Mispricing of Index Options with Respect to Stochastic Dominance Bounds?

Martin Wallmeier*

October 2015

Abstract

For one-month S&P 500 index options, Constantinides, Jackwerth and Perrakis (2009) report widespread and substantial violations of stochastic dominance bounds. According to the subsequent study of Constantinides et al. (2011), the violations can be exploited to generate abnormal trading profits. The reported mispricing, which is far more extreme than known from the pricing kernel puzzle, calls into question that option markets meet the most basic requirements of rational pricing. We argue that this analysis is seriously flawed, and provide evidence that options on the S&P 500, EuroStoxx 50 and DAX index are priced almost perfectly in line with stochastic dominance bounds. Our results indicate that index option markets are much more efficient than previous literature suggests.

JEL classification: G11; G14; G24

Keywords: Index options, stochastic dominance, volatility smile, implied volatility.

*Prof. Dr. Martin Wallmeier, Department of Finance and Accounting, University of Fribourg / Switzerland, Bd. de Pérolles 90, CH-1700 Fribourg, martin.wallmeier@unifr.ch, Tel. +41 26 300 8294

1 Introduction

European index options seem to provide an ideal setting for option valuation: their payoff function is simple, the underlying asset and the characteristics of its (historical) price processes are well-known, and trading in these options has been very active for many years. Despite this, empirical evidence on the market pricing of index options is still puzzling. The ongoing debate centers around the questions of whether options are generally too expensive, whether the smile is too steep and which factors determine the cross-section of option returns.[1] Here, the "smile" or "skew" refers to an illustration of the strike price pattern of option prices in terms of implied volatilities.

A related but more fundamental question is whether option prices at least fulfill the minimum requirement of respecting the stochastic dominance bounds put forth by Constantinides and Perrakis (2002). Strikingly, Constantinides et al. (2009) (henceforth: CJP) report widespread and substantial violations of stochastic dominance by one-month S&P 500 index options over the period 1986 to 2006. The violations decrease in the 1988-1995 period, but then increase in 1997-2003, remaining at a high level until the end of the sample period. Observed deviations are substantial: scatterplots for 2000 to 2006 show quotes that are widely dispersed around the stochastic dominance bounds, partly with a majority of quotes outside the bounds.[2] The initial decrease followed by a substantial increase in violations "is a novel finding and casts doubts on the hypothesis that the options market is becoming more rational over time, particularly after the crash" (CJP, 1268f.). However, definite conclusions are difficult to draw due to concerns about data quality. The Option Metrics Database used for 1997-2006 provides more noisy data (end-of-day quotes) than the Berkeley Options Database used over the 1986-1995 period (minute-by-minute quotes and trades). Thus, the increase in violations "may be due to the lower quality of the data" (CJP, 1247), although the authors argue that the distribution of violations does not support this conjecture (CJP, 1268).

[1] For the last question, see Constantinides et al. (2013). Literature on the other research questions is briefly reviewed later in this paper.

[2] See CJP, Fig. 3 Panel F (Feb. 2000 to May 2003), where approximately three quarters of the quotes lie outside the bounds.

These results have important implications for the understanding of option markets in general. If index options are mispriced in this extreme way, the pricing of more complex options on less well-known underlying assets will presumably also be distorted. If the pricing quality of one of the most heavily traded options deteriorates over time, it seems implausible to expect a positive learning curve in other, less popular derivative markets. We might also draw the conclusion that the limits of arbitrage are extremely tight, possibly due to indirect transaction costs, low liquidity and other market frictions (Santa-Clara and Saretto (2009)). Otherwise, we would expect hedge funds and other investors to exploit and eliminate substantial violations.

However, we argue that the study of CJP suffers from two main problems which strongly distort the results. The first is to assume a constant conditional volatility for extended periods in which the level of implied volatility varied substantially. If option pricing bounds are not adjusted to the current volatility level, inevitably, many false deviations will be found. The second problem is that CJP only consider calls and do not use put-call parity to accurately match option prices with the corresponding index levels. We show that the violations disappear when these issues are corrected.

An additional concern is about data quality, because for most of the price history of S&P 500 options, transaction data are not widely used. This is why we show detailed results for options on the EuroStoxx 50 and DAX indices. Both options are of European type. In terms of trading volume, they are among the most liquid options in the world, with the EuroStoxx 50 option ranking second after the Kospi 200 option.[3] Most importantly, all time-stamped transactions for several years are available (DAX option: 1995-2014; EuroStoxx 50 option: 2000-2014), and matching with index futures prices can be conducted by milliseconds. We include all trades in put and call options on the sample days (CJP: one minute per day, only calls). Including both puts and calls is important because trading is strongly concentrated on out-of-the-money (OTM) options and the implied volatility of in-the-money (ITM) options is sensitive to small errors in the index level. In addition, we use simultaneously traded puts and calls to adjust the index levels implied in futures prices. As a result, we are able to eliminate the main sources of error and noise in implied volatilities. Based on these transaction data instead of settlement

[3] Source: Futures Industry, Annual Volume Survey, March 2009, p. 27.

data from Option Metrics, we can give more precise insight into pricing of index options with respect to the stochastic dominance bounds of Constantinides and Perrakis (2002).

Examining more than 550,000 transactions in one-month options in 1995-2014, we find almost no violations of stochastic dominance. Approximately 98% of all transactions lie within the bounds. The remaining cases can naturally be explained by a slightly different shape of the one-month index return distribution after strong increases in volatility (e.g., after the bankruptcy of Lehman Brothers in Sept. 2008). Our analysis of settlement data for S&P 500 options confirms this picture.

We point out that we do not address the question of whether the general level of option prices is appropriate.[4] Following CJP, we only examine whether the shape of the skew fits into the stochastic dominance bounds when the general level of option prices is taken as given.[5] Thus, our analysis is related to the line of research that is concerned about the slope (as opposed to the level) of the smile, building on the observation of Rubinstein (1994) and Jackwerth and Rubinstein (1996) that OTM puts are expensive compared to at-the-money (ATM) puts. Jones (2006) confirms that deep-OTM puts on S&P500 index futures are overpriced, generating negative abnormal returns even after taking volatility and jump risk premia into account. In contrast, Broadie et al. (2009) note that very high returns of deep-OTM puts alone are not inconsistent with standard option valuation models because individual option returns are extremely dispersed and highly skewed. Thus, they propose a different test approach based on market-neutral option portfolios. The main finding is that stochastic volatility alone is insufficient to explain returns of S&P 500 futures options, but models including estimation risk and jump risk premia are consistent with the data. In contrast to these studies, we focus on the more general concept of stochastic dominance without examining specific asset pricing models.

We also do not question the pricing kernel puzzle of Jackwerth (2000), Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002). The puzzle is that the empirical pricing kernel implied by

[4] Several studies show that the ATM implied volatility is an upward-biased predictor of realized volatility (see, e.g., Jackwerth and Rubinstein (1996)). Other studies find evidence of a strongly negative volatility risk premium (e.g. Chernov and Ghysels (2000), Driessen and Maenhout (2013), Santa-Clara and Yan (2010)). Selling variance swaps therefore appears to be a profitable strategy, see Carr and Wu (2009) and Hafner and Wallmeier (2007).

[5] CJP note (1266): "Since the bounds are adjusted by the implied volatility, [...] we can draw inferences about the shape of the skew but not about the general level of option prices."

estimates of the risk-neutral and objective index distributions is often found to be nondecreasing in the index level, which is inconsistent with standard economic theory. The pricing kernel is typically hump-shaped with an increase around a final index level equal to the current level. This shape is often found when the risk-neutral distribution is strongly left-skewed while the objective distribution is more symmetrical. Therefore, one possible interpretation is that the skew in option prices is too pronounced. However, this puzzle is subtle compared to the bound violations reported in CJP and studied in this paper. Many smile patterns which give rise to the pricing kernel puzzle will still fit into the stochastic dominance bounds of Constantinides and Perrakis (2002). This makes the results of CJP so surprising: if these bounds are violated, mispricing of index options must be very substantial. This is why we re-examine the case.

The next section reports the stochastic dominance bounds which are taken from prior literature. Section 3 presents our analysis of transaction data for ESX and DAX options. Section 4 discusses flaws in CJP and presents corrected results for S&P 500 options. The implications for the subsequent study of Constantinides et al. (2011) are discussed in Section 5. Section 6 concludes.

