	13 33 47	31 55 69	21 43 71	0 4 35	0 11 17	5 11 11	6 17 53
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	20 40 40	38 52 79	36 61 89	42 85 100	3 17 34	8 24 38	6 19 58
Conditional index return distribution, 1972–2006, based on implied vol.	67 100 100	48 79 93	54 89 96	62 88 96	14 37 74	19 38 78	6 28 72
Conditional index return distribution, 1972–2006, based on VIX	N/A	N/A	21 75 93	15 54 96	6 17 46	3 19 73	0 28 72

^bThe panel displays the percentage of months in which stochastic dominance violations are absent in the cross section of option prices. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are fixed. In each row, the one-way transaction costs rate on the calls is 5 bps of the index price (top entries), 10 bps (bold, middle entries), or 20 bps (bottom entries).

on 200 samples of the 1928–1986 historical returns. The standard deviations are small and, therefore, comparisons of the table entries across the rows and columns can be made with some confidence. We need to be careful when comparing across panels A–D (Berkeley Options Database) and panels E–G (Option Metrics).

Most table entries are well below 100%, indicating that there are a number of months in which the risk-free rate, the price of the index, and the prices of the cross section of calls are inconsistent with a market in which there is even *one* risk-averse trader who is marginal in these securities, net of generous transaction costs.

In the top-left cell, the middle entry of 73% refers to the index return distribution over the 1928–1986 period and option prices over the precrash period from May 1986 to October 16, 1987. In 27% of these months, conditions (1)–(9) are infeasible and the prices imply stochastic dominance violations despite the generous allowance for transaction costs. The next six entries to the right, panels B–G, refer to call prices over the six postcrash periods. Violations increase in panels C–G. In panel D, all but 4% of the cross sections violate the stochastic dominance restrictions. The option prices in panel D are drawn

Table 2A
A percentage of months without stochastic dominance violations with (A) proportional transaction costs—ITM and OTM calls separately and (B) fixed transaction costs—ITM and OTM calls separately

	A: 8605– 8710	B: 8807– 9103	C: 9104– 9308	D: 9309– 9512	E: 9702– 9912	F: 0002– 0305	G: 0306– 0605
(A) Proportional transaction costs—ITM and OTM calls separately ^a							
Number of months	15	29	28	26	35	37	36
Unconditional index return distribution, 1928–1986	87	90	79	38	69	46	42
Unconditional index return distribution, 1972–1986	73	79	79	19	29	30	17
Unconditional index return distribution, 1987–2006	73	52	89	77	26	16	53
Unconditional index return distribution, 1972–1986	53	45	75	46	11	14	11
Unconditional index return distribution, 1988–2006	53	66	89	88	20	16	50
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	20	59	89	62	9	11	17
Conditional index return distribution, 1972–2006, based on implied vol.	53	62	93	77	17	14	47
Conditional index return distribution, 1972–2006, based on VIX	20	59	89	54	3	8	14
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	53	69	86	100	31	38	44
Conditional index return distribution, 1972–2006, based on implied vol.	27	59	79	69	11	30	11
Conditional index return distribution, 1972–2006, based on VIX	100	97	93	100	80	81	50
Conditional index return distribution, 1972–2006, based on implied vol.	87	83	96	77	34	59	19
Conditional index return distribution, 1972–2006, based on VIX	N/A	N/A	96	96	69	76	56
Conditional index return distribution, 1972–2006, based on VIX			93	69	34	49	19

^aThe panel displays the percentage of months in which stochastic dominance violations are absent in the cross section of ITM calls (top entry) and OTM calls (bottom entry). The one-way transaction costs rate on the index is 50 bps. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are proportional; for the ATM calls they are 10 bps of the index price.

Table 2B
A percentage of months without stochastic dominance violations with (A) proportional transaction costs—ITM and OTM calls separately and (B) fixed transaction costs—ITM and OTM calls separately

	A: 8605– 8710	B: 8807– 9103	C: 9104– 9308	D: 9309– 9512	E: 9702– 9912	F: 0002– 0305	G: 0306– 0605
(B) Fixed transaction costs—ITM and OTM calls separately ^b							
Number of months	15	29	28	26	35	37	36
Unconditional index return distribution, 1928–1986	80	55	7	0	54	43	11
Unconditional index return distribution, 1972–1986	80	90	86	27	57	38	39
Unconditional index return distribution, 1987–2006	60	59	39	4	14	16	14
Unconditional index return distribution, 1972–1986	67	66	93	81	29	19	58
Unconditional index return distribution, 1988–2006	47	62	54	15	17	14	22
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	47	62	93	96	23	16	44
Conditional index return distribution, 1972–2006, based on implied vol.	40	59	50	15	17	11	22
Conditional index return distribution, 1972–2006, based on VIX	40	62	96	92	20	16	33
Conditional index return distribution, 1972–2006, based on GARCH (1,1)	40	59	64	88	29	30	28
Conditional index return distribution, 1972–2006, based on implied vol.	47	62	82	96	31	41	53
Conditional index return distribution, 1972–2006, based on VIX	100	93	89	88	77	65	44
Conditional index return distribution, 1972–2006, based on implied vol.	100	93	100	96	69	81	61
Conditional index return distribution, 1972–2006, based on VIX	N/A	N/A	79	77	26	38	33
Conditional index return distribution, 1972–2006, based on VIX			93	96	60	76	64

^bThe panel displays the percentage of months in which stochastic dominance violations are absent in the cross section of ITM calls (top entry) and OTM calls (bottom entry). The one-way transaction costs rate on the index is 50 bps. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are fixed as 10 bps of the index price.

Table 3
Percentage of months without stochastic dominance violations using conditional implied-volatility-based index return distributions with $\pm 2\%$ offset

	A: 8605–8710	B: 8807–9103	C: 9104–9308	D: 9309–9512	E: 9702–9912	F: 0002–0305	G: 0306–0605
Number of months	15	29	28	26	35	37	36
Implied vol. -2%	13	55	71	50	0	5	0
Implied vol. -1%	47	72	93	69	29	30	6
Implied vol.	87	83	96	73	29	38	19
Implied vol. $+1\%$	93	76	96	65	26	32	19
Implied vol. $+2\%$	100	72	86	65	23	30	19
Best of above	100	83	96	73	29	38	19

The table displays the percentage of months in which stochastic dominance violations are absent in the cross section of option prices. The one-way transaction costs rate on the index is 50 bps. The one-way transaction costs on the options are proportional. The one-way transaction costs rate on the ATM calls is 10 bps of the index price. All results use the conditional implied-volatility-based index return distribution over the 1972–2006 period. Four offsets are used to change the implied ATM volatility by -2 , -1 , 1 , or 2% annualized return. The bold results “Best of above” count a monthly cross section as feasible if feasibility is established either without implied volatility offset or with any of the four offsets.

