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

Abstract

We document violations of stochastic dominance in the one-month S&P 500 index options market in the period 1986-1995 when the unconditional index return distribution is taken to be that of the historical index sample and two forward-looking samples that include or exclude the crash. Stochastic dominance means that a trader can improve her expected utility by engaging in a zero-net-cost trade. The violations persist when we allow for realistic transactions costs. Even though pre-crash option prices follow the Black-Scholes-Merton (BSM) model reasonably well, they are incorrectly priced, if index return expectations are based on the historical experience. Furthermore, some of these prices are below the bounds, contrary to received wisdom that historical volatility generally underprices options in the BSM model, dispelling the common misconception that the observed smile is too steep after the crash.

Current draft: October 21, 2004

JEL classification: G13

Keywords: Derivative pricing; volatility smile, incomplete markets, transaction costs; index options; stochastic dominance bounds

We thank workshop participants at the Bachelier 2004 Congress, the University of Chicago, and Concordia, Laval, NYU and St. Gallen Universities and, in particular, Menahem Brenner and Matthew Richardson for their insightful comments and constructive criticism. 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. E-mail addresses: gmc@ChicagoGSB.edu, Jens.Jackwerth@uni-konstanz.de, and SPerrakis@mercato.concordia.ca.

1 Introduction and Summary

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 striking prices. Rubinstein (1994) tested this prediction on the S&P 500 index options traded on the Chicago Board Options Exchange, an exchange that comes close to the dynamically complete and perfect market assumptions underlying the BSM model. As a function of the strike price, the implied volatility is flat from the start of the exchange-based trading in April 1986 until the October 1987 stock market crash. Thereafter, it is downward-sloping, a pattern referred to as the "volatility smile" that is also observed in international markets and to a lesser extent on individual-stock options.[1]

Ait-Sahalia and Lo (1998) Jackwerth and Rubinstein (1996), and Jackwerth (2000), among others, refined the result by estimating the *risk-neutral* stock price distribution from the cross section of option prices. Jackwerth and Rubinstein (1996) confirmed that, prior to the October 1987 crash, the risk-neutral stock price distribution is close to lognormal, consistent with flat implied volatility. Thereafter, the distribution is systematically skewed to the left, consistent with downward-sloping implied volatility. Jackwerth (2000) found that the pricing kernel implied by the observed cross section of prices of S&P 500 index options is everywhere decreasing prior to the October 1987 crash but that widespread violations are observed thereafter.

These findings raise several important questions. Does the BSM model work well prior to the crash? If it does, is it because the risk-neutral probability of a stock market crash was low and consistent with a lognormal distribution? Or, is it because the risk-neutral probability of a stock market crash was erroneously perceived to be low by the market participants? Why does the BSM model fail after the crash? Is it because the risk neutral probability of a stock market crash

---

[1] Jackwerth (2004), Brown and Jackwerth (2004), and Whaley (2003) review the literature and potential explanations. Jackwerth (2004) also reviews the parametric and non-parametric methods for estimating the risk-neutral distribution.

increased after the crash and became inconsistent with a lognormal distribution? Or, is it because the risk neutral probability of a stock market crash was erroneously perceived to do so? These are some of the questions that we address in this paper.

Whereas a downward sloping implied volatility is inconsistent with the BSM model, it is important to realize that this pattern is not inconsistent with economic theory in general. Two fundamental assumptions of the BSM model are that the market is frictionless and dynamically complete. We empirically investigate whether the observed cross section of S&P 500 index option prices are consistent with various economic models that explicitly allow for a dynamically incomplete market and also recognize trading costs and bid-ask spreads.

Both dynamic incompleteness and frictions have drastic consequences for option pricing. In both cases arbitrage methods alone are incapable of producing a unique option price. Market incompleteness is generally handled either by employing an equilibrium model that prices risk, or by testing bounds within which the option price should lie, along the lines explored in this paper.[2] Bounds containing the option price also arise when there are trading costs and bid-ask spreads.[3]

We avoid any a priori assumptions about the form of the real unconditional distribution of the S&P 500 index returns and use histograms extracted from three different data samples as estimates of this distribution: the historical index sample (1928-1985) and two forward-looking samples, one that includes the October 1987 crash (1986-1995) and one that excludes it (1988-1995). Based on the index return distributions extracted from these samples, we test the compliance of option prices to the predictions of models that sequentially introduce market incompleteness, transactions costs, and intermediate trading over the life of the options.

The paper is organized as follows. In Section 2, we present a general model for pricing options that incorporates market incompleteness, transactions costs, and intermediate trading over the life of the options. We specialize this model into a series of models that sequentially introduce market incompleteness, transactions

---

[2] For models that price risk, see, for instance, Bailey and Stulz (1989) and Amin (1993).

[3] See Leland (1985) and Bensaid *et al* (1992).

costs, and intermediate trading over the life of the options. Each model imposes restrictions on the cross section of the prices of options. In Section 3, we test the compliance of bid and ask one-month index options to the theoretical restrictions and discuss the results. The results are summarized in Section 4.

2 Theoretical Restrictions on Option Prices

2.1 A Model for Trading Equity and Bonds

We consider a market with heterogeneous agents and investigate the restrictions on option prices imposed by a particular class of utility-maximizing traders that we simply refer to 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 unaffected by the presence of agents with beliefs, endowments, preferences, trading restrictions, and transaction cost schedules that differ from those of the utility-maximizing traders.

Trading occurs at a finite number of trading dates, $t = 0,1,...,T,...,T'$.[4] The utility-maximizing traders are allowed to hold only two primary securities in the market, a bond and a stock. The stock has the natural interpretation as the market index. Derivatives are introduced in the next section. The bond is risk free and pays constant interest $R-1$ each period. The traders may buy and sell the bond without incurring transactions costs. At 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, $(1+\delta_t)S_{t+1}/S_t$, is identically and independently distributed over time. The assumption of i.i.d. returns is not innocuous and, in particular, rules out stochastic

[4] The calendar length of the trading horizon is N years and the calendar length between trading dates is N/T' years. Later on we vary T' and consider the mispricing of options under different assumptions regarding the calendar length between trading dates.

volatility. We deliberately rule out stochastic volatility in order to explore the extent to which market incompleteness and market imperfections (trading costs) alone explain the prices of index options. [5]

Stock trades incur proportional transaction costs charged to the bond account as follows. At 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 transactions cost rate satisfies the restriction $0 \le k < 1$. Note that there is no presumption that all agents in the economy face the same schedule of transactions costs as the traders do.

A trader enters the market at date t with dollar holdings x_t in the bond account and y_t / S_t *ex dividend* shares of stock. The endowments are stated net of any dividend payable on the stock at time t.[6] The trader increases (or, decreases) the dollar holdings in the stock account from y_t to $y_t' = y_t + \upsilon_t$ by decreasing (or, increasing) the bond account from x_t to $x_t' = x_t - \upsilon_t - k\,|\,\upsilon_t\,|$. The decision variable υ_t is constrained to be measurable with respect to the information at date t. The bond account dynamics is

$$x_{t+1} = \{x_t - \upsilon_t - k\,|\,\upsilon_t\,|\}R + (y_t + \upsilon_t)\frac{\delta_t S_{t+1}}{S_t}, \quad t \le T'-1 \tag{2.1}$$

and the stock account dynamics is

$$y_{t+1} = (y_t + \upsilon_t)\frac{S_{t+1}}{S_t}, \quad t \le T'-1. \tag{2.2}$$

---

[5] The results in Sections 2.3 and 2.4 hold without the i.i.d. returns assumption. The assumption may also be relaxed in Sections 2.5 and 2.6 for special classes of stock return distributions.

[6] We elaborate on the precise sequence of events. The trader enters the market at date t with dollar holdings $x_t - \delta_t y_t$ in the bond account and y_t / S_t *cum dividend* shares of stock. Then the stock pays cash dividend $\delta_t y_t$ and the dollar holdings in the bond account become x_t. Thus, the trader has dollar holdings x_t in the bond account and y_t / S_t *ex dividend* shares of stock.

At the terminal date, the stock account is liquidated, $\upsilon_{T'} = -y_{T'}$, and the net worth is $x_{T'} + y_{T'} - k \mid y_{T'} \mid$. At each date t, the trader chooses investment υ_t to maximize the expected utility of net worth, $E\left[u\left(x_{T'} + y_{T'} - k \mid y_{T'} \mid\right) \mid S_t\right]$.[7] We make the plausible assumption that the utility function, $u(\cdot)$, is increasing and concave, and is defined for both positive and negative terminal net worth.[8] Note that even this weak assumption of monotonicity and concavity of preferences is not imposed on all agents in the economy but only on the subset of agents that we refer to as traders.

We define the value function recursively as

$$V\left(x_t, y_t, t\right) = \max_\upsilon E\left[V\left(\left\{x_t - \upsilon - k \mid \upsilon \mid\right\} R + \left(y_t + \upsilon\right)\frac{\delta_t S_{t+1}}{S_t}, \left(y_t + \upsilon\right)\frac{S_{t+1}}{S_t}, t+1\right) \mid S_t\right] \tag{2.3}$$

for $t \le T' - 1$ and

$$V\left(x_{T'}, y_{T'}, T'\right) = u\left(x_{T'} + y_{T'} - k \mid y_{T'} \mid\right). \tag{2.4}$$

We assume that the parameters satisfy appropriate technical conditions such that the value function exists and is once differentiable.

Equations (2.1)-(2.4) define a dynamic program that can be numerically solved for given utility function and stock return distribution. We shall not solve this dynamic program because our goal is to derive restrictions on the prices of

[7] The results extend routinely to the case that consumption occurs at each trading date and utility is defined over consumption at each of the trading dates and over the net worth at the terminal date. See Constantinides (1979) for details. The model with utility defined over terminal net worth alone is a more realistic representation of the objective function of financial institutions.

[8] If utility is defined only for non-negative net worth, then the decision variable is constrained to be a member of a convex set that ensures the non-negativity of net worth. See, Constantinides (1979) for details. However, the derivation of bounds on the prices of derivatives requires an entirely different approach and yields weaker bounds. This problem is studied in Constantinides and Zariphopoulou (1999, 2001).

options that are independent of the specific functional form of the utility function but solely depend on the plausible assumption that the traders' utility function is monotone increasing and concave in the terminal wealth.

Without explicitly solving the above dynamic program, we identify certain necessary properties of the value function. The value function, $V(t) \equiv V(x_t, y_t, t)$, is increasing and concave in (x_t, y_t), properties that it inherits from the assumed monotonicity and concavity of the utility function, as proved in Constantinides (1979). We state the monotonicity of the value function as

$$V_x(t) > 0, \quad V_y(t) > 0, \qquad t = 0,\ldots,T,\ldots,T'. \tag{2.5}$$

and its concavity with respect to (x_t, y_t) as[9]

$$V_{xx}(t) < 0, \quad V_{yy}(t) < 0, \quad V_{xx}(t)V_{yy}(t) - \left(V_{xy}(t)\right)^2 > 0, \qquad t = 0,\ldots,T,\ldots,T'. \tag{2.6}$$

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

$$(1-k)V_x(t) \le V_y(t) \le (1+k)V_x(t), \quad t = 0,\ldots,T,\ldots,T'. \tag{2.7}$$

Marginal analysis on the bond holdings leads to the following condition on the marginal rate of substitution between the bond holdings at dates t and $t+1$:

$$V_x(t) = R\,E_t[V_x(t+1)], \quad t = 0,\ldots,T,\ldots,T'-1. \tag{2.8}$$

---

[9] The second-order partial derivatives of the value function exist everywhere except on the boundaries of the region of no transactions.