2 Stochastic dominance bounds

Constantinides and Perrakis (2002) derive bounds on call and put options in a multiperiod economy with intermediate trading and proportional transaction costs. The bounds are based on the assumption that at least one marginal investor exists whose utility of wealth is state-independent and who has a positive net exposure to the market (monotonicity of wealth condition). The upper call price bound is given by:[6]

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

where S_t is the stock price at time t, K the strike price, T the option's time to maturity, R_S the expected stock return and k the (one-way) transaction cost rate when buying and selling

[6] See Constantinides and Perrakis (2002), Proposition 1.

the index. The upper boundary of the put price, which is generally less tight, is given by:[7]

$$\overline{p}\left(S_t,t\right)=\frac{K}{R^{T-t}}+\frac{1-k}{1+k}\frac{\mathrm{E}\left[\left(K-S_T\right)^{+}-K\mid S_t\right]}{R_S^{T-t}}, \tag{2}$$

where R is the risk-free rate of return.

The lower bounds rely on the additional assumption that the investment horizon of at least one marginal investor coincides with the option's maturity date. The lower bounds are then independent of transaction costs and related by put-call parity:[8]

$$\underline{c}\left(S_t,t\right)=\frac{S_t}{\left(1+d\right)^{T-t}}-\frac{K}{R^{T-t}}+\left[\frac{\mathrm{E}\left[\left(K-S_T\right)^{+}\mid S_t\right]}{R_S^{T-t}}\right], \tag{3}$$

and:

$$\begin{aligned}\underline{p}\left(S_t,t\right)&=\frac{\mathrm{E}\left[\left(K-S_T\right)^{+}\mid S_t\right]}{R_S^{T-t}} &(4)\\ &=\underline{c}\left(S_t,t\right)+\frac{K}{R^{T-t}}-\frac{S_t}{\left(1+d\right)^{T-t}}, &(5)\end{aligned}$$

with d as dividend yield. CJP assume a transaction cost rate of 50 basis points ($k=0.005$). In reality, transaction costs for the main indices are often smaller because traders use futures as an index trading instrument. In EuroStoxx 50 and DAX futures, one-way transaction costs are typically below 10 basis points, which means that they do not strongly affect the option price bounds. For this reason, we assume $k=0$ in the empirical analysis, which has the advantage that the upper bounds are related by put-call parity in the same way as are the lower bounds (see Eq. (5)):

$$\overline{p}\left(S_t,t\right)=\overline{c}\left(S_t,t\right)+\frac{K}{R^{T-t}}-\frac{S_t}{\left(1+d\right)^{T-t}}. \tag{6}$$

Therefore, in terms of implied volatility, the bounds are identical for calls and puts. The assumption of zero transaction costs biases the results towards more frequent and more substantial violations of the (upper) bounds and is, therefore, conservative. For $k=0$, the price range be-

[7] See Constantinides et al. (2008), 584.

[8] See CJP, 1256; Constantinides and Perrakis (2002), Proposition 6; Constantinides and Perrakis (2007), 112.

tween the upper and lower bounds is determined by the market risk premium (R_S versus R):

$$\begin{aligned} & \overline{c}\left(S_t,t\right)-\underline{c}\left(S_t,t\right) \\ = \; & \overline{p}\left(S_t,t\right)-\underline{p}\left(S_t,t\right)=\frac{K}{R^{T-t}}-\frac{K}{R_S^{T-t}}. \end{aligned} \tag{7}$$

In the empirical analysis, we assume a market risk premium ($R_S - R$) of 6%. This is higher than the 4% premium of CJP, but well within the range of common estimates for the market risk premium. In most months, both rates lead to nearly identical results. In the few months with a small but discernible difference, market volatility is typically high, which suggests that the market risk premium might also be relatively high. The main conclusions are identical for a premium of 4%.

3 Violations of stochastic dominance bounds: evidence from transaction data for ESX and DAX options

3.1 Data and methodology

For a study examining option mispricing, it is crucially important to measure implied volatilities with great precision. Hentschel (2003, 788) describes the main sources of measurement error as follows: "For the index level, a large error typically comes from using closing prices for the options and index that are measured 15 minutes apart. This time difference can be reduced by using transaction prices, but such careful alignment of prices is not typical. Even when option prices and published index levels are perfectly synchronous, large indexes often contain stale component prices." We address these concerns in the following way. To overcome stale prices in the index, we derive the appropriate index level from transaction prices of the corresponding index future, which is the most common index trading instrument. We match each option trade with the previous future trade and require that the time difference does not exceed 30 seconds. In fact, the median time span between matched future and option trades in 2014 is smaller than 200 milliseconds. Even with perfect matching, the index level might still be flawed because it is not adjusted for dividends during the option's lifetime. This is particularly relevant for the ESX which is a price index, while the DAX is a performance index. Because dividend expectations of option traders are not directly observable, following Han (2008), we use put-call parity to derive a market estimate of the appropriate index adjustment.

More specifically, our procedure to measure implied volatilities is as follows (see Hafner and Wallmeier (2007)). To obtain the index level S_{mt} corresponding to an observed futures market price F_{mt} at time m on day t, we solve the futures pricing model $F_{mt} = S_{mt}e^{r(T-t)}$ for S_{mt}, where r is the risk-free rate of return and T the futures contract maturity date. We only consider the nearest available contract, which is always the one most actively traded. The futures implied index level S_{mt} is then adjusted such that transaction prices of pairs of ATM puts and calls traded within 30 seconds are consistent with put-call parity. The adjusted index level is $S_{mt}^{adj} = S_{mt} + A_t$, where A_t is the same adjustment value for all index levels observed intraday. Empirically, the adjustment is negligible with the exception of ESX options traded in March and April. The reason is that for these options, the maturity months (April and May) are different from the next maturity date of the future (June). Between the two maturity dates, most ESX firms pay out dividends, which are therefore considered differently in options and futures prices.

Following CJP, we consider options with a time to maturity of 30 calendar days.[9] In each month, there is exactly one day (a Wednesday) with this time to maturity. Thus, the sample period from 1995 to 2014 for the DAX option and from 2000 to 2014 for the ESX option consists of 240 and 180 trading days, respectively.

The stochastic dominance bounds are based on an assumed probability distribution of the underlying asset. Thus, observed violations can be explained either by option mispricing or by errors in estimating the probability distribution. In principle, any violation could be eliminated by picking the "right" distribution. To avoid this type of data snooping, we adopt the approach of CJP to estimate the shape of the unconditional distribution as the smoothed historical distribution of index returns over 1972-2006. For ESX, we use the shorter period 1987-2006 because the index was introduced only in 1998 and calculated backwards up to 1987. The historical dis-

[9] In the data section, CJP state that the retained options have a time to expiration of 30 days (p. 1257), in Appendix B the time to expiration is specified as 29 days (p. 1274).

tribution includes all intervals of 21 trading days during the estimation period. The conditional distributions are then obtained by scaling returns to be consistent with the current volatility level. More specifically, the volatility parameter is chosen such that the observed ATM implied volatility at the beginning of the day lies in the middle of the bounds implied by the conditional distribution. In this way, we control for the general level of option prices so that violations of stochastic dominance can be clearly attributed to the shape of the smile pattern. Following CJP, we de-mean the sample returns and add back the risk-free rate plus the market risk premium.

Fig. 1 (left graph) shows the conditional distribution of log DAX returns over 21 trading days, the smoothed distribution and the normal distribution with the same volatility on the last day of the sample period (Dec. 17, 2014). The distribution is skewed to the left (skewness of -1.01) and leptokurtic (excess kurtosis of 1.54). For the same day, the graph on the right shows the scatterplot of implied volatility versus moneyness for all trades in one-month DAX options, where moneyness is defined as the ratio of discounted strike price and contemporaneous index level. All trades occur within the stochastic dominance bounds indicated by the outer lines. The graph also shows the estimated regression line of a regression of implied volatility on moneyness according to:

$$IV = b_0 + b_1 M + b_2 M^2 + b_3 D M^3, \tag{8}$$

where M denotes moneyness, b_i are regression coefficients and D is a dummy variable which is one for $M > 0$ and zero otherwise. The last term is introduced to capture possible asymmetries of the smile profile for positive and negative moneyness. The mean adjusted R^2 of this regression model is 96.5% (median 97.8%). Because it precisely reflects the smile pattern, we will refer to the position of the regression line instead of single trades in one part of the empirical analysis.