Table 4
Percentage of months without stochastic dominance violations in the two-period case

	A: 8605– 8710	B: 8807– 9103	C: 9104– 9308	D: 9309– 9512	E: 9702– 9912	F: 0002– 0305	G: 0306– 0605
Number of months	15	29	28	26	35	37	36
Unconditional index return distribution, 1928–1986	60 (73)	52 (66)	39 (46)	8 (0)	17 (23)	24 (22)	14 (14)
Unconditional index return distribution, 1972–1986	33 (53)	41 (48)	43 (54)	19 (15)	11 (9)	3 (8)	8 (11)
Unconditional index return distribution, 1987–2006	13 (27)	45 (48)	46 (61)	23 (27)	3 (0)	0 (8)	8 (8)
Unconditional index return distribution, 1988–2006	13 (20)	45 (55)	57 (75)	31 (31)	6 (6)	5 (8)	8 (8)
Conditional index return distribution, 1972–2006, Based on GARCH (1,1)	20 (40)	34 (41)	50 (82)	58 (58)	3 (9)	14 (11)	6 (6)
Conditional index return distribution, 1972–2006, based on implied vol.	87 (93)	66 (66)	61 (75)	54 (58)	14 (14)	27 (30)	8 (14)
Conditional index return distribution, 1972–2006, based on VIX	N/A	N/A	64 (75)	42 (50)	9 (9)	14 (16)	8 (11)

The table displays the percentage of months in which stochastic dominance violations are absent in the cross section of option prices *when one intermediate trading date is allowed* over the life of the 1-month options. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are proportional; for the ATM calls they are 10 bps of the index price. In parentheses, the table displays the percentage of months in which stochastic dominance violations are absent in the case when *no intermediate trading is allowed* over the life of the 1-month options. Two periods of 15 days are used.

from the reliable Berkeley Options Database. The high incidence of violations cannot be attributed to data problems.

We investigate the robustness of the historical estimate of the index return distribution over the 1928–1986 period by reestimating the historical distribution of the index return over the more recent period of 1972–1986. The incidence of violations increases in all panels except in panels C and D, where it decreases but remains high.

When we use the forward-looking index sample 1987–2006 that includes the stock market crash (third row) or the forward-looking index sample 1988–2006 that excludes it (fourth row), the precrash options exhibit substantially more violations. Our interpretation is that, before the crash, option traders were using average historical volatility to price options and were not actively forecasting volatility changes. This interpretation is reinforced in row 5, panel A. The GARCH method in forecasting volatility does worse than the first two rows and only marginally better than the third and fourth rows.

In the last three rows, we use the GARCH-based, the IV-based, and the revised VIX-based conditional index distribution, all based on index returns over 1972–2006, as explained in Section 2.1. Of these three methods, GARCH is the only one that, in the spirit of this paper, uses information exclusively from the time series of index returns to impose restrictions on the prices of options. By contrast, the IV-based and the revised VIX-based methods use the volatility implied in the option prices themselves, irrespective of whether this volatility is rational or not. In particular, implied volatility tends to be higher than realized volatility.

The IV-based method performs better than the revised VIX method. This is surprising because the revised VIX is meant to be a theoretically motivated refinement of the IV method. The IV-based method performs well in pricing precrash options. Nevertheless, violations with the IV-based method remain surprisingly severe, particularly over the 1997–2006 period.

2.4 Robustness in the single-period case

Floor traders, institutional investors, and broker-assisted investors face different transaction cost schedules in trading options. Are the results robust under different transaction cost schedules? The pattern of violations remains essentially the same. In Table 1A, the number at the top of each cell is the percentage of nonviolations when the combined one-half bid-ask spread and one-way trading fee on one option are based on 5 bps of the index price. We observe a large percentage of violations for all index and option price periods. The number at the bottom of each cell is the percentage of nonviolations when the combined one-half bid-ask spread and one-way trading fee on one option are based on 20 bps of the index price. Predictably, we observe fewer violations for all index and option price periods but there are still many violations.

Table 1B displays the percentage of violations but with fixed instead of proportional transaction costs. The pattern of violations is similar to the pattern displayed in Table 1A with proportional transaction costs. In some cells, violations increase, and in others, they decrease. This is surprising because we would expect that fixed transaction costs, which imply larger transaction costs for OTM calls, would result in fewer violations across the board. This begs the question of whether it is the OTM or the ITM calls that are responsible for the majority of violations.

Table 2A displays violations by ITM calls (top entry) and OTM calls (bottom entry) under the proportional transaction costs regime. In almost all cases, there is a higher percentage of violations by OTM calls than by ITM calls.¹⁵ Table 2B displays the violations due to ITM calls (top entry) and OTM calls (bottom entry) under the fixed transaction costs regime. Since fixed transaction costs imply higher transaction costs for OTM calls than in Table 2A, the violations by OTM calls substantially decrease. Since fixed transaction costs imply lower transaction costs for ITM calls than in Table 2A, the violations by ITM calls substantially increase. In Table 2B, there are fewer violations by OTM calls than by ITM calls. The violations persist when we employ the GARCH-based conditional return distribution but substantially decrease for both the ITM and OTM calls when we employ the IV-based distribution. However, the fact remains that there are substantial violations by OTM calls. This observation is novel and contradicts the common inference drawn from the observed implied volatility smile that the problem lies primarily with the left-hand tail of the index return distribution.

We further investigate the observation made in Table 1A that, when we use the IV-based conditional index return distribution, violations remain severe. We therefore entertain the possibility that the ATM IV is a biased measure of the volatility of the index return distribution.