Finally, marginal analysis on the stock holdings leads to the following condition on the marginal rate of substitution between the stock holdings at date t and the bond and stock holdings at date $t+1$:

$$V_y(t) = E_t\left[\frac{S_{t+1}}{S_t}V_y(t+1) + \frac{\delta_t S_{t+1}}{S_t}V_x(t+1)\right], \quad t = 0,...,T,...,T'-1. \quad (2.9)$$

Below we employ these conditions on the value function to derive restrictions on the prices of options.

2.2 Restrictions on the Prices of Options

We consider $J, \quad j = 1,2,...,J$, European call and put options on the index, with random cash payoff X_j at their common expiration date T, $T \leq T'$. At time zero, the trader can buy the j^{th} derivative at price $P_j + k_j$ and sell it at price $P_j - k_j$, net of transactions costs. Thus $2k_j$ is the bid-ask spread plus the round-trip transactions cost that the trader incurs in trading the j^{th} derivative. In our empirical investigation we consider both the case where k_j is common across derivatives and the case where k_j is proportional to the price of the j^{th} derivative. Note that there is no presumption that all agents in the economy face the same bid-ask spreads and transactions costs as the traders do.

We assume that the traders are marginal in all the J derivatives. Furthermore, we assume that, if a trader holds a finite (positive or negative) number of the derivatives, these positions are sufficiently small relative to her holdings in the bond and stock that the monotonicity and concavity conditions (2.5) and (2.6) on the value function remain valid. (Conditions (2.7)-(2.9) remain valid even if the holdings of the derivatives are not small.)

Marginal analysis leads to the following restrictions on the prices of options:

$$(P_j - k_j)V_x(0) \leq E_0[X_j V_x(T)] \leq (P_j + k_j)V_x(0), \quad j = 1,2,...,J. \quad (2.10)$$

Similar restrictions on the prices of options apply at dates $t = 1,...,T-1$.

Below, we illustrate the implementation of the restrictions on the prices of options in a number of important special cases. First, we consider the case $T = 1$ which rules out trading between the bond and stock accounts over the lifetime of the options. We refer to this case as the *single-period case*. Note that the single-period case does not rule out trading over the trader's horizon after the options expire; it just rules out trading over the lifetime of the options. We discuss the single-period case both with and without transactions costs. Then we consider the more realistic case $T = 2$ which allows for one intermediate trading between the bond and stock accounts over the lifetime of the options.

A useful way to identify the options that cause infeasibility or near-infeasibility of the problem is to single out a "test" option, say the n^{th} option, and solve the following problem:

$$\max_{\{V_x(t),V_y(t)\}_{t=0,...,T}} \left(\min_{\{V_x(t),V_y(t)\}_{t=0,...,T}} \right) E_0\left[X_n \frac{V_x(T)}{V_x(0)} \right] \tag{2.11}$$

subject to the conditions (2.5)-(2.10), where the n^{th} option is removed from the set of the J options in conditions (2.10). If this problem is feasible, then the attained maximum and minimum have the following interpretation. If one can buy the test option for less than the minimum attained in this problem, then there is stochastic dominance: at least one investor, *but not necessarily all investors*, increases her expected utility by trading the test option. Likewise, if one can write the test option, for more than the maximum attained in this problem, then again there is stochastic dominance.

2.3 Special Case: Zero Transactions Costs and no Trading over the Life of the Options

The stock market index has price S_0 at the beginning of the period; *ex dividend* price S_i with probability π_i in state i, $i = 1,...,I$ at the end of the period; and *cum* *dividend* price $(1+\delta)S_i$ at the end of the period. We order the states such that S_i is increasing in i. The j^{th} derivative, $j=1,...,J$, has price P_j at the beginning period and cash payoff X_{ij} at the end of the period in state i.

Since the transactions cost rate is assumed to be zero, we have $V_x(0)=V_y(0)$ and $V_x(1)=V_y(1)$. We denote by m_i the marginal rate of substitution in state i, $m_i \equiv V_x(1)/V_x(0)$. The conditions (2.8)-(2.10) become:

$$1 = R\sum_{i=1}^{I}\pi_i m_i \tag{2.12}$$

$$S_0 = \sum_{i=1}^{I}\pi_i m_i (1+\delta) S_i \, . \tag{2.13}$$

and

$$P_j = \sum_{i=1}^{I}\pi_i m_i X_{ij}, \quad j=1,...,J \, . \tag{2.14}$$

The random variable $m: m_i,\ i=1,...,I$ is the *stochastic discount factor* or *pricing kernel*. Absence of arbitrage implies and is implied by the existence of a strictly positive pricing kernel that satisfies (2.12)-(2.14). Non-existence of a strictly positive pricing kernel signifies arbitrage such as violations of the Merton (1973) no-arbitrage restrictions on the prices of options.

The concavity of the value function implies additional restrictions on the pricing kernel. Historically, the expected premium of the return on the stock over the bond is positive. Under the assumption of positive expected premium, the trader is long in the stock. Since the assumption in the single-period model is that there is no trading between the bond and stock accounts over the life of the option, the trader's wealth at the end of the period is increasing in the stock return. Note that this conclusion critically depends on the assumption that there is no intermediate trading in the bond and stock. Since we employed the convention that the stock return is increasing in the state *i*, the trader's wealth on date *T* is increasing in the state *i*. Then the concavity of the value function implies that the marginal rate of substitution is decreasing in the state *i*:

$$m_1 \geq m_2 \geq ... \geq m_I > 0 \, . \tag{2.15}$$

Existence of a pricing kernel that satisfies the restrictions (2.12)-(2.15) is said to rule out *stochastic dominance*. Ruling out stochastic dominance means that there exists at least one trader with increasing and concave utility that supports the observed prices. If we cannot rule out stochastic dominance, then *any* trader with increasing and concave utility can improve her expected utility through trading.

We emphasize that the restriction on option prices imposed by the criterion of the absence of stochastic dominance is motivated by the economically plausible assumption that there exists at least *one* agent in the economy with the properties that we assigned to a trader. This is a substantially weaker assumption than requiring that *all* agents have the properties that we assigned to traders. Stochastic dominance then implies that at least one agent, *but not necessarily all agents*, increases her expected utility by trading.[10]

In our empirical investigation, we find that in none of the months a pricing kernel exists that satisfies the restrictions (2.12)-(2.15), if we rule out bid-ask spreads and transactions costs. Therefore, in all months in our sample there is stochastic dominance.

A useful way to identify the options that cause infeasibility or near-infeasibility of the problem is to single out a "test" option, say the n^{th} option and derive bounds that signify infeasibility if the price of the test option lies outside the bounds. The general form of this problem was stated in equation (2.11). In the special case of no trading over the life of the option and zero transactions costs, the bounds on the test option with payoff X_{in} in state i are given by

$$\max_{\{m_i\}} \left(\min_{\{m_i\}} \right) \sum_{i=1}^{I} \pi_i m_i X_{in} \tag{2.16}$$

[10] We also emphasize that the restriction of the absence of stochastic dominance is weaker than the restriction that the capital asset pricing model (CAPM) holds. The CAPM requires that the pricing kernel be linearly decreasing in the index price. The absence of stochastic dominance merely imposes that the pricing kernel be monotone decreasing in the index price.

subject to the conditions (2.12)-(2.15), where the n^{th} option is removed from the set of the J options in conditions (2.14).[11]

2.4 Special Case: Transactions Costs and no Trading over the Life of the Options

We denote by $V_x(i)$, and $V_y(i)$, the partial derivatives of the value function at date one and state i with respect to the bond and stock respectively. Conditions (2.8)-(2.10) become

$$V_x(0) = R\sum_{i=1}^{I}\pi_i V_x(i). \tag{2.17}$$

$$V_y(0) = \sum_{i=1}^{I}\pi_i\left[\frac{S_i}{S_0}V_y(i) + \frac{\delta S_i}{S_0}V_x(i)\right] \tag{2.18}$$

and

$$\left(P_j - k_j\right)V_x(0) \le \sum_{i=1}^{I}\pi_i X_{ij} V_x(i) \le \left(P_j + k_j\right)V_x(0), \quad j = 1,\dots,J. \tag{2.19}$$

Conditions (2.5)-(2.7) become[12]

$$V_x(0) > 0,\ V_y(0) > 0,\ V_x(i) > 0,\ V_y(i) > 0, \quad i = 1,\dots,I \tag{2.20}$$

$$V_y(1) > V_y(2) > \dots > V_y(I) > 0 \tag{2.21}$$

and

$$(1-k)V_x(i) \le V_y(i) \le (1+k)V_x(i), \quad i = 1,\dots,I. \tag{2.22}$$

---

[11] Perrakis and Ryan (1984), Levy (1985), Ritchken (1985), and Ryan (2000, 2003) were the first to derive upper and lower bounds on the prices of European options in this case and under the assumption of zero transaction costs. See also Perrakis (1986) and Ritchken and Kuo (1988).

[12] Since the value of the bond account at the end of the period is independent of the state i, the concavity conditions $V_{xx}(t) < 0$ and $V_{xx}(1)V_{yy}(1) - \left(V_{xy}(1)\right)^2 > 0$ cannot be imposed. Only the concavity condition $V_{yy}(t) < 0$ is imposed as in equation (2.21).

In our empirical investigation, we report the percentage of months for which the problem defined by equations (2.17)-(2.22) is feasible. These are months for which stochastic dominance is ruled out. As before, a useful way of pinpointing the options that cause infeasibility or near-infeasibility of the problem is to single out a "test" option and solve the problem (2.11) subject to the restrictions (2.17)-(2.22), where the test option is removed from the set of the J options in conditions (2.19).

2.5 Intermediate Trading over the Life of the Options

In the next major step towards realism, we allow for intermediate trading, relaxing the implausible assumption of the single-period model that, over the life of the options, markets for trading are open only at the initial trading date and at the common expiration date of the options.

Recall that the single-period stochastic-dominance model without transactions implies that the wealth at the end of the period is an increasing function of the stock price at the end of the period and, therefore, the pricing kernel is a decreasing function of the stock price at the end of the period. Likewise, with transactions costs, the value of the stock account at the end of the period is an increasing function of the stock price at the end of the period and, therefore, the marginal utility of wealth out of the stock account is a decreasing function of the stock price at the end of the period.

Constantinides and Zariphopoulou (1999) pointed out that intermediate trading invalidates these implications because the wealth at the end of the period (or, the value of the stock account at the end of the period) becomes a function not only of the stock price at the option's expiration but also of the entire sample path of the stock price.[13]

[13] In the special case of i.i.d. returns, power utility and zero transactions costs, the wealth at the end of the period is a function only of the stock price. However, these assumption would considerably diminish the generality of the present paper.