The setup of our empirical study is as follows. Our sample days are those on which index options have a time to maturity of exactly 30 calendar days (one day in each month). For each sample day, we estimate implied volatilities and stochastic dominance bounds as illustrated in the right graph of Fig. 1. We analyze this information in three steps. First, by pooling all sample days together, we give an overview of the number and size of bound violations by option type (put or call) and moneyness range. Second, we examine the occurrence of violations over time. Third, we take a closer look at the days with the most significant violations.

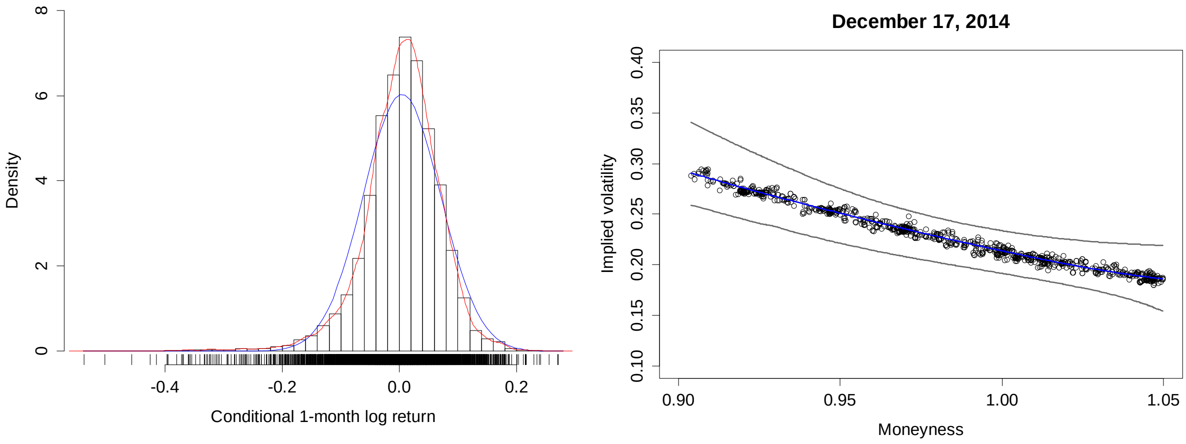

Figure 1: Conditional one-month return distribution, stochastic dominance bounds and trades on December 17, 2014

3.2 Overview of results

Table 1 reports summary statistics for the pricing of DAX options during the 1995 to 2014 period. Panel A includes all trades, while Panels B to D are based on subsamples defined by different moneyness intervals. The upper part of each panel shows the number and the percentage of trades inside and outside the stochastic dominance bounds. The lower part shows the mean size of the deviations in terms of implied volatility (column "Mean") and as a percentage of the upper or lower bound (column "in %").

As seen in Panel A, 262,504 transactions with moneyness between 0.9 and 1.05 were executed on the 240 sample days. Puts are more often traded than calls (share of 56%). The vast majority of put and call transactions (97.4% and 98.8%) are located within the bounds. Among the remaining trades, upper bound violations occur about as often as lower bound violations. The mean of the upper deviations is one percentage point, corresponding to 2.4% of the upper bound implied volatility. The lower bound deviations tend to be even smaller.

Panels B to D show that trading in low moneyness options is heavily concentrated on puts (42,673 of 45,998 transactions in Panel B), while call option trades prevail at high moneyness levels (84,119 of 105,331 transactions in Panel D). The largest share of violations is observed for OTM puts (Panel B), where 5% of trades occur outside the bounds. Lower bound violations occur more often than violations of the upper bound, which suggests that, in general, the slope of the skew pattern is not excessively high. The extent of the lower bound deviations is small, as indicated by the mean of 0.51 percentage points (2.3% of the lower bound). In the other moneyness intervals (Panels C and D), the percentage of trades violating the bounds is always below 3%.

The empirical results are very similar for ESX options, as seen in Table 2. The ESX sample includes 288,065 transactions[10] on 180 days from 2000 to 2014, of which 97.6% are located within the bounds. Most of the outside cases for OTM puts or lower bound violations. They occur when the smile is almost flat. The size of the lower bound violations is even smaller than before (mean of 0.34 percentage points).

Due to the similarity of DAX and ESX option pricing, we report the following detailed results only for the option with the longer sample period, which is the DAX option.

3.3 Timeline of violations

To illustrate the periods in which significant deviations from the stochastic dominance bounds occur, we resort to the estimated regression function of model (8) as it provides a precise description of the smile pattern. More specifically, we analyze the position of the regression function with respect to the upper and lower stochastic dominance bounds at the two moneyness levels 0.9 and 1.05. We do not choose more extreme moneyness values because, outside this range, trading becomes thin and the bounds are often uninformative (lower bounds zero and upper bound for low moneyness very high). The relative position of the regression function with respect to the bounds is measured by:

$$relPos\,(M^*) = \frac{IV_R\,(M^*) - LB\,(M^*)}{UB\,(M^*) - LB\,(M^*)}, \tag{9}$$

where $M^* \in \{0.9, 1.05\}$ is moneyness, $IV_R\,(M^*)$ is the implied volatility of the estimated regression function (8) at moneyness M^*, and $UB\,(M^*)$ and $LB\,(M^*)$ are the upper and lower bounds corresponding to moneyness M^*. The stochastic dominance bounds are respected if

[10] Trading in ESX options was thin during the first five years of the product's lifetime but then increased substantially. Since 2008, there are more transactions in ESX than DAX options. In 2014, the number of transactions in ESX options was even four times higher than that of DAX options. In spite of this development, the market for DAX options remains active with more than 1,000 transactions per sample day in 2014.

	Puts		Calls		All	
	Panel A: All transactions					
	N	in %	N	in %	N	in %
Upper violation	2'012	1.4	568	0.5	2'580	1.0
Inside bounds	142'800	97.4	114'478	98.8	257'278	98.0
Lower violation	1'827	1.2	819	0.7	2'646	1.0
Sum	146'639	100.0	115'865	100.0	262'504	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0104	2.4	0.0078	2.3	0.0098	2.4
Lower deviation IV	0.0058	2.4	0.0064	2.0	0.0060	2.3
	Panel B: 0.9 <= Moneyness < 0.95					
	N	in %	N	in %	N	in %
Upper violation	923	2.1	98	2.8	1'021	2.1
Inside bounds	42'673	95.0	3'325	94.7	45'998	95.0
Lower violation	1'326	3.0	89	2.5	1'415	2.9
Sum	44'922	100.0	3'512	100.0	48'434	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0123	2.5	0.0111	2.2	0.0122	2.5
Lower deviation IV	0.0052	2.3	0.0049	2.4	0.0051	2.3
	Panel C: 0.95 <= Moneyness < 1.0					
	N	in %	N	in %	N	in %
Upper violation	922	1.2	185	0.7	1'107	1.0
Inside bounds	78'915	98.7	27'034	99.0	105'949	98.8
Lower violation	104	0.1	92	0.3	196	0.2
Sum	79'941	100.0	27'311	100.0	107'252	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0093	2.3	0.0071	1.9	0.0089	2.3
Lower deviation IV	0.0035	1.4	0.0077	2.4	0.0055	1.9
	Panel D: 1.0 <= Moneyness <= 1.05					
	N	in %	N	in %	N	in %
Upper violation	167	0.8	285	0.3	452	0.4
Inside bounds	21'212	97.4	84'119	98.9	105'331	98.6
Lower violation	397	1.8	638	0.8	1'035	1.0
Sum	21'776	100.0	85'042	100.0	106'818	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0063	2.0	0.0070	2.5	0.0067	2.3
Lower deviation IV	0.0083	3.0	0.0064	1.9	0.0072	2.3

Table 1: DAX option pricing with respect to stochastic dominance bounds, 1995-2014.