In Table 3, we offset the IV by -2 , -1 , 1 , or 2% annualized return. In the last row, “Best of above,” we report the maximum percentage of feasible month in each panel, either without IV offset or with *any* of the four offsets, allowing the offset to be different in each panel. The one-way transaction costs rate on the index is 50 bps. The one-way transaction costs on the options are proportional. The one-way transaction costs rate on the ATM calls is 10 bps of the index price. All results use the conditional implied-volatility-based index return distribution over the sample period 1972–2006. In the precrash sample, violations disappear if we increase the implied volatility by 2% , consistent with received wisdom that sample volatility is lower than implied volatility. In the other subperiods, violations persist even under the “Best of above” category. This is surprising because this heavy-handed adjustment of the IV lacks theoretical justification and is explicitly designed to eliminate violations. Furthermore, it is no longer consistently the case that sample volatility is lower than implied volatility.

2.5 Stochastic dominance violations in the two-period model

In the previous sections, we considered feasibility in the context of the single-period model. We established that there are stochastic dominance violations in a significant percentage of the months. Does the percentage of stochastic

¹⁵ This inference is moderated by the fact that the sample of OTM calls is larger than the sample of ITM calls. Other things equal, the larger the sample, the harder it is to find a monotone decreasing pricing kernel that prices the calls. However, the figures (discussed later in Section 2.6) are not subject to this reservation and are consistent with the observation that there is substantial mispricing of OTM calls.

dominance violations increase or decrease as the allowed frequency of trading in the stocks and bonds over the life of the option increases? In the special case of zero transaction costs, *i.i.d.* returns, and constant relative risk aversion, it can be theoretically shown that the percentage of violations should increase as the allowed frequency of trading increases. However, we cannot provide a theoretical answer if we relax any of the above three assumptions. Therefore, we address the question empirically.

We compare the percentage of stochastic dominance violations in two models: one with one intermediate trading date over the 1-month life of the calls and another with no intermediate trading dates over the life of the calls. To this end, we partition the 30-day horizon into two 15-day intervals and approximate the 15-day return distribution by a 61-point kernel density estimate of the 15-day returns. In this instance, we base our kernel method on the 15-day returns instead of the 30-day returns. The assumed transaction costs are as in the base case presented in Table 1A. The one-way transaction costs rate on the index is 50 bps. The transaction costs on the options are proportional; for the ATM calls, they are 10 bps of the index price. The results are presented in Table 4.

We cannot investigate the effect of intermediate trading by directly comparing the results in Tables 1A and 4 because the return-generating process differs in the two tables since the time horizons are different. Recall that the results in Table 1A are based on a state space of 61 values for 30-day returns. By contrast, the results in Table 4 are based on a state space of 61 values of the 15-day returns. The 30-day return then is the product of two 15-day returns treated as *i.i.d.* With this process of the 30-day return, we calculate the percentage of months without stochastic dominance violations and report the results in Table 4 in parentheses.

The effect of allowing for one intermediate trading date over the life of the 1-month options is shown by the top entries in Table 4. These entries are contrasted with the bracketed entries, which represent the percentage of months without stochastic dominance violations when intermediate trading is forbidden.¹⁶ In most cases, intermediate trading increases the incidence of violations. We conclude that allowance for intermediate trading strengthens the earlier systematic evidence of stochastic dominance violations. In the next section, we obtain further insights on the causes of infeasibility, by displaying the options that violate the upper and lower bounds on option prices.

2.6 Stochastic dominance bounds

The stochastic dominance violations reported thus far are based on the nonexistence of a trader who is simultaneously marginal in the entire cross section of call prices at the beginning of each month. This requirement effectively

¹⁶ However, we find that the middle entries in Table 1A are rather similar to the bracketed entries in Table 4. This is an indication that the 30-day return used in Table 1A can be reasonably well approximated by convoluting two 15-day returns as in Table 4, bracketed entries.

rules out the possibility that the call options market is *segmented*. We entertain the possibility of segmented markets by examining violations of stochastic dominance through violations of the stochastic dominance bounds (10)–(11) discussed in Section 1.4. These bounds are derived from the perspective of a trader who is marginal in the index, the risk-free rate, and *only one call option at a time*. Therefore, these bounds allow for the possibility that the market is segmented. A second advantage of examining violations through these bounds is that the bounds apply irrespective of the permitted frequency of trading in the bond and stock accounts over the life of the option.

We calculate these bounds and translate them into bounds on the implied volatility of option prices. In Figures 1–4, we present the upper IV bound based on (10) and the lower bound based on (11). We present both the bid (circles) and ask (crosses) option prices, translated into IVs. A violation occurs whenever an observed call bid price lies above the upper bound or an observed call ask price lies below the lower bound.

In Figures 1 and 2, panels A–D are based on the Berkeley Options Database, 1986–1995; in Figures 3 and 4, panels E–G are based on the Option Metrics Database, 1997–2006. In all cases, the transaction costs rate on the index is 50 bps.

The bounds are based on the conditional index return distribution. First, we estimate the unconditional distribution over the 1972–2006 period. Then, for each subperiod, we calculate the average IV and rescale the volatility of the unconditional distribution in each panel accordingly. Since the bounds are adjusted by the implied volatility, irrespective of whether this volatility is rational or not, we can draw inferences about the shape of the skew but not about the general level of option prices.

In Figure 1, panel A, bid prices of some OTM calls lie above the upper bound and ask prices of some ITM calls lie below the lower bound.¹⁷ These findings are consistent with the results reported in Tables 2A and B, panel A, that both OTM and ITM precrash calls violate stochastic dominance. The shape of the upper and lower bounds in Figure 1, panel A, suggests that *if call prices exhibited a smile before the crash, there would be fewer violations*. This is a novel finding because precrash option prices have been documented to follow the BSM model reasonably well and this has been interpreted as evidence that they are correctly priced.

In Figure 2, panels B–G dispel another common misconception, namely, that the observed smile is too steep after the crash. In fact, panel G illustrates that there is hardly a smile in the 2003–2006 period. Postcrash violations are due to both ITM and OTM calls; sometimes bid prices are above the upper bound and ask prices are below the lower bound. These findings are consistent with the results reported in the tables.

¹⁷ Bids with zero implied volatility (not asks, which are always positive) imply that the price is so low that there does not exist a positive implied volatility solving the Black-Scholes equation. These bids do not violate the bounds as they do not present utility-improving opportunities.