In the empirical section, we test for violations of stochastic dominance by testing for the feasibility of the conditions (2.5)-(2.10) in the case $T = 2$. For T large, the numerical implementation becomes tedious and potentially explosive. This consideration motivates the development of bounds that are independent of the allowed frequency of trading of the stock and bond over the life of the option. These bounds are presented below.

2.6 The Constantinides-Perrakis Option Bounds

Constantinides and Perrakis (2002) recognized that it is possible to recursively apply the single-period approach and derive stochastic dominance bounds on option prices in a market with intermediate trading over the life of the options. The task of computing these bounds is easy compared to the full-fledged investigation of the feasibility of conditions (2.5)-(2.10) for large T for two reasons. First, the derivation of the bounds takes advantage of the special structure of the payoff of a call or put option, specifically the convexity of the payoff as a function of the stock price. Second, the set of assets is limited to three assets: the bond, stock and one option, the test option.

The upper and lower bounds on a test option have the following interpretation. If one can buy the test option for less than the lower option bound, then there is stochastic dominance between the bond, stock and the test option. Likewise, if one can write the test option for more than the upper option bound, then again there is stochastic dominance between the bond, stock and the test option. Below, we state these bounds without proof.[14]

At any time t prior to expiration, the following is an upper bound on the price of a call:

[14] These bounds may not be the tightest possible bounds for any given frequency of trading. However, they are presented here because of their universality in 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). See also Constantinides and Perrakis (2004) for extensions to American-style options and futures options.

$$\overline{c}(S_t,\ \mathrm{t}) = \frac{(1+k)}{(1-k)R_S^{T-t}} E\left[\left[S_T - K\right]^+ \left|S_t\right.\right], \tag{2.23}$$

where R_S is the expected return on the stock.

A partition-independent 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, $T = T'$. In such a case, transactions costs become irrelevant in the put-call parity and the following is a lower bound:[15]

$$\underline{c}(S_t,\ \mathrm{t}) = (1+\delta)^{\mathrm{t-T}}\, \mathrm{S_t} - K / R^{T-t} + E[(K - S_T)^+ \left|\mathrm{S_t}\right.]/\mathrm{R_S^{T-t}} \tag{2.24}$$

where R is one plus the risk free interest rate.

Put option upper and lower bounds also exist that are independent of the frequency of trading. They are given as follows:

$$\overline{p}(S_t,\ t) = \frac{K}{R^{T-t}} + \frac{1-k}{1+k}\left(R_S^{T-t}\right)^{-1}\left[E\left[\left[K - S_T\right]^+\right] - K \middle| S_t\right], \tag{2.25}$$

and

$$\begin{aligned} \underline{p}(S_t, t) &= \left(R_S^{T-t}\right)^{-1} \frac{1 - k}{1+k} E\left[\left[K - S_T\right]^+ \mid S_t\right], \quad t \le T-1 \\ = \underline{p}(S_T, T) &= \left[K - S_T\right]^+, \quad t = T. \end{aligned} \tag{2.26}$$

In the empirical section, we test for violations of stochastic dominance by testing for violations of these bounds, first without transactions costs and second with transactions costs.

[15] In the special case of zero transactions costs, the assumption $T = T'$ is redundant because the put-call parity holds.

3 Empirical Results

3.1 The Data

We use the historical daily record of the S&P 500 index and its daily dividend record over the period 1928-1995. 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. The unconditional distribution of the index is extracted from three alternative samples of thirty-day index returns: the *historical* sample uses returns over the period 1928-1986; the *forward-looking* sample *inclusive of the crash* uses the returns over the period 1987-1995 and includes the 1987 stock market crash; finally, the *forward-looking* sample *exclusive of the crash* uses the returns over the period 1988-1995 and excludes the stock market crash.

For the S&P 500 index options we use the tick-by-tick Berkeley Options Database of all quotes and trades over the years 1986-1995. We focus on the most liquid options by excluding options that are deeper than 10% out of the money or deeper than 5% in the money. For 108 months we retain only the call option quotes for the day corresponding to options thirty days to expiration.[16] For each day retained in the sample, we aggregate the quotes to the minute and pick the minute between 9:00-11:00 AM with the most quotes as our cross-section for the month. We present these quotes in terms of their bid and ask implied volatilities. These are the volatilities which would be needed in the BSM formula to price the option exactly at the bid or ask quote, respectively. Details on the databases can be found in the appendix, in Jackwerth and Rubinstein (1996), and in Jackwerth (2000).

[16] We lose some month for which we do not have sufficient data, namely after the crash of October 1987 and before the introduction of S&P 500 index options in April 1986.

3.2 Assumptions on Bid-Ask Spreads and Trading Fees

We introduce bid-ask spreads and trading fees as follows. For the index, we model the combined one-half bid-ask spread and one-way trading fees as a one-way proportional transactions cost rate equal to 0.5% of the index price.

For the options, we present results under a variety of assumptions regarding bid-ask spreads and trading fees. In our standard case, we set the combined one-half bid-ask spread and one-way trading fees on one option as equal to fraction 0.002 (or 0.0005 or 0.005) of the index price. This corresponds to about 75 (or 19 or 188) cents one-way fee per option, irrespective of the moneyness and, therefore, the price of the option. We also present results in which the transaction cost is proportional to the option price and calculated as follows. The combined one-half bid-ask spread and one-way trading fees on the at-the-money option is set equal to fraction 0.002 of the index price. The combined one-half bid-ask spread and one-way trading fees on some other option equals the fee on the at-the-money option multiplied by the ratio of the price of the said option and the price of the at-the-money option.

3.3 Stochastic Dominance in the Single-Period Case

We check each month for feasibility of the conditions (2.12)-(2.15). Infeasibility of these conditions implies stochastic dominance: an investor can improve her utility by trading in these assets without incurring any out-of-pocket costs. If the conditions (2.12)-(2.15) are infeasible, we check for their feasibility by using option prices one hour later and one day later. It is only after we establish that the conditions are infeasible at all three dates that we pronounce infeasibility of the conditions in a given month.

If we rule out bid-ask spreads and trading fees, we find that in none of the months the conditions (2.12)-(2.15) are feasible. Therefore, in all months in our sample there is stochastic dominance.

We introduce bid-ask spreads and trading fees as described in Section 3.2. The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transactions cost on each option is 20 bps of the index price. The number of options in each monthly cross-section fluctuates between 3 and 34 with a median of 10. The percentage of months without stochastic dominance violations are displayed in Table 1. The non-violations in the cases with one-way transaction cost on each option equal to 5 bps and 50 bps of the index price are displayed in parentheses.

The time series of option prices is divided over four subperiods and stochastic dominance violations are reported in different columns for each subperiod. The first subperiod extends from May 1986 to October 16 1987, just prior to the crash. The other three samples are all post-crash and span July 1988 to March 1991, April 1991 to August 1993, and September 1993 to December 1995.

The time series of index returns is divided into three subperiods and stochastic dominance violations are reported in different rows for each subperiod. The first subperiod covers 1928-1986. Since there are too many observations, only every 6th return is recorded in building the empirical return distribution. The second subperiod covers 1987-1995, including the crash. The third subperiod covers 1988-1995, excluding the crash. In all three subperiods, the mean premium of the index return over the risk free return is adjusted to be 4% annually.[17]

[17] We make this adjustment in order to eschew the issue of the predictability of the equity premium. Our results remain practically unchanged if we do not make this adjustment. Essentially, the prices of one-month options are insensitive to the expected return on the stock.

Table 1. Percentage of Months without Stochastic Dominance Violations in the Single-Period Case

The table displays the percentage of months in which stochastic dominance is absent in the cross-section of option prices. The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transactions cost on each option is 20 bps of the index price. In parentheses, the table displays the percentage of months in the cases with one-way transactions cost on each option equal to 5 and 50 bps of the index price.

	Panel A: 860516-871016	Panel B: 880715-910315	Panel C: 910419-930820	Panel D: 930917-951215
Number of Months in each Period	18	33	28	28
Feasibility in the Historical Index Sample 1928-1986	73% (47, 100)	90% (41, 97)	79% (18, 96)	19% (0, 96)
Feasibility in the Forward-Looking Index Sample (including the crash) 1987-1995	13% (13, 73)	55% (31, 83)	100% (57, 100)	96% (42, 100)
Feasibility in the Forward-Looking Index Sample (excluding the crash) 1988-1995	13% (0, 53)	34% (21, 72)	89% (29, 100)	100% (15, 100)

Most entries in the table are well below 100%, suggesting 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 options are inconsistent with a market where none of the securities are stochastically dominated, net of generous transactions costs. The pattern of violations differs considerably when using the historical index sample as opposed to using the forward-looking sample.

The top left entry of 73% refers to the index return distribution over the period 1928-1986 and the option prices over the pre-crash period from May 1986 to October 16, 1987. In 27% of the months the conditions (2.11)-(2.14) are infeasible and the prices imply stochastic dominance despite the generous allowance for transactions costs. In the same row to the right, the entries refer to the index return distribution over the period 1928-1986 and the option prices over the post-crash sub-periods. More violations occur in the post-crash sub-periods, consistent with the evidence in Jackwerth (2000) who finds that the estimated pricing kernel is monotonically decreasing in the pre-crash period, but locally increasing during the post-crash period.[18]

Using the forward-looking index sample 1987-1995 that includes the crash, or the forward-looking index sample 1988-1995 that excludes the crash, the pattern reverses itself and there are now more stochastic dominance violations in the pre-crash period than in the post-crash period. Comparing the top left entries and the bottom right entries that match the index period with the option period, we observe fewer cases of infeasibility with recent index and option prices than with more distant index and option prices. Thus, options are more rationally priced in the post-crash than the pre-crash period. This, despite the fact that the volatility smile is pronounced in the post-crash period and hardly present in the pre-crash period.

Since it is unclear what are appropriate transactions costs for options, in the same table the first entries in brackets present the percentage of non-violations for the case that the combined one-half bid-ask spread and one-way trading fees on one option is 5 bps of the index price. We observe a huge percentage of violations for all index and option price sub-periods. The second entries in brackets present the percentage of non-violations for the case that the combined one-half bid-ask spread and one-way trading fees on one option is 50 bps of the index price. Violations in the top left and bottom right entries of the table disappear but only because the transactions costs now are quite large.

Table 2 displays the percentage of months in which stochastic dominance is absent in the cross-section of in-the-money calls (top entry) and out-of-money calls (bottom entry). The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 0.5%. The one-way transactions

---

[18] The patterns in Jackwerth (2000) do not exactly match with Table 1 as he is using a different technique when he estimates the smoothed risk-neutral and actual distributions and divides them into each other.

cost on each option is 20 bps of the index price. For the pre-crash sample of index and option prices, there are fewer violations for ITM than OTM calls. Nonetheless, both ITM and OTM display virtually identical patterns of results, which are broadly similar to those in Table 1. For both forward-looking samples, feasibility is poor in the pre-crash and the early post-crash periods, and good to excellent in the two late post-crash periods. In the historical sample, feasibility deteriorates substantially in the late post-crash period. More significantly, feasibility improves dramatically for the more recent period in which the index and option quotes are matched: the violations disappear in the bottom right panel, while their numbers are substantial in the top left panel.