	Puts		Calls		All	
	Panel A: All transactions					
	N	in %	N	in %	N	in %
Upper violation	766	0.4	294	0.3	1'060	0.4
Inside bounds	173'080	96.6	108'200	99.4	281'280	97.6
Lower violation	5'414	3.0	311	0.3	5'725	2.0
Sum	179'260	100.0	108'805	100.0	288'065	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0066	1.7	0.0052	1.5	0.0062	1.6
Lower deviation IV	0.0028	1.2	0.0057	1.9	0.0030	1.2
	Panel B: 0.9 <= Moneyness < 0.95					
	N	in %	N	in %	N	in %
Upper violation	248	0.4	17	2.6	265	0.4
Inside bounds	65'392	92.4	606	92.1	65'998	92.4
Lower violation	5'142	7.3	35	5.3	5'177	7.2
Sum	70'782	100.0	658	100.0	71'440	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0065	1.4	0.0079	1.8	0.0066	1.4
Lower deviation IV	0.0026	1.1	0.0041	1.6	0.0026	1.1
	Panel C: 0.95 <= Moneyness < 1.0					
	N	in %	N	in %	N	in %
Upper violation	418	0.4	78	0.5	496	0.4
Inside bounds	94'222	99.4	17'104	99.3	111'326	99.4
Lower violation	121	0.1	45	0.3	166	0.1
Sum	94'761	100.0	17'227	100.0	111'988	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0069	1.9	0.0065	1.7	0.0069	1.8
Lower deviation IV	0.0031	1.1	0.0043	1.4	0.0034	1.2
	Panel D: 1.0 <= Moneyness <= 1.05					
	N	in %	N	in %	N	in %
Upper violation	100	0.7	199	0.2	299	0.3
Inside bounds	13'466	98.2	90'490	99.5	103'956	99.3
Lower violation	151	1.1	231	0.3	382	0.4
Sum	13'717	100.0	90'920	100.0	104'637	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0053	1.5	0.0044	1.4	0.0047	1.5
Lower deviation IV	0.0091	2.9	0.0063	2.0	0.0074	2.3

Table 2: ESX option pricing with respect to stochastic dominance bounds, 2000-2014.

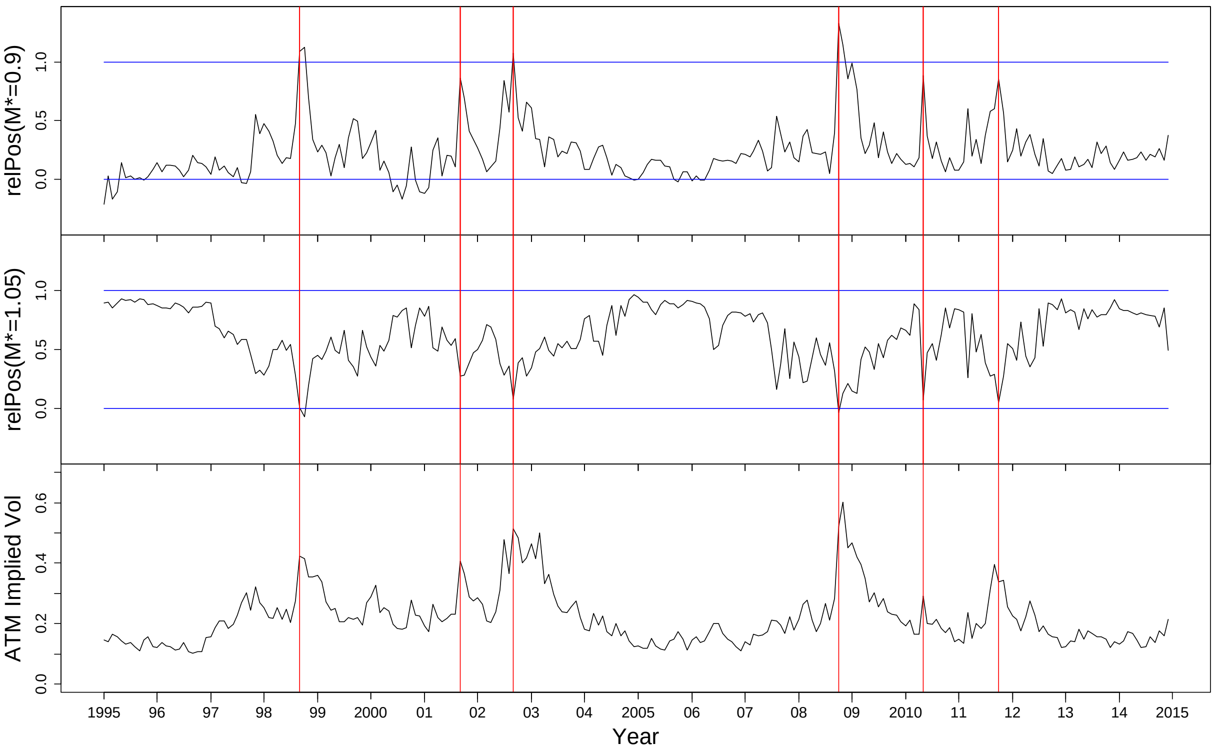

Figure 2: Position of smile regression with respect to stochastic dominance bounds at moneyness 0.9 (upper panel) and 1.05 (middle panel). The bottom panel shows the ATM implied volatility. Data are monthly, with one sample day in each month from Jan. 1995 to Dec. 2014.

$0 \leq relPos\,(M^*) \leq 1$. The cases $relPos\,(M^*) < 0$ and $relPos\,(M^*) > 1$ indicate violations of the lower and upper bound, respectively.

Fig. 2 shows the measure $relPos\,(M^*)$ over time for $M^* = 0.9$ (upper panel) and $M^* = 1.05$ (middle panel). For comparison, the bottom panel shows the concurrent ATM implied volatility as an indicator of the degree of uncertainty in the market.

Most of the time, $relPos\,(M^*)$ moves within the bounds of zero to one. The upper bound of OTM put options ($M^* = 0.9$) is violated three times. In three further events, prices come close to the upper bound. These six events, which are marked by vertical lines in Fig. 2, refer to:

1. the Russian crisis of Sept./Oct. 1998;

2. the September 11, 2001 terrorist attacks;

3. the sharp market decline of Sept. 2002;

4. the financial crisis after the bankruptcy of Lehman Brothers (Oct./Nov. 2008);
5. the high level of uncertainty in May 2010 related to the European sovereign debt crisis;
6. market movements in Oct. 2011 related to the European sovereign debt crisis.

In the middle panel for $M^* = 1.05$, these events are recognizable as downward swings towards the lower bound, although the bound is broken in only two of the events (the Russian crisis and the bankruptcy of Lehman Brothers), and only to a negligible extent. If we compare both panels, it becomes obvious that the skew pattern became more pronounced during each crisis, with OTM puts priced near the upper bound and OTM calls priced near the lower bound. It is interesting to note that option prices stayed well in-between the stochastic dominance bounds in other turbulent months during the sample period, in particular the Asian crisis of 1997, the end of the Dot-com boom in 2000 and the Iraq war in 2003.

Apart from the crisis months, the upper panel indicates that OTM puts are mostly priced close to the *lower* bound, which suggests that the smile is generally not too steep, given the historical distribution of one-month index returns. During the second half of 2000, the smile is almost flat so that the lower bound is slightly violated.[11] We show more details about this phase in the next section, after a closer look at the crisis events.

3.4 A closer look at the most significant deviations

3.4.1 Russian crisis, 9/11, Lehman bankruptcy, European sovereign debt crisis

When excluding eight months (out of 240) related to the six crisis events presented in section 3.3, more than 99% of all transactions lie within the stochastic dominance bounds, and the remaining transactions deviate by less than 0.5 percentage points of implied volatility, on average. Thus, almost all observed violations are related to the crisis events. We illustrate the corresponding smile patterns for the most significant events in more detail in Fig. 3. The left graph in each row refers to the month prior to the crisis, the right graph to the crisis month itself. The four

[11] We leave out the first few months of DAX option trading at the beginning of 1995, which were characterized by very low volatility and almost no skew. These deviations are very small in terms of implied volatility.

rows represent the Russian crisis of 1998, the 9/11 attacks, the collapse of Lehman Brothers and the European sovereign debt crisis.

In each event, implied volatilities jump upwards (higher level of the skew in the right graphs compared to the left). The structure of implied volatilities across moneyness remains highly regular in the crisis months, but the skew becomes steeper, and at both ends it protrudes beyond the bounds range. Therefore, OTM puts appear to be too expensive and OTM calls too cheap, but the deviations remain so small that a higher-than-usual downside risk could easily explain the observed patterns.

Given the uncertainty about the conditional index return distribution it is natural to find a certain number of deviations from bounds which are based on a specific distributional assumption. In crisis events, skewness and kurtosis are presumably different than on average.[12] In addition, we still lose precision in our analysis by holding conditional volatility constant during the day. By updating volatility following intraday changes of ATM implied volatility, the number of violations would further decrease.[13] Given these considerations, we interpret the empirical evidence as almost perfectly in line with stochastic dominance bounds.