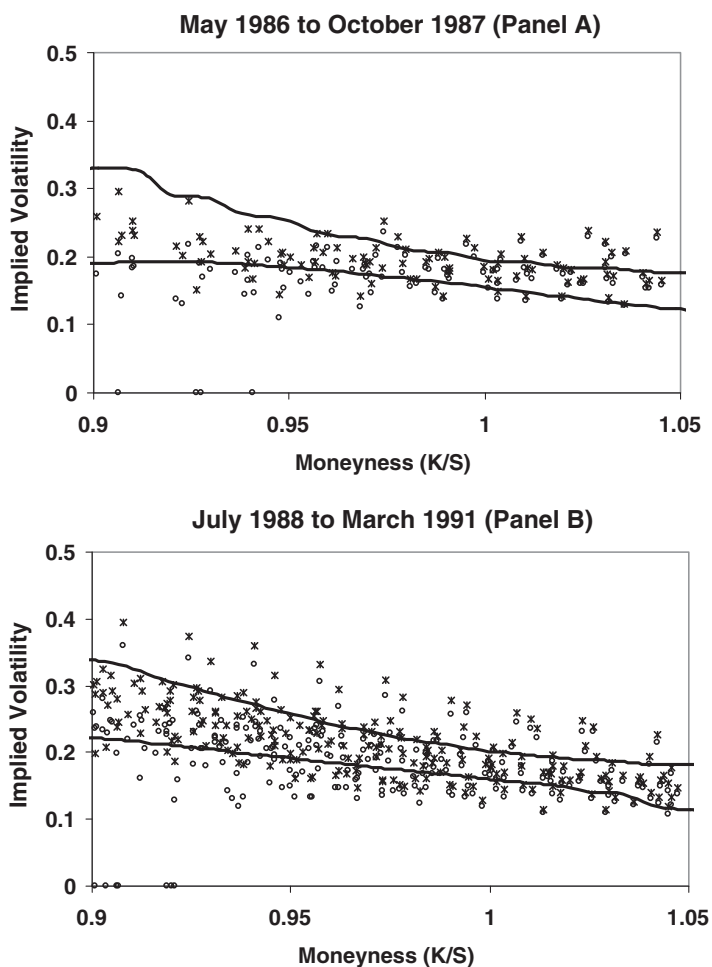


Figure 1
Bound violations over May 1986 to October 1987 and July 1988 to March 1991
 The observed bid (circles) and ask (crosses) call prices, as implied volatilities, are plotted as functions of the moneyness. The upper and lower option bounds are based on the index sample distribution for 1972–2006, rescaled with the conditional volatility of the relevant panel. The transaction costs rate on the index is 50 bps.

In Figure 2, panels C and D, there are very few violations, consistent with the results in Table 1A and B, when the conditional index return distribution is based on IV. The good performance does not carry over into subsequent periods. In Figure 3, panels E and F, several bid call prices over the 1997–2003 period lie way above the bounds. This is true for both ITM and OTM calls. This is an altogether different pattern of violations than in the earlier panels A–D. In interpreting the high incidence of violations of option prices over the 1997–2003 period in the tables, we were conservative because of concerns

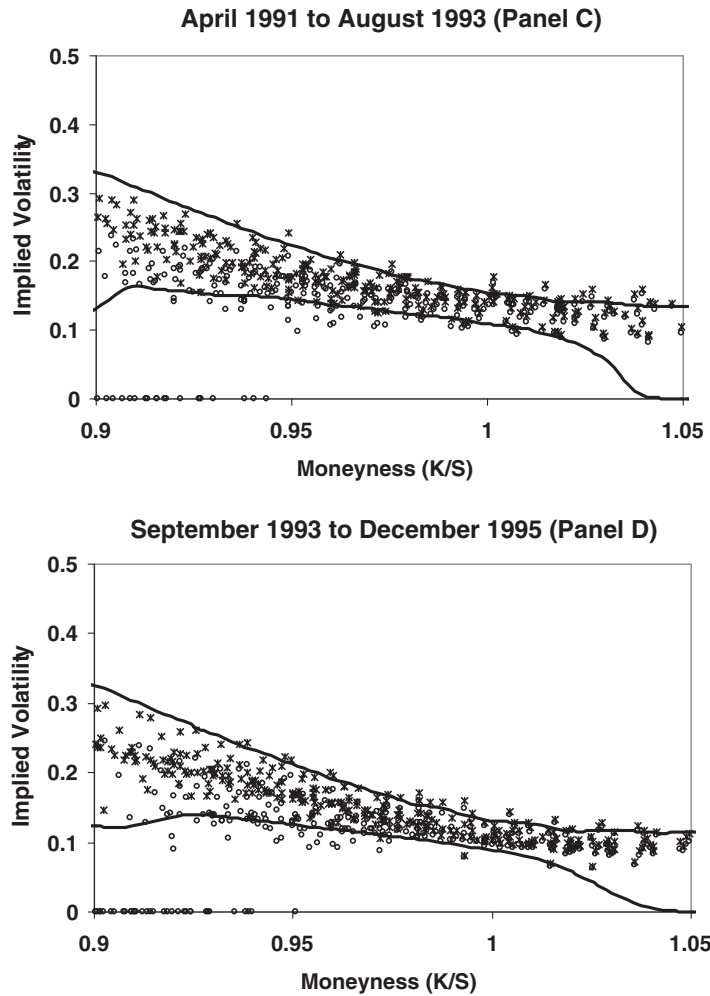


Figure 2
Bound violations over April 1991 to August 1993 and September 1993 to December 1995
 The observed bid (circles) and ask (crosses) call prices, as implied volatilities, are plotted as functions of the moneyness. The upper and lower option bounds are based on the index sample distribution for 1972–2006, rescaled with the conditional volatility of the relevant panel. The transaction costs rate on the index is 50 bps.

regarding the quality of the Option Metrics Database. The figures provide a clearer picture. If the violations were the result of low-quality data, then we would observe roughly as many violations of the lower bound as we do in the upper bound. This is not the case. Most of the violations are violations of the upper bound. The decrease in violations over the 1988–1995 postcrash period (panels B–D) is followed by a substantial increase in violations over 1997–2003 (panels E and F). This is a novel finding and casts doubts on the hypothesis that