Table 2. Percentage of Months without Stochastic Dominance Violations in the Single-Period Case—ITM and OTM Calls Separately

The table displays the percentage of months in which stochastic dominance is absent in the cross-section of in-the-money calls (top entry) and out-of-money calls (bottom entry). The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transaction cost on each option is 20 bps of the index price.

	Panel A: 860516-871016	Panel B: 880715-910315	Panel C: 910419-930820	Panel D: 930917-951215
Number of Months in each Period	18	33	28	28
Feasibility in the Historical Index Sample 1928-1986	87% 80%	90% 90%	86% 100%	46% 81%
Feasibility in the Forward-Looking Index Sample (including the crash) 1987-1995	53% 27%	59% 66%	100% 100%	96% 100%
Feasibility in the Forward-Looking Index Sample (excluding the crash) 1988-1995	20% 13%	41% 48%	96% 89%	100% 100%

Table 3 displays the percentage of months in which stochastic dominance is absent in the cross-section of option prices but now with proportional instead of fixed transactions costs. The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 0.5%. The one-way transaction cost on each option is proportional to the index price, as explained in Section 3.2.

The top left entry of 73% refers to the percent feasibility in case that the index return distribution is over the period 1928-1986 and the option prices are over the pre-crash period from May 1986 to October 16, 1987. The feasibility with proportional transactions costs is identical to the feasibility with fixed transaction costs. However, with the forward-looking index sample, both including and excluding the crash, and with the options in all the post-crash sub-periods, the feasibility with proportional transactions costs is substantially lower than the feasibility with fixed transaction costs. With proportional transactions costs, options are more rationally priced in the post-crash than the pre-crash period, but in both periods the percentage of months with feasible pricing is well below 100%.

Table 3. Percentage of Months without Stochastic Dominance Violations in the Single-Period Case and Proportional Transactions Costs

The table displays the percentage of months in which stochastic dominance is absent in the cross-section of option prices. The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transactions cost on each option is proportional to the index price, as explained in Section 3.2.

	Panel A: 860516-871016	Panel B: 880715-910315	Panel C: 910419-930820	Panel D: 930917-951215
Number of Months in each Period	18	33	28	28
Feasibility in the Historical Index Sample 1928-1986	73%	90%	86%	50%
Feasibility in the Forward-Looking Index Sample (including the crash) 1987-1995	13%	48%	89%	92%
Feasibility in the Forward-Looking Index Sample (excluding the crash) 1988-1995	7%	31%	71%	92%

In the next section, we examine the implications of relaxing the assumption that trading in the stock and bond over the life of the option is forbidden.

3.4 Stochastic Dominance in the Multiperiod Case

In the previous section, we considered feasibility in the context of the single-period model. We established that there is infeasibility, and therefore stochastic dominance violations, in a significant percentage of the months. Does the percentage of stochastic dominance violations increase or decrease as the allowed frequency of trading in the stock and bond over the life of the option increases? In the very special case of zero transactions costs, i.i.d. returns and constant relative risk aversion, we can prove that the percentage of stochastic dominance violations increases 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 life of the options and another with no intermediate trading dates over the life of the options. To this end, we partition the 30-day horizon into two 15-day intervals and approximate the 15-day return distribution by an 11-point kernel density estimate of the 15-day returns. We use the standard Gaussian kernel of Silverman (1986, pp. 15, 43, and 45). The assumed transactions costs are as in the base case presented in Table 1. The one-way transactions costs rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transactions cost on each option is 20 bps of the index price. The results are presented in Table 4.

We may not investigate the effect of intermediate trading by directly comparing the results in Tables 1 and 4 because the return generating process differs in the two tables. Recall that the results in Table 1 are based on a 30-day stock return generating process that has as many different returns as the different observed realizations and frequency equal to the observed frequency.[19] By contrast, the results in Table 4 are based on a simplified 15-day 11-point kernel density estimate of the 15-day returns. The coarseness of the grid is dictated by the need to keep the problem computationally manageable. 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 brackets.

The effect of allowing for one intermediate trading date over the life of the one-month options is illustrated by the main 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. Intermediate trading generally decreases the number of feasible months in our subperiods.[20] We conclude that intermediate trading strengthens the single-period systematic evidence of stochastic dominance violations.

[19] For the historic sample of stock returns we take only one every six monthly returns.

[20] In only one case does this effect go the other way and we suspect a convergence problem with the MATLAB optimization code, pending further investigation.

Table 4. Percentage of Months without Stochastic Dominance Violations, in the 2-Period Case

The table displays the percentage of months in which stochastic dominance is absent in the cross section of option prices *when one intermediate trading date is allowed* over the life of the one-month options. The one-way transactions cost rate (one-way trading fees plus half the bid-ask spread) on the index is 50 bps. The one-way transactions cost on each option is 20 bps of the index price. In parentheses, the table displays the percentage of months in which stochastic dominance is absent in the case when *no intermediate trading is allowed* over the life of the one-month options. Two periods of 15 days and a kernel density of 15-day returns is used (discretized to 11 values from $e^{-0.20}$ to $e^{0.20}$, spaced 0.04 apart in log spacing).

	Panel A: 860516-871016	Panel B: 880715-910315	Panel C: 910419-930820	Panel D: 930917-951215
Number of Months in each Period	18	33	28	28
Feasibility in the Historical Index Sample 1928-1986	53% (80%)	45% (62%)	14% (21%)	0% (0%)
Feasibility in the Forward-Looking Index Sample (including the crash) 1987-1995	13% (13%)	17% (31%)	64% (64%)	50% (58%)
Feasibility in the Forward-Looking Index Sample (excluding the crash) 1988-1995	13% (13%)	41% (41%)	93% (86%)	96% (96%)

In the next section, we obtain further insights as to which options cause infeasibility by displaying the options that violate the upper and lower bounds on option prices.

3.5 Stochastic Dominance Bounds in the Single-Period and Multiperiod Cases

In Section 2.6, equations (2.23)-(2.26), we stated a set of stochastic dominance bounds on option prices that are valid irrespective of whether trading in the bond and stock over the life of the option is allowed or not. In this section, we calculate these bounds on option prices and translate them as bounds on the implied volatility of option prices. In the figures that follow, a violation occurs whenever an observed option bid price (translated into implied volatility) lies above the tighter of the two upper bounds or an observed option ask price (translated into implied volatility) lies below the tighter of the two lower bounds.

In Figures 1-3, we set the transactions costs rate on the stock equal to zero and also set the trading fees on the options equal to zero. The bid-ask spread on the option price is taken under consideration, as we present both the bid and ask option prices. With zero transactions costs, the two upper bounds on the implied volatility coincide because the put upper bound in equation (2.25) may be obtained from the call upper bound in equation (2.23) through the put-call parity. Likewise, with zero transactions costs, the two lower bounds on the implied volatility coincide because the put lower bound in equation (2.26) may be obtained from the call lower bound in equation (2.24) through the put-call parity.

In Figure 1, the bounds are based on the historical sample of stock returns, 1928-1986. In Figure 2, the bounds are based on the forward-looking sample of stock returns, 1987-1995, inclusive of the crash. In Figure 3, the bounds are based on the forward-looking sample of stock returns, 1988-1995, exclusive of the crash. The options data confirm what we already know from the literature: before the crash, the smile is largely flat and, after the crash, it is downward sloping.

We find that the bounds are similar across the samples (namely, they are downward sloping) but the historical stock sample induces bounds that are altogether higher as the volatility is higher in the historical sample than in the forward looking samples. The pattern of violations follows quite naturally. The flat pre-crash smile fits reasonably well within the rather high historical bounds even though these are downward sloping. The post-crash smiles are too low for the rather high location of the historical sample bounds. Going to the forward-looking sample bounds, these are located somewhat lower than the historical sample bounds. Therefore, they match the also downward-sloping post-crash option prices rather well because they are located somewhat lower too. However, they do not match very well the higher horizontal smile of the pre-crash options.

In Figure 1, we observe that both upper and lower bounds exhibit a clear smile pattern, which is present in similar forms in all four panels of Figure 1. In particular, the smile in the bounds is very much present in the pre-crash panel A, in spite of the fact that the observed option prices conform approximately to the BSM model, with a horizontal smile, which lies almost entirely within the bounds. By contrast, in post-crash panels B-D, the observed option prices show progressively more marked departures from horizontality, which still lie within the bounds in panel B but violate strongly the bounds in panels C and D, even around at-the-money. This conforms closely to the observation originally made by Rubinstein (1994), that option prices behave differently pre- and post-crash, with the former following the BSM model and the latter not.

In panel A, several pre-crash ask prices of OTM calls fall below the lower bound. Even though pre-crash option prices follow the BSM model reasonably well, it does not follow that these options are correctly priced. Our novel finding is that pre-crash option prices are incorrectly priced, if index return expectations are formed based on the historical experience. Furthermore, some of these prices are below the bounds, contrary to received wisdom that historical volatility generally underprices options in the BSM model.

Figure 1. Bounds Based on the Historical Sample without Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 without transaction costs on the option prices during those periods, based on the historical sample of stock returns, 1928-1986. Since there are no transaction costs, the two upper bounds coincide and the two lower bounds coincide.