In a related paper on the pricing of American-type S&P 500 *futures* options, Constantinides et al. (2011) take estimation errors of the return distribution into account so that stochastic dominance in a strict sense can no longer be identified. However, the bounds can still serve as a means to identify potential mispricing. Constantinides et al. (2011) find that a corresponding trading strategy actually provides significant abnormal returns. In our case, as Fig. 3 illustrates, such a strategy would imply selling OTM put options in the most extreme market situations. This strategy will be high-risk, no matter how it is implemented. To make things worse, during the sample period of 20 years, there are fewer than ten independent trading opportunities, namely the crisis events, with implied volatility deviations above one percentage point. In this setting, it is clearly beyond the power of any statistical test to find evidence of significant abnormal returns. The observed violations are far too small and too rare to be able to devise a profitable trading strategy.

[12] Kozhan et al. (2013) show that skew risk is closely related to variance risk.

[13] Intrady movements of the skew profile generally take the form of parallel shifts, see Wallmeier (2015).

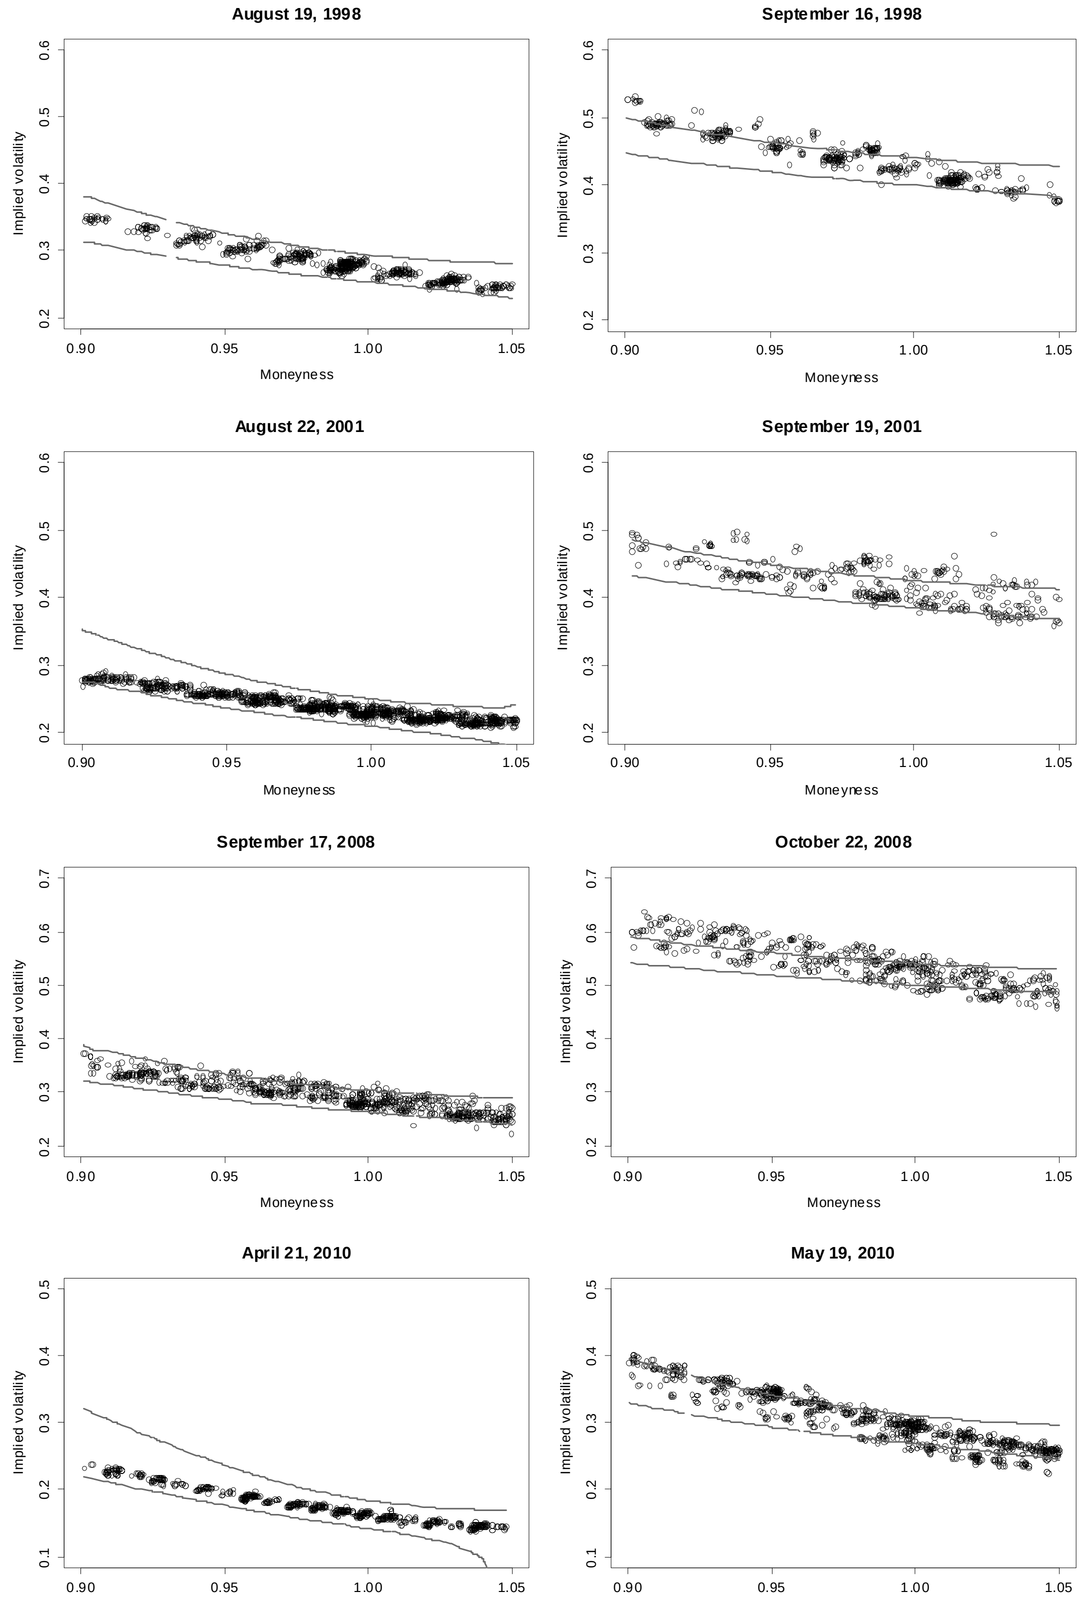

Figure 3: Smile scatterplots in times of crisis. Left: previous month, right: month of crisis event. The four rows represent the Russian crisis of 1998, the 9/11 attacks, the collapse of Lehman Brothers and events related to the European debt crisis.

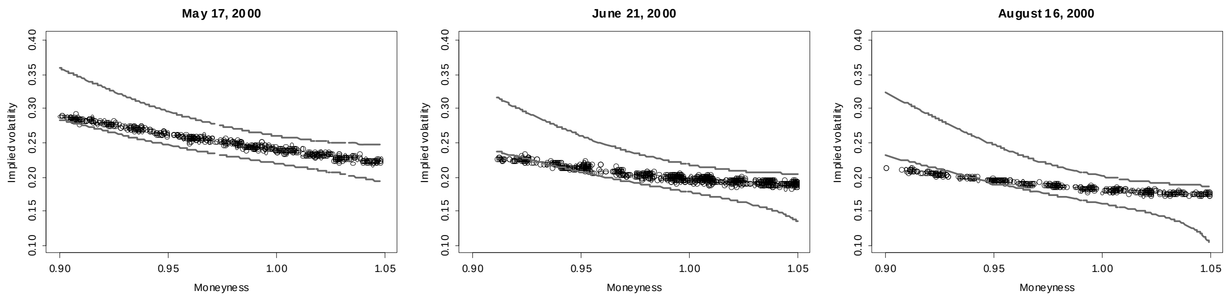

Figure 4: Flattening smile in the second half of 2000.