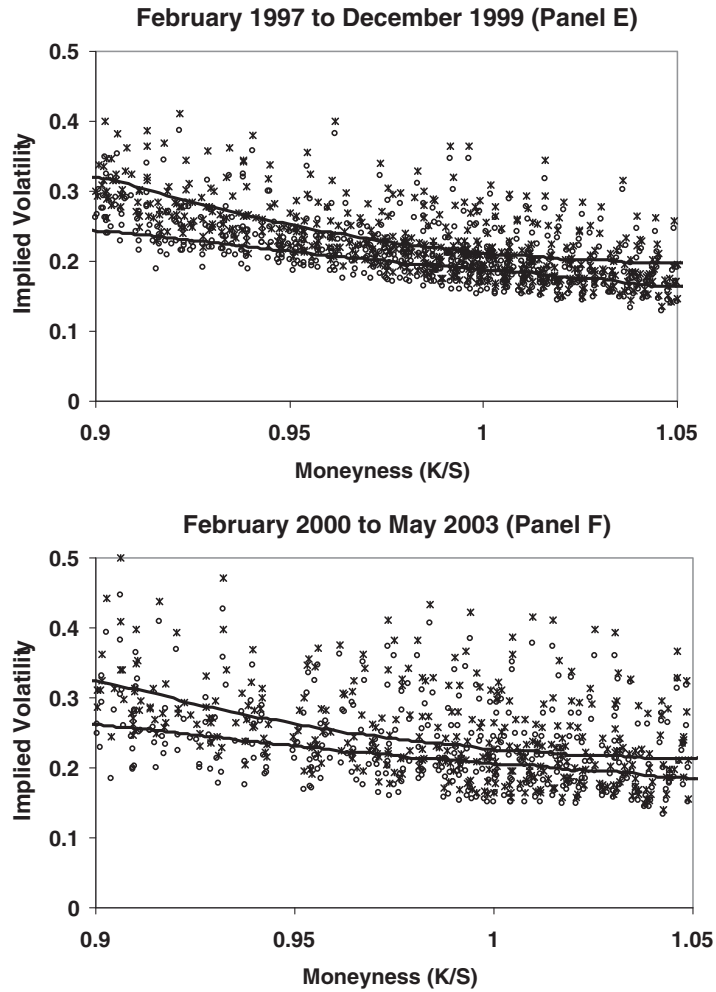


Figure 3
Bound violations over February 1997 to December 1999 and February 2000 to May 2003
 The observed bid (circles) and ask (crosses) call prices, as implied volatilities, are plotted as functions of the moneyness. The upper and lower option bounds are based on the index sample distribution for 1972–2006, rescaled with the conditional volatility of the relevant panel. The transaction costs rate on the index is 50 bps.

the options market is becoming more rational over time, particularly after the crash.

In results not reported, we estimated the bounds of Figures 1–4 with the *unconditional* distributions, both historical and forward-looking, and compared the bounds to the observed option prices. The pattern is broadly similar to the one exhibited by the conditional distributions presented in the paper. In particular, the estimated bounds exhibited a smile in the precrash period as in

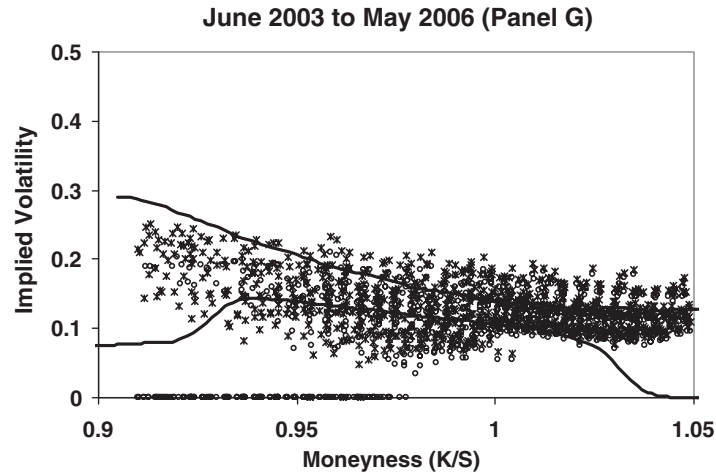


Figure 4
Bound violations over June 2003 to May 2006
The observed bid (circles) and ask (crosses) call prices, as implied volatilities, are plotted as functions of the moneyness. The upper and lower option bounds are based on the index sample distribution for 1972–2006, rescaled with the conditional volatility of the relevant panel. The transaction costs rate on the index is 50 bps.

Figure 1, panel A. We also observed a decrease in violations over the 1988–1995 postcrash period, followed by an increase in violations over the 1997–2003 period.

3. Concluding Remarks

We document widespread violations of stochastic dominance in the 1-month S&P 500 index options market over the 1986–2006 period, before and after the October 1987 stock market crash. We do not impose a parametric model on the index return distribution but estimate it as the (smoothed) histogram of the sample distribution, using seven different index return samples: two samples before the crash, one long and one short; two forward-looking samples, one that includes the crash and one that excludes it; one sample adjusted for GARCH-forecasted conditional volatility; one adjusted for implied volatility; and one sample adjusted for VIX-forecasted conditional volatility. We allow the market to be incomplete and imperfect by introducing generous transaction costs in trading the index and the options.

Evidence of stochastic dominance violations means that any trader can increase expected utility by engaging in a zero-net-cost trade. We consider a market with heterogeneous agents and investigate the restrictions on option prices imposed by a particular class of utility-maximizing economic agents that we refer to simply as traders. We do not make the restrictive assumption that all agents belong to the class of the utility-maximizing traders. Thus, our results are robust and unaffected by the presence of agents with beliefs,

endowments, preferences, trading restrictions, and transaction costs schedules that differ from those of the utility-maximizing traders modeled in this paper.

Our empirical design allows for three implications associated with state dependence. First, each month we search for a pricing kernel to price the cross section of 1-month options without imposing restrictions on the time-series properties of the pricing kernel, month by month. Thus, we allow the pricing kernel to be state dependent. Second, we allow for intermediate trading; a trader's wealth on the expiration date of the options is generally a function not only of the price of the market index on that date but also of the entire path of the index level, thereby rendering the pricing kernel state dependent. Third, we allow the volatility of the index return to be state dependent and employ the estimated conditional volatility.