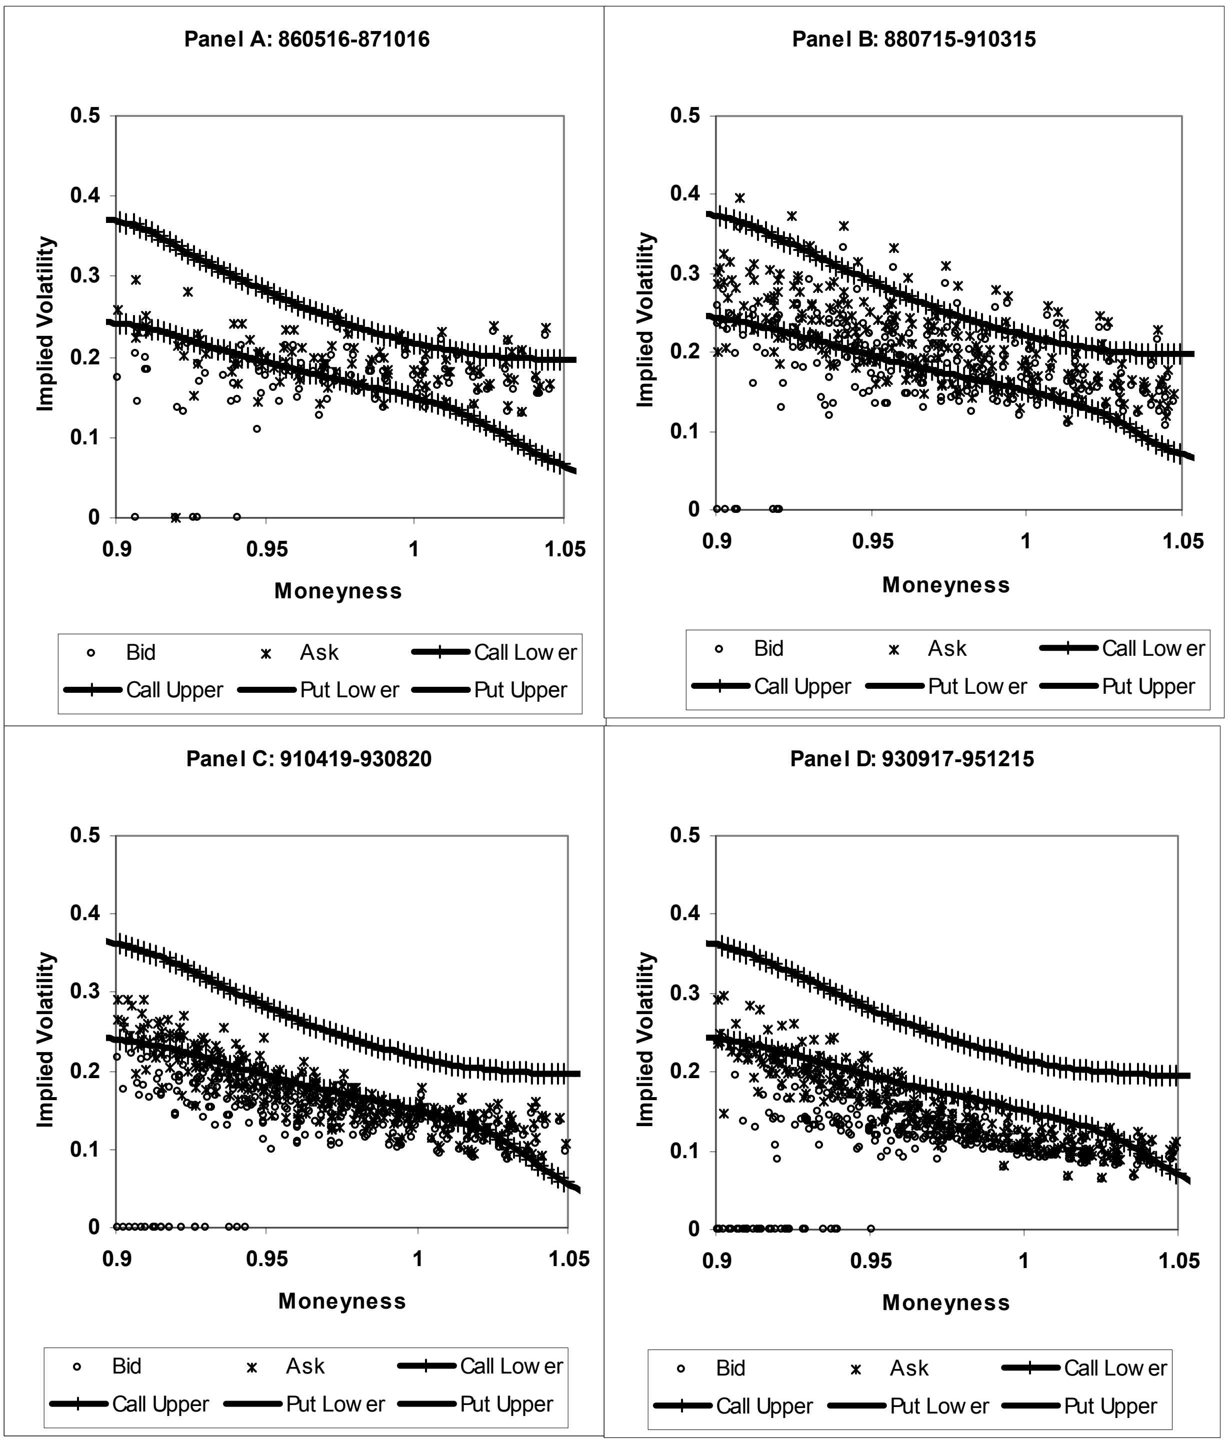

Figures 2 and 3 display the same case as Figure 1, except that the bounds are now estimated with the forward-looking sample of stock returns, one including the crash (Figure 2) and the other excluding the crash (Figure 3). We observe very different estimates of the bounds, in which the smile is much steeper in all periods and creates a very different pattern of bounds violations. Thus, it is in the pre-crash panel A that we now observe the sharpest violations of the bounds, while the violations are weaker in the first post-crash panel B, and disappear completely in the later panels C and D. The shift in the underlying index distribution, clearly visible in the change of the bounds estimates when we move from the historical to the forward-looking sample, seems to have been reflected in the option market prices as a type of "learning process" for the option investor, albeit with a certain lag. This shift is particularly pronounced in Figure 3, where the crash period has been excluded and for which the fit of the bounds with the observed option prices is better.

The following plausible scenario emerges from this pattern. In the early years since April 1986, when index options were introduced on the Chicago Board Options Exchange, investors applied the BSM model to price options without questioning the empirical validity of the constant-volatility assumption of the model. Hence prices largely adhere to the BSM model during this period. Following the crash of 1987, investors seem to have recognized the limitations of the BSM model.

Figures 2 and 3, panels B-D, dispel another common misconception, that the observed smile is too steep after the crash. Our novel finding is that most of the bound violations post-crash are due to the option smile not being steep enough relative to expectations on the index price formed post-crash. Even though the BSM model assumes that there is no smile, an investor who properly understood the post-crash distribution of index returns should have priced the options with a steeper smile than the smile reflected in the actual option prices.

Figure 2. Bounds Based on the Forward Looking Sample (1987-1995) without Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 without transaction costs on the option prices during those periods, based on the forward looking sample of stock returns, 1987-1995. Since there are no transaction costs, the two upper bounds coincide and the two lower bounds coincide.

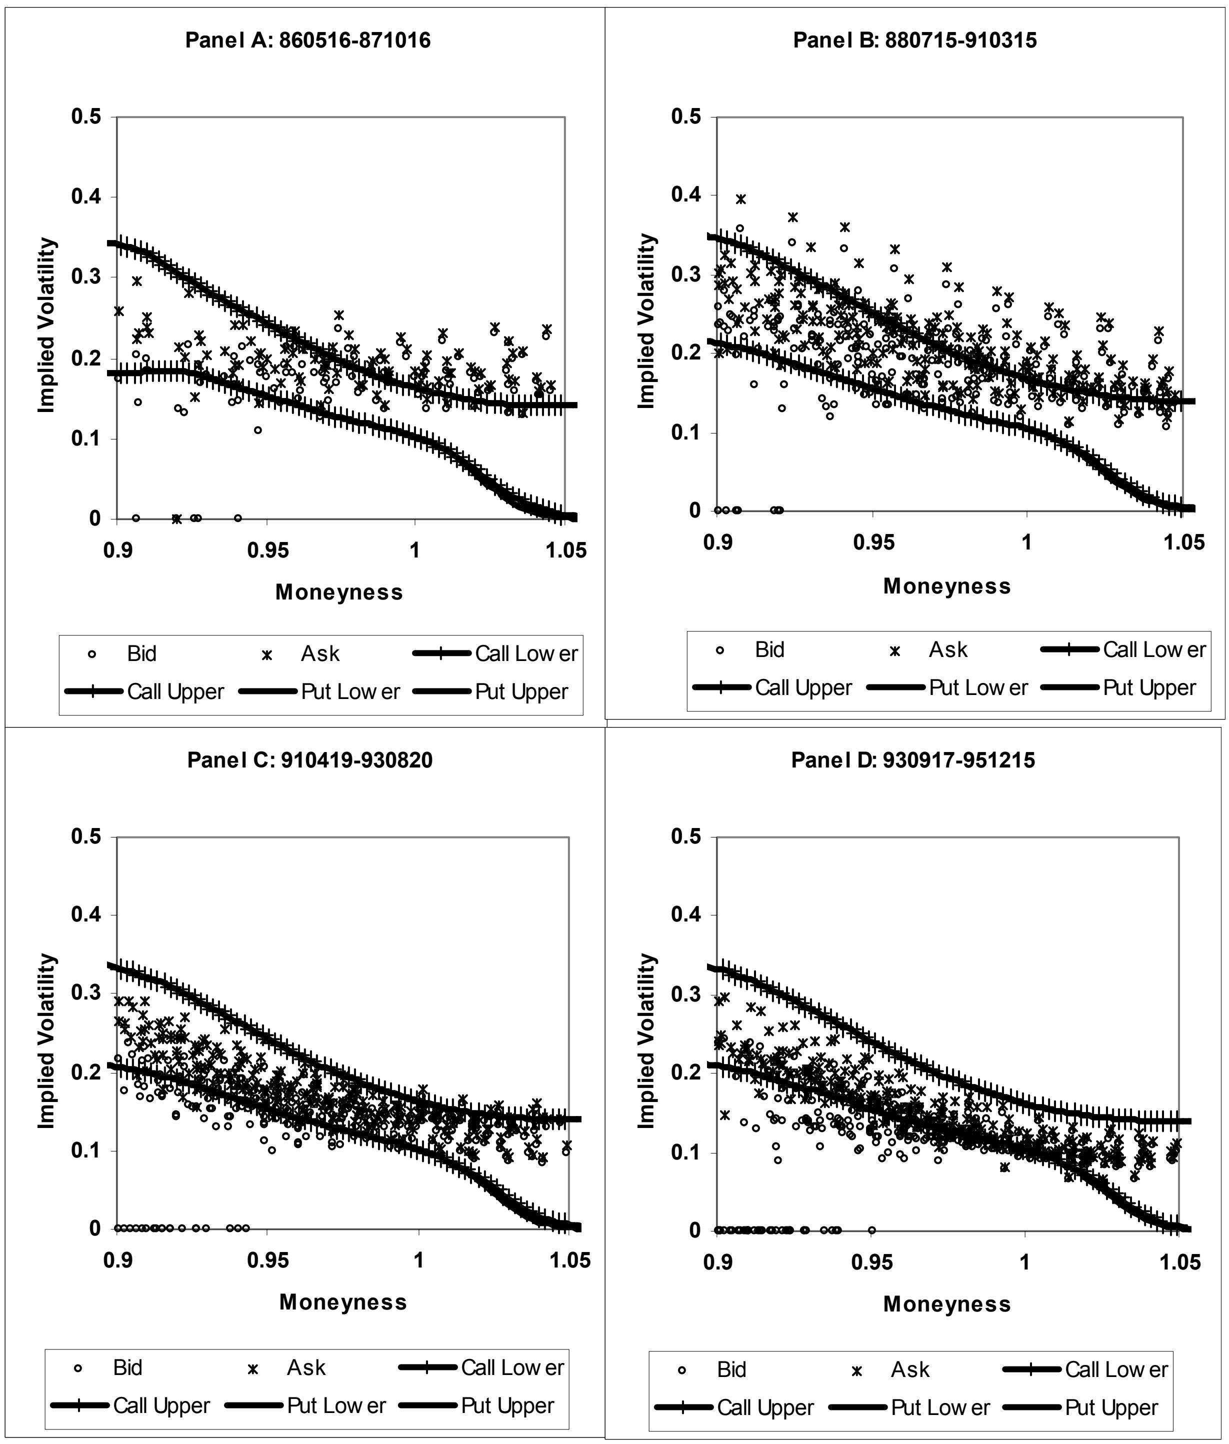

Figure 3. Bounds Based on the Forward Looking Sample (1988-1995) without Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 without transaction costs on the option prices during those periods, based on the forward looking sample of stock returns, 1988-1995. Since there are no transaction costs, the two upper bounds coincide and the two lower bounds coincide.