3.4.2 Periods without a (pronounced) skew pattern

In the second half of 2000 until February 2001, the *lower* bound of OTM puts is violated (see upper panel in Fig. 2). Fig. 4 illustrates the transactions from May, June and August 2000. In the scatterplot for May 2000, OTM puts are priced close to the lower bound, but all trades stay within the bounds range. Over the next three months, volatility decreases further and the smile continues to flatten out. The OTM put premium appears to be too low but again, deviations are small. One obvious possibility is that market participants in this period expected the return distribution over the next month to be close to normal, so that the implied volatilities were almost flat.

4 Critique of Constantinides/Jackwerth/Perrakis (2009)

To understand why the results of CJP are so different, we replicate their analysis for S&P 500 options over the last two subperiods (February 2000 to May 2003; June 2003 to May 2006). As in CJP, our data are end-of-day bid and ask quotes for call and put options from Option Metrics. We only consider options with positive trading volume on that day.

The results of CJP are shown in their Figures 3 and 4. In Figure 4 of CJP for 2003-2006, three properties stand out: First, there is a large number of bound violations. Second, the pattern is strikingly irregular compared to smile graphs shown in this paper so far; in particular, a cluster of observations with moneyness between 0.95 and 1 and implied volatility below 10% does not seem to fit into any familiar smile pattern. Third, there are many cases of arbitrage violations in which implied volatility could not be computed (marked on the horizontal axis). We argue that all three properties are non-existent: in our own analysis, there are no bound violations, the irregularities disappear, and arbitrage violations do not occur.

The reason for these differences is that the analysis of CJP suffers from two main problems. The first is that the underlying index level is not adjusted to be consistent with put-call parity. The second and most important problem is that the bounds are not adjusted to the current level of conditional volatility. Details on both aspects follow.

Put-call parity

Settlement data for option prices and index levels are typically not perfectly synchronous. In addition, the index level has to be adjusted for expected dividend payments during the option's lifetime.[14] Small adjustment errors will produce substantial errors in implied volatilities. The standard approach is to infer the underlying index level from put-call parity (see Aït-Sahalia and Lo (1998), Fan and Mancini (2009), Chen and Xu (2014)). CJP, however, try to adjust the *interest rates*, which is difficult to justify per se and also does not work well.[15]

We illustrate the importance of this step in Fig. 5 for the last day of the sample period of CJP, which is May 17, 2006. The left graph shows the smile pattern based on the closing index level on that day (1270.32), the middle graph shows implied volatilities provided by Option Metrics, and the right graph shows the smile for the underlying index level which is consistent with put-call-parity (1264.10). The differences are large, especially when only call options are included as in CJP.[16] The smiles in the left and middle panels are clearly misestimated. Therefore, it is

[14] CJP infer the closing index levels from closing futures prices. In this way, the index level is adjusted for expected dividends until the futures maturity date. In some months, a mismatch occurs because the maturity dates of options and futures deviate (e.g. option maturity April, next future maturity June).

[15] For data from the Berkeley Options Database (1986-1995), CJP "compute implied interest rates embedded in the European put-call parity relation" (p. 1273). For data from the Option Metrics Database, the authors "cannot arrive at a consistently positive interest rate implied by option prices [...] and use T-bill rates instead" (p. 1274).

[16] The distorted patterns for calls in the left and middle panels of Figure 5 are characterized by: inconsistent quotes (marked on the x-axis); partly decreasing implied volatilities for moneyness between 0.95 and 1; and an overall flat pattern. These characteristics are present in CJP but not in our analysis. CJP state: "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." We find a significant smile in each month, as in the right panel of Figure 5.

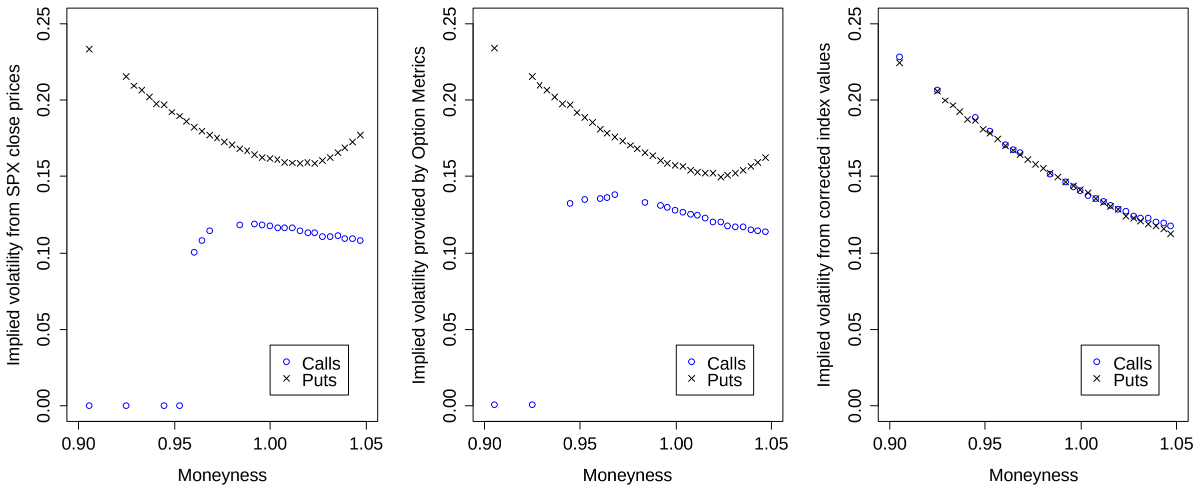

Figure 5: Smile of S&P500 options on May 17, 2006 (time to maturity: 30 days). Option values are midpoints of bid and ask quotes. Left panel: implied volatilities based on the closing index level of 1270.32. Middle panel: implied volatilities provided by Option Metrics (Option Price Files). Right panel: implied volatilities based on an adjusted underlying index value of 1264.10; the adjustment of −6.22 corresponds to −0.49% of the closing index level.

crucially important to ensure put-call parity in estimating the smile.

Taking put-call parity into account,[17] our versions of Figures 3 and 4 in CJP are shown in the upper two graphs of Fig. 6. For moneyness below (above) 1, we use bid and ask quotes of put (call) options. A comparison of our scatterplot for the period 2003 to 2006 (upper right graph of Fig. 6) with Figure 4 in CJP reveals that the irregularities and arbitrage violations have disappeared.

Conditional volatility

The objective (in CJP and this study alike) is to adjust the bounds to the current volatility level so that the test is on the *shape* of the skew instead of its *level*: "Since the bounds are adjusted by the implied volatility, irrespective of whether this volatility is rational or not, we can draw

[17] More specifically, the put-call parity-consistent underlying index level for a given trading day is determined as follows: For each strike K_i with $0.95 \leq K_i \leq 1.02$, we define $A_i = (C_i - P_i) + K_i \cdot \exp(-rT)$, where C_i is the mid quote for a call option with strike K_i, P_i is the corresponding put option mid quote, r is the riskless rate of return and T the options' time to maturity. We use the mean A_i value as the adjusted underlying index level. All implied volatilities for puts and calls are based on this adjusted level.