Even though precrash option prices conform to the BSM model reasonably well, once the constant volatility input to the BSM formula is judiciously chosen, this does not speak on the rationality of option prices. Our novel finding is that precrash options are incorrectly priced if the distribution of the index return is estimated from time-series data even with a variety of statistical adjustments. Our derived option bounds exhibit a smile, which suggests that precrash option prices would violate these bounds less frequently if they exhibited a smile too. Our interpretation of these results is that, before the crash, option traders were extensively using the BSM pricing model and the dictates of this model were imposed on the option prices even though these dictates were not necessarily consistent with the time-series behavior of index prices.

There are substantial violations by OTM calls under both the fixed and proportional transaction costs regimes. This observation is novel and contradicts the common inference drawn from the observed implied volatility smile that the problem lies primarily with the left-hand tail of the index return distribution. We do not find evidence that the observed smile is too steep after the crash.

If the violations by ITM and OTM calls were the result of low-quality data, then we would observe roughly as many violations of the lower bound as we do of the upper bound. This is not the case. Most of the violations are violations of the upper bound. The decrease in violations over the 1988–1995 postcrash period is followed by a substantial increase in violations over the 1997–2003 period. This is a novel finding and casts doubts on the hypothesis that the options market is becoming more rational over time, particularly after the crash.

By providing an integrated approach to the pricing of options that allows for incomplete and imperfect markets, we provide testable restrictions on option prices that include the BSM model as a special case. We reviewed the empirical evidence on the prices of S&P 500 index options. The economic restrictions are violated surprisingly often, suggesting that the mispricing of these options cannot be attributed entirely to the fact that the BSM model does not allow for market incompleteness and realistic transaction costs.

In this paper, we allowed for some implications associated with *nonpriced* state variables. Several extant models addressed the inconsistencies with the

BSM model and the violations of monotonicity of the pricing kernel by introducing *priced* state variables and/or exploring alternative specifications of preferences. For example, Brennan, Liu, and Xia (2006) rejected an explanation of index option prices based on a pricing kernel that is a nonlinear function of the market return, the interest rate, and the Sharpe ratio. It remains an open and challenging topic for future research to *endogenously* generate the process of the risk premia associated with these state variables in the context of an equilibrium model of the macro economy and explain on a month-by-month basis the cross section of S&P 500 index option prices.

Our search for a trader who is simultaneously marginal in the stock, risk-free rate, and the entire cross section of 1-month call options does not address the possibility that equilibrium exists but in a segmented market. In Figures 1–4, we partially allowed for the possibility that equilibrium exists but the market is segmented by searching for a trader that is simultaneously marginal in the stock, risk-free rate, and just *one* 1-month call option at a time. Even in this case, we report several violations. In practice, individual investors (our “traders”) may face additional restrictions imposed by their brokers in writing options, beyond the restrictions that we imposed through trading costs and bid-ask spreads. It remains an open and challenging topic for future research to investigate the extent to which more severe market segmentation or imperfections can reconcile the results presented in this paper.

Appendix A

We allow for three trading dates, $t = 0, 1, 2$, at the beginning, middle, and end of the month-long period ending with the expiration of the options. We define the stock returns over the first subperiod as $z_{1i} \equiv (1 + \delta) S_{1i}/S_0$, corresponding to the I states on date 1, $i = 1, \dots, I$. We assume that the returns over the two subperiods are independent. Thus, the stock returns over the second subperiod, $z_{2k} \equiv (1 + \delta) S_{2ik}/S_{1i}$, $k = 1, \dots, I$, are independent of i . There are I^2 states on date 2, $i = 1, \dots, I$, $k = 1, \dots, I$.

We define the state-dependent marginal utility of wealth out of the bond account on each of the three trading dates as $M^B(0)$, $M_i^B(1)$, and $M_{ik}^B(2)$. Likewise, we define the state-dependent marginal utility of wealth out of the stock account on each of the three trading dates as $M^S(0)$, $M_i^S(1)$, and $M_{ik}^S(2)$. The conditions on positivity and monotonicity of the marginal utility of wealth out of the bond and stock accounts at $t = 0, 1$ are given by equations (1)–(4). The corresponding conditions at $t = 2$ are

$$M_{ik}^B(2) > 0, \quad i, k = 1, \dots, I \quad (\text{A1})$$

and

$$M_{i1}^S(2) \geq M_{i2}^S(2) \geq \dots \geq M_{ik}^S(2) \geq \dots \geq M_{iI}^S(2) > 0, \quad i = 1, \dots, I. \quad (\text{A2})$$

On each date, the trader may transfer funds between the bond and stock accounts and incur transaction costs. Conditions (5) and (6) hold. The corresponding condition at $t = 2$ is

$$(1 - k) M_{ik}^B(2) \leq M_{ik}^S(2) \leq (1 + k) M_{ik}^B(2), \quad i, k = 1, \dots, I. \quad (\text{A3})$$

Conditions (7) and (8) on the marginal rate of substitution between dates 0 and 1 hold. The corresponding conditions between dates 1 and 2 are as follows:

$$M_i^B(1) = R \sum_{k=1}^I \pi_k M_{ik}^B(2), \quad i = 1, \dots, I \quad (\text{A4})$$

and

$$M_i^S(1) = \sum_{k=1}^I \pi_k [z_{2k} M_{ik}^S(2) + \delta z_{2k} M_{ik}^B(2)], \quad i = 1, \dots, I. \quad (\text{A5})$$

Condition (9) is replaced by

$$(P_j - k_j)M^B(0) \leq \sum_{i=1}^I \sum_{k=1}^I \pi_i \pi_k M_{ik}^B(2) X_{ikj} \leq (P_j + k_j)M^B(0),$$

$$j = 1, \dots, J. \quad (\text{A6})$$

The probability of state (i,k) is $\pi_i \pi_k$ because, by assumption, the stock returns are independent over the two subperiods.

In our empirical analysis, we report the percentage of months in which conditions (1)–(8) and (A1)–(A6) are feasible and, therefore, stochastic dominance is ruled out.

Appendix B

B.1 Berkeley Options Database

The Berkeley Options Database contains all minute-by-minute quotes and trades of the European options and futures on the S&P 500 index from April 2, 1986, to December 29, 1995. Details on this database are found in Jackwerth and Rubinstein (1996), Jackwerth (2000), and below.