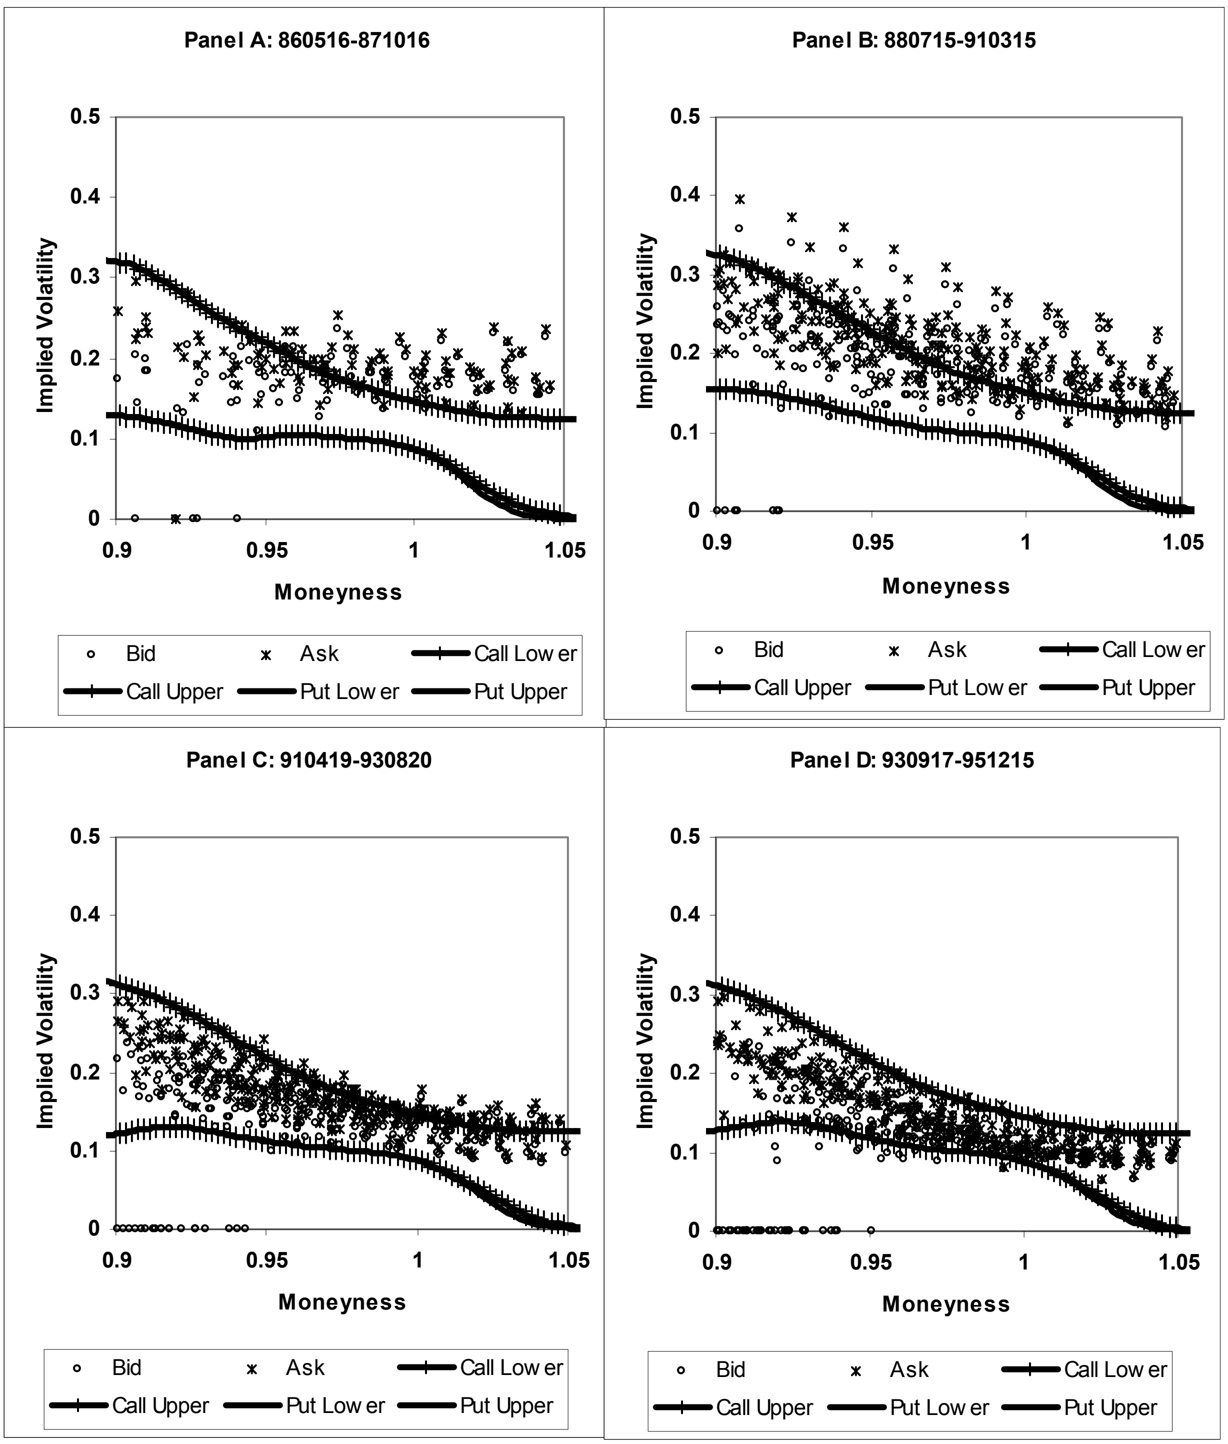

Figures 4-6 are the counterparts of Figures 1-3 with a generous allowance of 0.5% transactions cost rate in trading the stock. With transactions costs, the two upper bounds no longer coincide. We present both upper bounds but focus on the tighter of the two. The two lower bounds almost coincide. Again, we present both lower bounds but focus on the tighter of the two. Transactions costs have relatively little impact on the tighter of the two upper bounds and the tighter of the two lower bounds. The pattern of violations pre-crash and post-crash remains unchanged, although their intensity in the post-crash periods decreases due to the moderate widening of the bounds.

When the forward-looking sample is used in the estimation of the bounds, the shift in the bounds results in a worse fit of the observed option prices in the pre-crash period, but in a much better fit in the post-crash periods. In fact, the observed option prices now lie almost entirely within the bounds in the last two post-crash periods. The fit is best in Figure 6, when the year of the crash is excluded from the sample with which the bounds are estimated.

Figure 4. Bounds Based on the Historical Sample with Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 with transaction costs on the index of 20 bps during those periods, based on the historical sample of stock returns, 1928-1986.