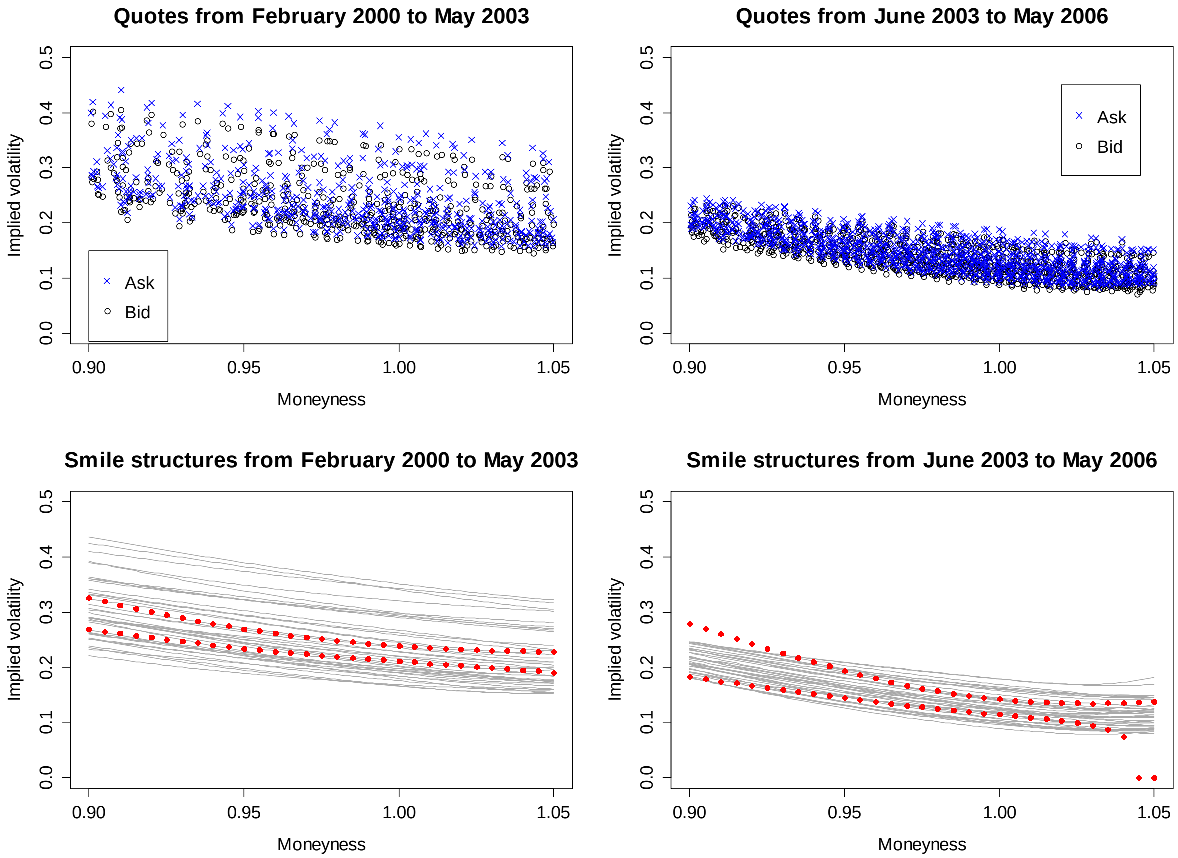

Figure 6: The upper graphs show implied volatilities based on bid and ask quotes of out-of-the money options (puts: moneyness ≤ 1, calls: moneyness > 1). The underlying index level is adjusted such that put-call parity holds for the mid quotes. The lower graphs show the smile regression lines for the months of the sample period. The red dotted lines show bounds based on an average conditional volatility as in CJP.

inferences about the shape of the skew but not about the general level of option prices." (p. 1266) However, in fact, CJP set the conditional volatility equal to the *average* implied volatility over subperiods of up to three years. Within these subperiods, the conditional volatility and the stochastic dominance bounds are assumed to be constant.

In reality, volatility varies considerably over such an extended period as is illustrated in the two lower graphs of Fig. 6. The lines show the estimated smile regressions for each month in the sample period (left graph: 2000-2003, right graph: 2003-2006).[18] The dotted (red) lines show bounds computed as in CJP (based on average volatility). It is apparent from the almost parallel skew patterns, that the volatility level shifted substantially during the subperiods, especially from 2000 to 2003. As the bounds used in CJP are not adjusted accordingly, many upper and lower deviations are observed. However, these clearly reflect the failure to adjust the bounds to the current volatility level rather than violations of stochastic dominance. CJP state: "The figures provide a clearer picture. [...] The decrease in violations over the 1988-1995 postcrash period [...] is followed by a substantial increase in violations over 1997-2003 [...]. This is a novel finding and casts doubts on the hypothesis that the options market is becoming more rational over time" (p. 1268f.) In contrast to this view, the higher incidence of apparent violations over 2000-2003 is merely due to the fact that volatility varied more strongly over this period, so that the false assumption of constant volatility had more serious consequences. In the analysis of CJP, even the high implied volatilities after the September 11 attacks are classified as upper bound violations because they do not fit into the average bounds range.

In another part of the paper, CJP adjust conditional volatility monthly. It is estimated by GARCH (1,1) or set equal to either ATM implied volatility or the VIX index. Our approach is different: we choose conditional volatility such that the ATM implied volatility lies in the middle between the upper and lower stochastic dominance bound.[19] This is the *only* estimate which guarantees that we reproduce the general level of option prices, as intended. In practice, this estimate can easily be updated at short intraday intervals (e.g., each minute) based on current ATM implied volatility.

[18] The regression function is the same as in Section 3.1. Implied volatilities (dependent variable) are based on mid quotes. We include all call and put options over the moneyness range from 0.9 to 1.05.

[19] This conditional volatility is closely related but not identical to the ATM implied volatility.

Revised results

In Fig. 8 and 9 in the Appendix, we show the smile graphs with upper and lower stochastic dominance bounds for each month of the last subperiod of CJP (June 2003 to May 2006). We also include the smile regression line according to equation (8). The implied volatilities are based on midpoints of bid and ask quotes, and the graphs include call and put options. As in Section 3, we use the conservative assumption of zero transaction costs and assume a risk premium of 6%. As before, the unconditional distribution is the smoothed historical distribution of S&P 500 returns for 1972-2006.

The main result is that all skew patterns fit perfectly into the bounds range, which is fundamentally different from the results reported by CJP.

Further tests in CJP

Two further tests of CJP examine if the empirical pricing kernel is a decreasing function of the index return. Previous studies had to reject this hypothesis, which gave rise to the *pricing kernel puzzle* (Jackwerth (2000), Aït-Sahalia and Lo (2000), Rosenberg and Engle (2002)). The pricing kernel tests of CJP rest on much more restrictive assumptions than the test of bound violations presented so far. One additional assumption is that there is at least one trader who is marginal in the entire cross section of option prices instead of one option at a time. More importantly, intermediate option trading is excluded or restricted to one intermediate point in time, which is a severe restriction given the continuous trading of index options. This is one reason why we do not replicate these specific tests. The more important reason, however, is that we do not question the phenomenon of non-monotonic empirical pricing kernels.[20] Quite the contrary: it is easy to verify that the typical smile patterns do not pass the pricing kernel test even if they fully respect the stochastic dominance bounds of Constantinides and Perrakis (2002).[21]

[20] Beare and Schmidt (2014) provide recent evidence that this phenomenon can be exploited to construct a portfolio of options whose return stochastically dominates the market return. This result does not contradict our finding that the stochastic dominance bounds of Constantinides and Perrakis (2002) hold.

[21] In this statement, we do not consider transaction costs, consistent with our analysis of stochastic dominance bounds.

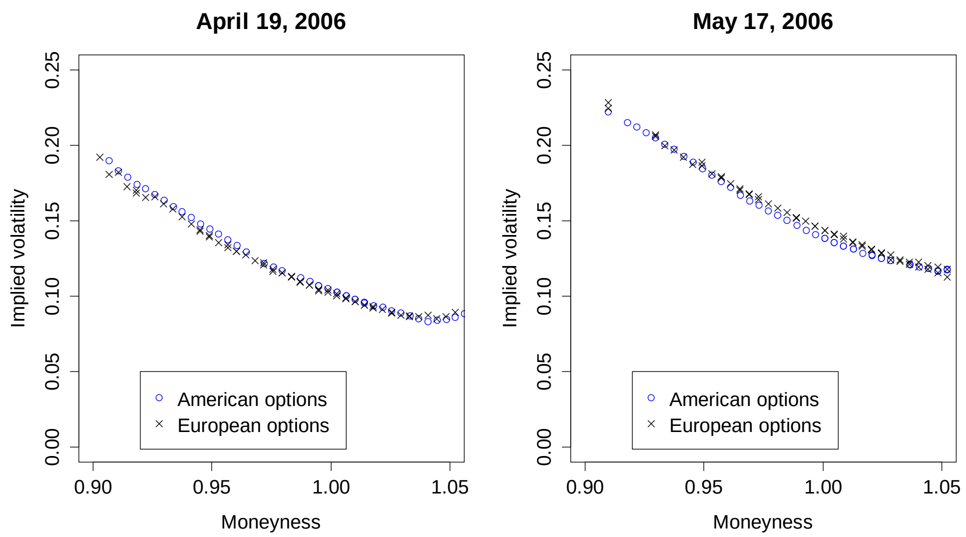

Figure 7: Comparison of smile patterns of American-type S&P 500 futures options (traded at CME) with SPX options (European type index options). The implied volatilities of the American options are based on settlement prices, while those of the European options are based on midpoints of bid and ask quotes. For American options, we apply the approximate valuation model of Barone-Adesi and Whaley (1987).