B.1.1 Index level. Traders typically use the index futures market rather than the cash market to hedge their option positions. The reason is that the cash market prices lag futures prices by a few minutes due to lags in reporting transactions of the constituent stocks in the index. We check this claim by regressing the index on each of the first 20-minute lags of the futures price. The single regression with the highest adjusted R^2 is assumed to indicate the lag for a given day. The median lag of the index over the 1,542 days from 1986 to 1992 is 7 minutes. Because the index is stale, we compute a futures-based index for each minute from the futures market as $S_0 = (1 + \delta)^{-1} R F$, where F is the futures price at the option expiration. For each day, we use the median interest rate R implied by all futures quotes and trades and the index level at that time. We approximate the dividend yield δ by assuming that the dividend amount and timing expected by the market were identical to the dividends actually paid on the S&P 500 index. However, some limited tests indicate that the choice of the index does not seem to affect the results of this paper.

B.1.2 Interest rate. We compute implied interest rates embedded in the European put-call parity relation. Armed with option quotes, we calculate separate lending and borrowing interest returns from put-call parity where we use the above future-based index. For each expiration date, we assign a single lending and borrowing rate to each day, which is the median of all daily observations across all strike prices. We then use the average of these two interest rates as our daily spot rate for the particular time to expiration. Finally, we obtain the interpolated interest rates from the implied forward curve. If there is data missing, we assume that the spot rate curve can be extrapolated

horizontally for the shorter and longer times-to-expiration. Again, some limited tests indicate that the results are not affected by the exact choice of the interest rate.

B.1.3 Option prices. We use only bid and ask prices on call options. For each day retained in the sample, we aggregate the quotes to the minute and pick the minute between 9:00 and 11:00 AM with the most quotes as our cross section for the month.

We use only call options with 30 days to expiration that occur once every month during our sample. We also trim the sample to allow for moneyness levels between 0.90 and 1.05. Cross sections with fewer than five option quotes are discarded. We also eliminate the cross sections right after the crash of 1987 as the data are noisy and restart the sample with the cross section expiring on July 15, 1988.

B.1.4 Arbitrage violations. In the process of setting up the database, we check for a number of errors that might have been contained in the original minute-by-minute transaction level data. We eliminate a few obvious data-entry errors, as well as a few quotes with excessive spreads—more than 200 cents for options and 20 cents for futures. General arbitrage violations are eliminated from the dataset. We also check for violations of vertical and butterfly spreads. Within each minute, we keep the largest set of option quotes that satisfies the restriction $S(1 + \delta) \geq C_i \geq \max[0, S(1 + \delta) - K_i/R]$.

Early exercise is not an issue as the S&P 500 options are European and the discreteness of quotes and trades only introduces a stronger upward bias in the midpoint implied volatilities for deep-out-of-the-money puts (moneyness less than 0.6), which we do not use in our empirical work. We start out with 107 raw cross sections and are left with 98 final cross sections.

B.2 Option Metrics Database

The Option Metrics Database contains indicative end-of-day European call and put option quotes on the S&P 500 index from January 2, 1997, to May 31, 2006. In merging the Option Metrics Database with the Berkeley Options Database, we follow the above procedure as much as possible, given the closing bid and ask prices that the Option Metrics Database provides. Therefore, only departures and innovations from the above procedure are noted.

B.2.1 Index level. As the closing (noon) index price, we use the price implied by the closing (noon) futures price.

B.2.2 Interest rate. As we cannot arrive at a consistently positive interest rate implied by option prices, we use T-bill rates instead, obtained from the Federal Reserve Bank of St. Louis Economic Research Database (FRED[®]).

B.2.3 Option prices. In the final sample, only call and put options with at least 100 traded contracts are included. We calculate a hypothetical noon option cross section from the closing cross section and the index observed at noon and the close. Here we assume that the implied volatilities do not change between noon and the close. We start out with 109 raw cross sections and are left with 108 final cross sections. The time to expiration is 29 days.

B.3 S&P 500 information bulletin

We obtain the historical daily record of the S&P 500 index and its daily dividend record over the 1928–2006 period from the S&P 500 information bulletin. Before April 1982, dividends are estimated from monthly dividend yields.

B.4 VIX

We use the *revised* CBOE Volatility Index (VIX). The revised VIX is 100 times the forecast of the annualized 30-day volatility of the S&P 500 index. It is a model-independent forecast based on S&P 500 index options with 1-month and 2-month expiration and a wide range of in-the-moneyness. The index has been back-filled by the CBOE and is currently available from 1990 to the present. Note that VIX is not available for panels A and B of Tables 1, 2, and 4.

Appendix C

The GARCH (1,1) special case of the Engle and Gonzalez-Rivera (1991) semiparametric model applied to the monthly S&P 500 index return, y_t , is described by equations (C1)–(C3):

$$y_t = \bar{y} + \varepsilon_t \quad (\text{C1})$$

$$h^{-1/2}\varepsilon_t \sim i.i.d. g(0, 1) \quad (\text{C2})$$

and

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}, \quad (\text{C3})$$

where $g(0, 1)$ is an unknown distribution with zero mean and unit variance.

The parameters (ω, α, β) are estimated by maximum likelihood under the (false) assumption that $h^{-1/2}\varepsilon_t \sim i.i.d. N(0, 1)$. Then the time series $\{h_t^{-1/2}\varepsilon_t\}$ is calculated and the true density $g(0, 1)$ is estimated as the histogram of all the time-series observations. The histogram is being smoothed by our kernel methods in the same way as all the other distributions in order to keep the procedures comparable.

One may consider reestimating the parameters (ω, α, β) by maximum likelihood, replacing the assumption that $h^{-1/2}\varepsilon_t \sim i.i.d. N(0, 1)$ with the assumption that $h^{-1/2}\varepsilon_t \sim i.i.d. \hat{g}(0, 1)$, where $\hat{g}(0, 1)$ is the estimated density in the last step above. Engle and Gonzalez-Rivera (1991) showed by simulation that this additional step is unnecessary in practice.

References

- Ait-Sahalia, Y., and A. W. Lo. 1998. Nonparametric Estimation of State-Price Densities Implicit in Financial Prices. *Journal of Finance* 53:499–547.
- Ait-Sahalia, Y., and A. W. Lo. 2000. Nonparametric Risk Management and Implied Risk Aversion. *Journal of Econometrics* 94:9–51.
- Ait-Sahalia, Y., and J. Duarte. 2003. Nonparametric Option Pricing under Shape Restrictions. *Journal of Econometrics* 116:9–47.
- Bates, D. S. 2006. The Market for Crash Risk. Working Paper, University of Iowa, Iowa City.
- Bensaid, B., J.-P. Lesne, H. Pagés, and J. Scheinkman. 1992. Derivative Asset Pricing with Transaction Costs. *Mathematical Finance* 2:63–86.
- Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein. 2007. Explaining Pre- and Post-1987 Crash Prices of Equity and Options within a Unified General Equilibrium Framework. Working Paper, University of Minnesota.
- Black, F., and M. S. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81:637–54.
- Bollen, N., and R. Whaley. 2004. Does Net Buying Pressure Affect the Shape of Implied Volatility Functions? *Journal of Finance* 59:711–53.