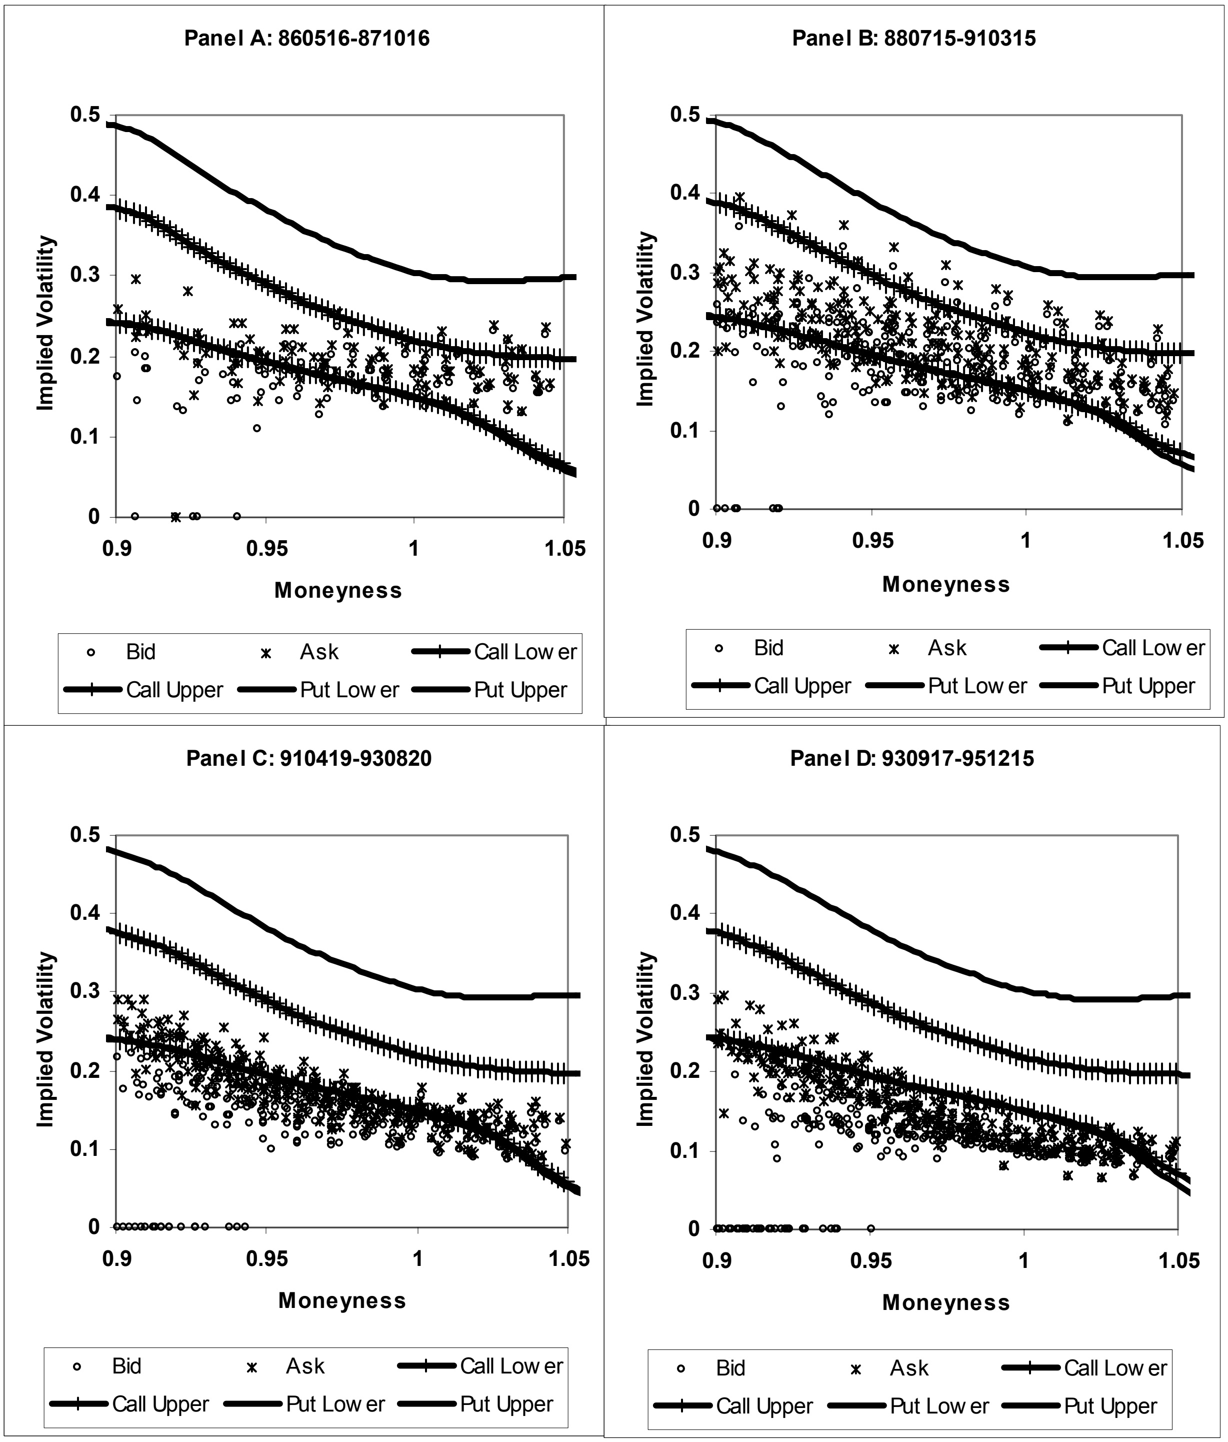

Figure 5. Bounds Based on the Forward Looking Sample (1987-1995) with Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 with transaction costs on the index of 20 bps during those periods, based on the forward looking sample of stock returns, 1987-1995.

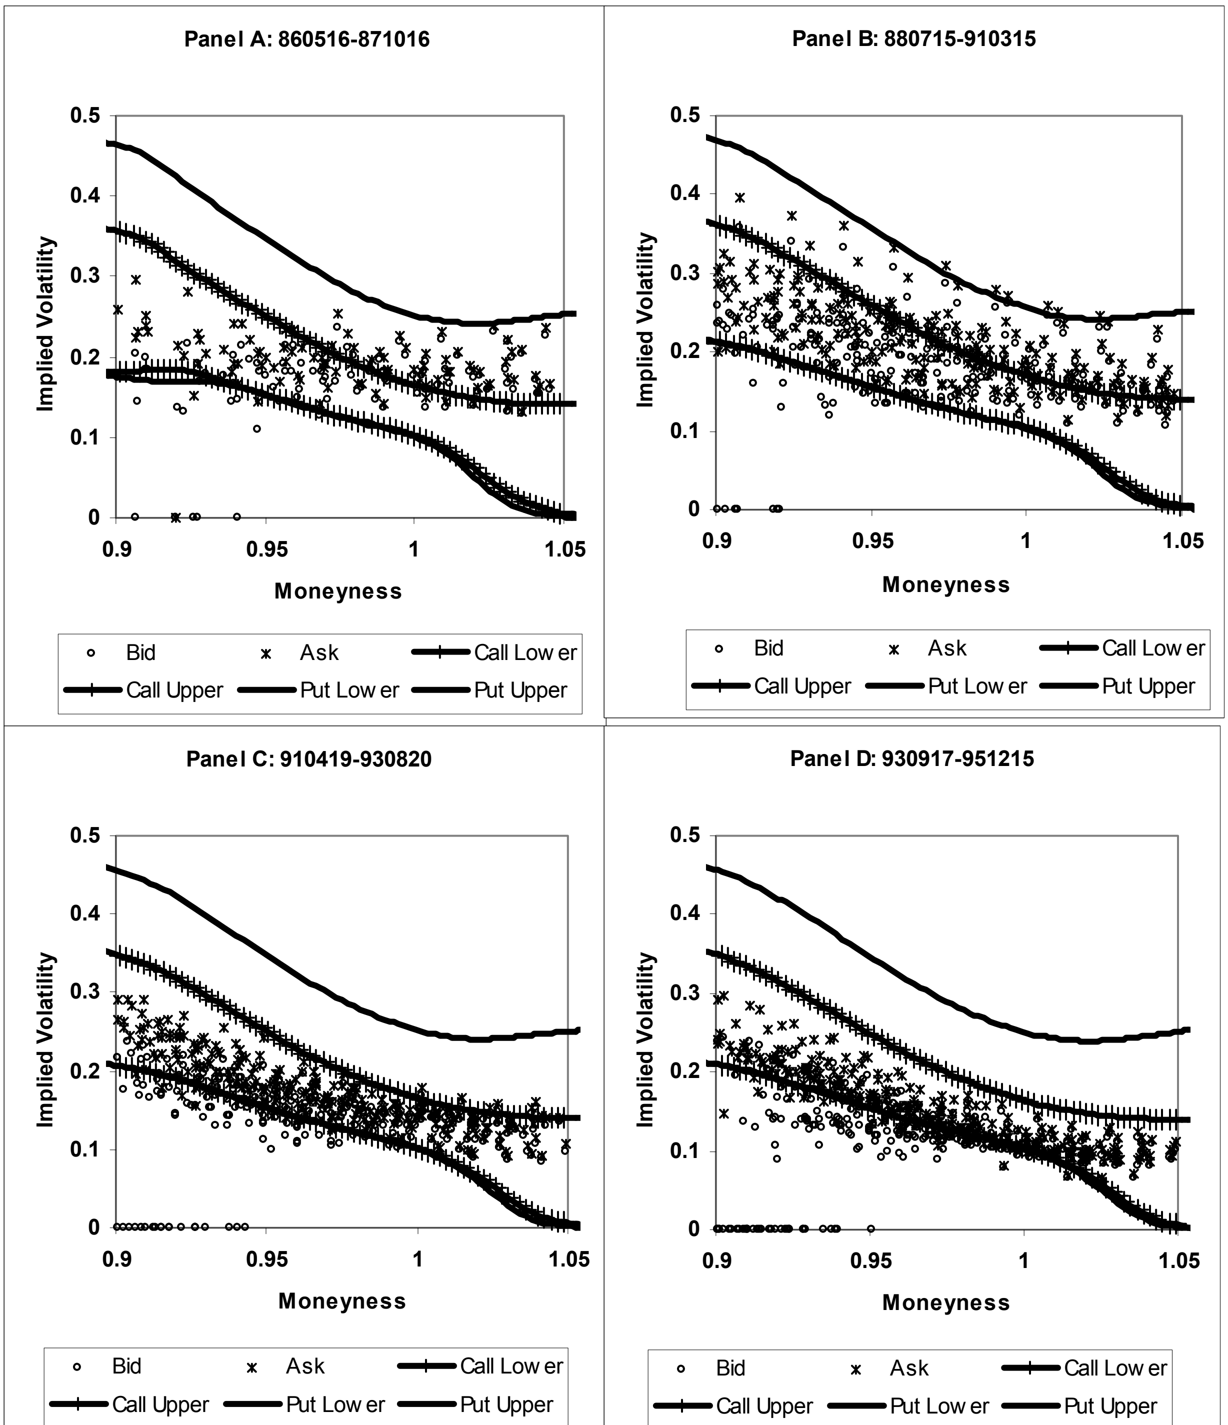

Figure 6. Bounds Based on the Forward Looking Sample (1988-1995) with Transactions Costs

The figure displays the observed bid and ask quotes during one pre-crash and three post-crash periods. It also shows the call and put upper and lower bounds from Section 2.6 with transaction costs on the index of 20 bps during those periods, based on the forward looking sample of stock returns, 1988-1995.