- Brennan, M. J., X. Liu, and Y. Xia. 2006. Option Pricing Kernels and the ICAPM. Working Paper, UCLA.
- Brown, D., and J. C. Jackwerth. 2004. The Pricing Kernel Puzzle: Reconciling Index Option Data and Economic Theory. Working Paper, University of Wisconsin, Madison.
- Christoffersen, P., S. Heston, and K. Jacobs. 2006. Option Valuation with Conditional Skewness. *Journal of Econometrics* 131:253–84.
- Constantinides, G. M. 1979. Multiperiod Consumption and Investment Behavior with Convex Transactions Costs. *Management Science* 25:1127–37.
- Constantinides, G. M., and S. Perrakis. 2002. Stochastic Dominance Bounds on Derivatives Prices in a Multi-period Economy with Proportional Transaction Costs. *Journal of Economic Dynamics and Control* 26:1323–52.
- Constantinides, G. M., and S. Perrakis. 2007. Stochastic Dominance Bounds on American Option Prices in Markets with Frictions. *Review of Finance* 11:71–115.
- Constantinides, G. M., M. Czerwonko, J. C. Jackwerth, and S. Perrakis. 2007. Are Options on Index Futures Profitable for Risk Averse Investors? Empirical Evidence. Working Paper, University of Chicago.
- Constantinides, G. M., and T. Zariphopoulou. 1999. Bounds on Prices of Contingent Claims in an Intertemporal Economy with Proportional Transaction Costs and General Preferences. *Finance and Stochastics* 3:345–69.
- Constantinides, G. M., and T. Zariphopoulou. 2001. Bounds on Derivative Prices in an Intertemporal Setting with Proportional Transaction Costs and Multiple Securities. *Mathematical Finance* 11:331–46.
- David, A., and P. Veronesi. 2002. Option Prices with Uncertain Fundamentals. Working Paper, University of Chicago.
- Engle, R. F., and G. Gonzalez-Rivera. 1991. Semiparametric ARCH Models. *Journal of Business and Economic Statistics* 9:345–59.
- Fama, E. F. 1970. Multiperiod Consumption-Investment Decisions. *American Economic Review* 60:163–74.
- Garcia, R., R. Luger, and E. Renault. 2003. Empirical Assessment of an Intertemporal Option Pricing Model with Latent Variables. *Journal of Econometrics* 116:49–83.
- Han, B. 2006. Limits of Arbitrage, Sentiment, and Index Option Smile. Working Paper, University of Texas, Austin.
- Hull, J. C. 2006. *Options, Futures, and Other Derivatives*. Englewood Cliffs, NJ: Prentice-Hall.
- Jackwerth, J. C. 2000. Recovering Risk Aversion from Option Prices and Realized Returns. *Review of Financial Studies* 13:433–51.
- Jackwerth, J. C. 2004. Option-Implied Risk-Neutral Distributions and Risk Aversion. Research Foundation of AIMR, Charlottesville, VA, ISBN 0-943205-66-2.
- Jackwerth, J. C., and M. Rubinstein. 1996. Recovering Probability Distributions from Option Prices. *Journal of Finance* 51:1611–31.
- Leland, H. E. 1985. Option Pricing and Replication with Transactions Costs. *Journal of Finance* 40:1283–301.
- Levy, H. 1985. Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach. *Journal of Finance* 40:1197–217.
- Liu, J., J. Pan, and T. Wang. 2005. An Equilibrium Model of Rare-Event Premia and Its Implications for Option Smirks. *Review of Financial Studies* 18:131–64.
- McDonald, R. L.. 2006. *Derivatives Markets*. Reading, MA: Addison-Wesley.
- Merton, R. C. 1973. Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science* 4:141–83.

Mispricing of S&P 500 Index Options

- Oancea, I. M., and S. Perrakis. 2007. Stochastic Dominance and Option Pricing in Discrete and Continuous Time: An Alternative Paradigm. Working Paper, Concordia University, Montreal.
- Perrakis, S. 1986. Option Bounds in Discrete Time: Extensions and the Pricing of the American Put. *Journal of Business* 59:119–41.
- Perrakis, S. 1993. Option Pricing for Multinomial Stock Returns in Diffusion and Mixed Processes. *Canadian Journal of Administrative Sciences* 10:68–82.
- Perrakis, S., and P. J. Ryan. 1984. Option Pricing Bounds in Discrete Time. *Journal of Finance* 39:519–25.
- Ritchken, P. H. 1985. On Option Pricing Bounds. *Journal of Finance* 40:1219–33.
- Ritchken, P. H., and S. Kuo. 1988. Option Bounds with Finite Revision Opportunities. *Journal of Finance* 43:301–8.
- Rosenberg, J. V., and R. F. Engle. 2002. Empirical Pricing Kernels. *Journal of Financial Economics* 64:341–72.
- Rubinstein, M. 1994. Implied Binomial Trees. *Journal of Finance* 49:771–818.
- Ryan, P. J. 2000. Tighter Option Bounds from Multiple Exercise Prices. *Review of Derivatives Studies* 4:155–88.
- Ryan, P. J. 2003. Progressive Option Bounds from the Sequence of Concurrently Expiring Options. *European Journal of Operational Research* 151:193–223.
- Santa-Clara, P., and S. Yan. 2004. Jump and Volatility Risk and Risk Premia: A New Model and Lessons from S&P 500 Options. Working Paper, UCLA.
- Shefrin, H. 2005. *A Behavioral Approach to Asset Pricing*. Amsterdam: Elsevier/North-Holland
- Silverman, B. W. 1986. *Density Estimation*. London: Chapman and Hall.
- Singleton, K. J. 2006. *Empirical Dynamic Asset Pricing*, Princeton, NJ: Princeton University Press.