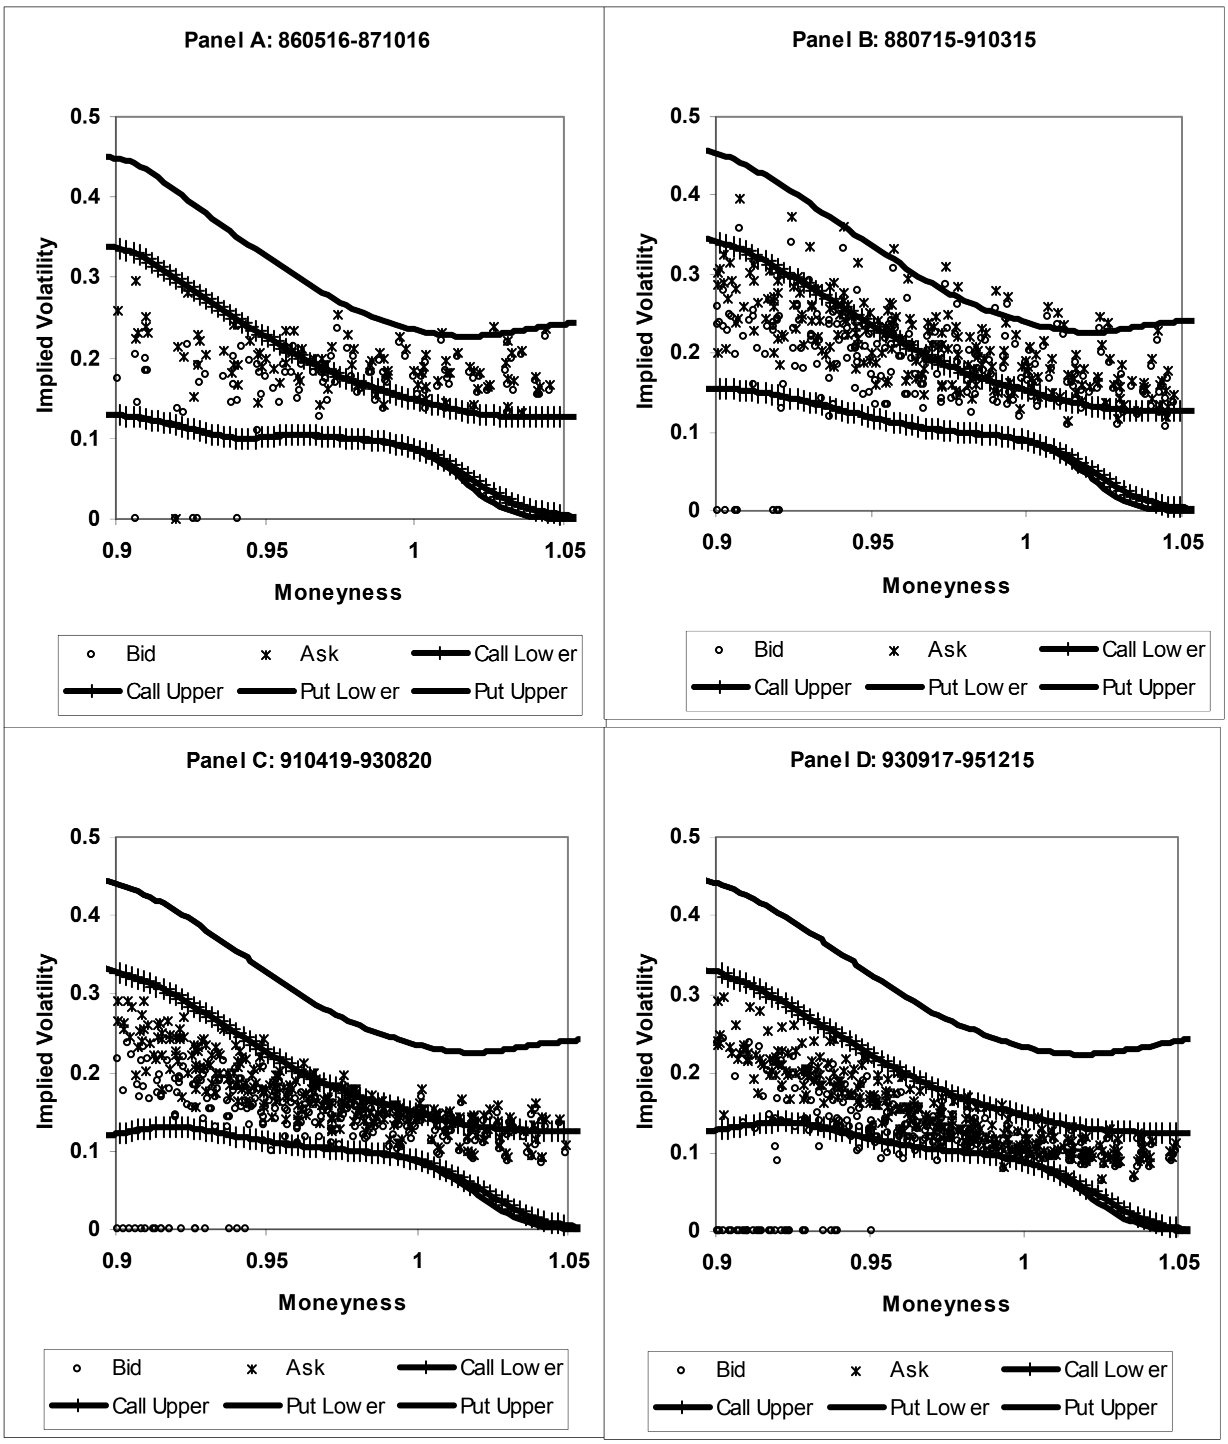

4 Concluding Remarks

We document violations of stochastic dominance in the one-month S&P 500 index options market in the period 1986-1995 when the unconditional index return distribution is taken to be that of the historical index sample (1928-1985), the forward-looking sample that includes the crash (1986-1995), or the forward-looking sample that excludes the crash (1988-1995).

Evidence of stochastic dominance means that a trader can improve her 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 traders that we simply refer to as traders. We do not make the restrictive assumption that all economic 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 transactions cost schedules that differ from those of the utility-maximizing traders modeled in this paper.

The violations increase when we allow for trading the index over the life of the options. Violations persist even when we allow for realistic trading fees in the index and options markets and recognize bid-ask spreads in the options market. The pattern of violations changes with the index sample: the historical sample generates fewer violations pre-crash and in the first sub-period after the crash and more violations in the later two subperiods of the post-crash period. The forward-looking samples generate the opposite pattern.

We also investigate violations of stochastic dominance through violations of upper and lower bounds on the bid and ask prices of options. This is a less demanding scenario in that it tests whether options, taken only one at a time, are correctly priced relative to the index. An advantage of this investigation is that the bounds are independent of assumptions regarding the allowed frequency of trading during the one-month life of each option. The observed violations of the bounds reinforce our earlier findings.

We find that the bounds are similar across the samples (namely, they are downward sloping) but the historical index sample induces bounds that are higher as the volatility is higher in the historical sample than in the forward looking samples. The flat pre-crash smile fits reasonably well within the rather high historical bounds even though these are downward sloping. The post-crash smiles are too low for the rather high location of the historical sample bounds. The forward-looking sample bounds are located somewhat lower than the historical sample bounds. Therefore, they match the also downward-sloping post-crash option prices rather well because they are located somewhat lower too. However, they do not match very well the higher horizontal smile of the pre-crash options.

One novel finding is that, even though pre-crash option prices follow the BSM model reasonably well, it does not follow that these options are correctly priced. Pre-crash option prices are incorrectly priced, if index return expectations are formed based on the historical experience. Furthermore, some of these prices are below the bounds, contrary to received wisdom that historical volatility generally underprices options in the BSM model.

Another novel finding dispels the common misconception that the observed smile is too steep after the crash. Most of the bound violations post-crash are due to the option smile not being steep enough relative to expectations on the index price formed post-crash. Even though the BSM model assumes that there is no smile, an investor who properly understood the post-crash distribution of index returns should have priced the options with a steeper smile than the smile reflected in the actual option prices.

The violations reported in this paper are based on the unconditional index return distribution. We recognize the non-stationarity of the index return distribution to the extent that we model the unconditional index return distribution after three different and representative index return samples, the historical index sample and the two forward-looking samples, one that includes the crash and one that that excludes it. It remains an open and challenging topic for future research to investigate whether the observed violations can be explained away with a more detailed model of the conditional index return distribution. It also remains a topic for future research to investigate whether priced state variables, omitted in our investigation, explain the observed violations.

Appendix

The empirical tests are based on a database containing all minute-by-minute European option quotes and trades on the S&P500 index from April 2, 1986 to December 29, 1995. We use only option quotes since we cannot know for actual trades where they occurred relative to the bid/ask spread and our results might be affected. The database also contains all futures trades and quotes on the S&P 500. Our goal is to obtain a panel of daily return observations on the index, the risk-free rate, and on several options with different strike price/index level ratios (moneyness) and constant maturity.

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 twenty minute lags of the futures price. The single regression with the highest adjusted R^2 was assumed to indicate the lag for a given day. The median lag of the index over the 1542 days from 1986 to 1992 was seven minutes. Because the index is stale, we compute a future-based index for each minute from the future market

$$S_0 = \frac{R}{1+\delta} F, \tag{A1}$$

where F is the futures price at option expiration.

For each day, we use the median interest rate implied by all futures quotes and trades and the index level at that time. We approximate the dividend yield 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.

Interest Rates. 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 used the above future-based index. We assign, for each expiration date, a single lending and borrowing rate to each day, which is the median of all daily observations across all striking prices. We then use the average of those 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.

Options with adjusted moneyness and constant maturity. It is important to use options with adjusted moneyness and constant maturity since our test statistics involve the conditional covariance matrix of option pricing errors. If the maturity of the options were not constant over time, then the conditional covariance matrix of the pricing errors would be time varying, too. This would require additional exogenous assumptions on the structure of the covariance matrix and the estimation of several additional parameters, which could lead to additional estimation error in our test statistics.

In our data set, all puts are translated into calls using European put-call parity. Then, we compute the implied volatilities where we use the Black-Scholes formula as a translation device only. We then adjust throughout each day for the movement of the stock price by assuming that the implied volatilities are independent of the underlying stock price. Then, we pick the stock price closest to 12 pm as our daily stock price and value all options from throughout the day as if they were call options with the implied volatilities measured above and struck at the moneyness level measured above. We do not eliminate any daily observations due to their level of moneyness.

Arbitrage Violations. In the process of setting up the database, we check for a number of errors, which 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 data set. We also check for violations of vertical and butterfly spreads. Within each minute we keep the largest set of option quotes which does not violate:

$$S(1+\delta) \geq C_i \geq \max[0, S(1+\delta) - K_i R] \tag{A2}$$

American early exercise is not an issue as the S&P 500 options are European in nature, 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.

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.

Amin, Kaushik 1993. "Jump Diffusion Option Valuation in Discrete Time." *Journal of Finance* 48, 1833-1863.

Bailey W. and R. M. Stulz, 1989. "The Pricing of Stock Index Options in a General Equilibrium Model." *Journal of Financial and Quantitative Analysis* 24, 1-12.

Bensaid, B., J-P. Lesne, H. Pagés and J. Scheinkman, 1992. "Derivative Asset Pricing with Transaction Costs." *Mathematical Finance* 2, 63-86.

Black, F. and M. S. Scholes, 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81, 637-654.

Brown, D. and J. C. Jackwerth, 2004. "The Pricing Kernel Puzzle: Reconciling Index Option Data and Economic Theory." Working paper, University of Wisconsin, Madison.

Constantinides, G. M., 1979. "Multiperiod Consumption and Investment Behavior with Convex Transactions Costs." *Management Science* 25, 1127-1137.

Constantinides, G. M. and S. Perrakis, 2002. "Stochastic Dominance Bounds on Derivatives Prices in a Multiperiod Economy with Proportional Transaction Costs." *Journal of Economic Dynamics and Control* 26, 1323-1352.

Constantinides, G. M. and S. Perrakis, 2004. "Stochastic Dominance Bounds on American Option Prices in Markets with Frictions." Working paper, University of Chicago and Concordia University.

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, 345-369.

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, 331-346.

Jackwerth, J. C., 2000. "Recovering Risk Aversion from Option Prices and Realized Returns." *Review of Financial Studies* 13, 433-451.

Jackwerth, J. C., 2004. "Option-Implied Risk-Neutral Distributions and Risk Aversion." ISBN 0-943205-66-2, Research Foundation of AIMR, Charlottesville, USA.

Jackwerth, J. C. and M. Rubinstein, 1996. "Recovering Probability Distributions from Option Prices." *Journal of Finance* 51, 1611-1631.

Leland, H. E., 1985. "Option Pricing and Replication with Transactions Costs." *Journal of Finance* 40, 1283-1301.

Levy, H., 1985. "Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach." *Journal of Finance* 40, 1197-1217.

Merton, R. C., 1973. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science* 4, 141-183.

Perrakis, S., 1986. "Option Bounds in Discrete Time: Extensions and the Pricing of the American Put." *Journal of Business* 59, 119-141.

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, 519-525.

Ritchken, P. H., 1985. "On Option Pricing Bounds." *Journal of Finance* 40, 1219-1233.

Ritchken, P.H. and S. Kuo, 1988. "Option Bounds with Finite Revision Opportunities." *Journal of Finance* 43, 301-308.

Rubinstein, M., 1976. "The Valuation of Uncertain Streams and the Pricing of Options." *Bell Journal of Economics* 7, 407-425.

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, No. 2, 155-188.

Ryan, P. J., 2003. "Progressive Option Bounds from the Sequence of Concurrently Expiring Options." *European Journal of Operational Research* 151, 193-223.

Silverman, B. W., 1986. *Density Estimation*, Chapman and Hall, London.

Whaley, R. E., 2003. "Derivatives." In G. M. Constantinides, M. Harris, and R. Stulz, eds., *Financial Markets and Asset Pricing: Handbook of the Economics of Finance,* Volume 1B. Handbooks in Economics 21 Elsevier/North-Holland, Amsterdam.