5 Implications for Constandinides et al. (2011)

In a subsequent study,[22] Constantinides et al. (2011) aim to address the concern that the reported violations in CJP do not account for potential errors in the estimation of the upper and lower stochastic dominance bounds.[23] The authors examine the significance of the violations by testing whether they can be exploited to generate abnormal trading profits.

Clearly, the finding that no bound violations are observed (as in the previous section) would remove the basis of the test. We have to consider, however, that Constantinides et al. (2011) examine *American*-type S&P 500 *futures* options instead of *European*-type S&P 500 *index* options as in CJP. One might argue that the former can be mispriced with respect to stochastic dominance bounds even if the European options are not. We argue, however, that this is highly implausible if not practically impossible. The difference in the underlying asset should be ir-

[22] The authors are the same with the inclusion of M. Czerwonko whose research assistance is acknowledged in CJP.

[23] See Constantinides et al. (2011), p. 1409f. The authors are the same with the addition of M. Czerwonko whose research assistance is mentioned in CJP.

relevant for the question at hand because the future value is tightly linked to the index level. Moreover, the valuation effect of the early exercise option, while non-negligible, is known to be small for short-term options (see Whaley (1986)). In economic terms, the two options are extremely similar, which will be reflected in similar smile patterns.

We confirm this expectation for the last two months of the sample period in CJP (April and May 2006) using settlement data for American S&P 500 futures options from the Chicago Mercantile Exchange (CME).[24] (We have data for this option only since April 2006). The options have a time to maturity of 30 days. We compute implied volatilities in the same way as before, but based on the Barone-Adesi and Whaley (1987)-approximation of American option values. This means that we again derive the underlying asset value from put-call parity for ATM options. While this put-call-parity relation is not a strict theoretical requirement as in the case of European options, it is a plausible approximation given the small valuation effect of the early exercise option on ATM options.[25] Fig. 7 shows the overlay of the smile patterns of European S&P 500 index options and American S&P 500 futures options. As expected, the patterns are so similar that we can safely conclude that the American smile fits into the bounds range just like the European smile (see Section 4). This is true even if we disregard the fact that the bounds for American options are much wider than the bounds for European options.[26] There is no reason to believe that the conclusion would be different in other months.[27]

Constantinides et al. (2011), however, report that a large proportion of call bid quotes violates the call upper bound for American options (e.g., 30.5% of all bid quotes in the moneyness range 1.01-1.03, when conditional volatility is assumed to be an adjusted implied volatility, see Constantinides et al. (2011), Table II). This clearly cannot be reconciled with the evidence presented in this paper. The estimation of implied volatilities in Constantinides et al. (2011) is similar to that of CJP: there is no comparison of put and call implied volatilities to account

[24] I would like to express my gratitude to Thomson Reuters (Markets) SA for providing the data.

[25] Ramaswamy and Sundaresan (1985) note: "Therefore, the value added by the 'American' feature is rather small, especially for options that are at-the-money." Our procedure ensures that the put-call parity relation for American futures options (see Whaley (1986), p. 52) is respected.

[26] See the illustration for one date (May 22, 1996) in Constantinides et al. (2011), Fig. 1. Even at-the-money, the span between the tightest bounds (call upper bound and put lower bound) is huge, with about 15 percentage points of implied volatility. (The actual quotes are not shown in the graph.)

[27] We can verify this for months after the sample period of CJP, but do not have data before May 2006.

for errors in the underlying asset value (e.g., caused by imprecise dividend estimates), and conditional volatility is not guaranteed to be consistent with the overall level of option prices.

6 Conclusion

For S&P 500 options, CJP report widespread violations of the stochastic dominance bounds put forth by Constantinides and Perrakis (2002). While it is well-known that index option pricing gives rise to the pricing kernel puzzle, the mispricing documented in CJP is far more extreme and calls into question that option markets meet even the most basic requirements of rational pricing. We provide new evidence based on a much more comprehensive and precise database of index options on the EuroStoxx 50 and the DAX index. Our results are opposite to the prior evidence of CJP for the S&P 500 index option. Even without considering transaction costs, option prices are almost perfectly in line with the stochastic dominance restrictions of Constantinides and Perrakis (2002). Approximately 98% of all 550,000 sample transactions in 1995-2014 are located within the bounds. The rare cases of systematic violations can be attributed to crisis events such as the bankruptcy of Lehman Brothers in September 2008. The pricing pattern in these months is still very regular and can naturally be explained by a slightly different shape of the one-month index return distribution.

There is no reason to believe that index options are priced so differently in Europe compared to the US. Therefore, we reconsider the case of S&P 500 options and find that the analysis in CJP is seriously flawed. The problems are not related to low quality of the data but to the data analysis. In our revised analysis based on the same data, no violations of the stochastic dominance bounds of Constantinides and Perrakis (2002) are observed. This has direct implications for the subsequent study of Constantinides et al. (2011): without bound violations, there is no basis for their test of whether violations can be exploited to generate abnormal trading profits.

References

Aït-Sahalia, Y. and Lo, A. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* **53**: 499–547.

Aït-Sahalia, Y. and Lo, A. (2000). Nonparametric risk management and implied risk aversion, *Journal of Econometrics* **94**: 9–51.

Barone-Adesi, G. and Whaley, R. (1987). Efficient analytic approximation of american option values, *Journal of Finance* **42**: 301–320.

Beare, B. and Schmidt, L. (2014). *Empirical implications of the pricing kernel puzzle for the return on contingent claims*, Working Paper University of California.

Broadie, M., Chernov, M. and Johannes, M. (2009). Understanding index option returns, *Review of Financial Studies* **11**: 4493–4529.

Carr, P. and Wu, L. (2009). Variance risk premiums, *Review of Financial Studies* **22**: 1311–1341.

Chen, S. and Xu, Z. (2014). On implied volatility for options - some reasons to smile and more to correct, *Journal of Econometrics* **179**: 1–15.

Chernov, M. and Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation, *Journal of Financial Economis* **56**: 407–458.

Constantinides, G., Czerwonko, M., Jackwerth, J. and Perrakis, S. (2011). Are options on index futures profitable for risk-averse investors? Empirical evidence, *Journal of Finance* **66**: 1407–1437.

Constantinides, G., Jackwerth, J. and Perrakis, S. (2008). Option pricing: Real and risk-neutral distributions, *in* J. Birge and V. Linetsky (eds), *Handbooks in OR & MS Vol. 15*, pp. 565–591.

Constantinides, G., Jackwerth, J. and Perrakis, S. (2009). Mispricing of S&P 500 index options, *Review of Financial Studies* **22**: 1247–1277.

Constantinides, G., Jackwerth, J. and Savov, A. (2013). The puzzle of index option returns, *Review of Asset Pricing Studies* **3**: 229–257.

Constantinides, G. and Perrakis, S. (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. and Perrakis, S. (2007). Stochastic dominance bounds on American option prices in markets with frictions, *Review of Finance* **11**: 71–115.

Driessen, J. and Maenhout, P. (2013). The world price of jump and volatility risk, *Journal of Banking & Finance* **37**: 518–536.

Fan, J. and Mancini, L. (2009). Option pricing with model-guided nonparametric methods, *Journal of the American Statistical Association* **104**: 1351–1372.

Hafner, R. and Wallmeier, M. (2007). Volatility as an asset class: European evidence, *European Journal of Finance* **13**(7): 621–644.

Han, B. (2008). Investor sentiment and option prices, *Review of Financial Studies* **21**: 387–414.

Jackwerth, J. (2000). Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* **13**: 433–451.

Jackwerth, J. and Rubinstein, M. (1996). Recovering probability distributions from option prices, *Journal of Finance* **51**: 1611–1631.

Jones, C. (2006). A nonlinear factor analysis of S&P 500 index option returns, *Journal of Finance* **61**: 2325–2363.

Kozhan, R., Neuberger, A. and Schneider, P. (2013). The skew risk premium in the equity index market. Working Paper.

Ramaswamy, K. and Sundaresan, S. (1985). The valuation of options on futures contracts, *Journal of Finance* **11**: 1319–1340.

Rosenberg, J. and Engle, R. (2002). Empirical pricing kernels, *Journal of Financial Economics* **64**: 341–372.

Rubinstein, M. (1994). Implied binomial trees, *Journal of Finance* **69**: 771–818.

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

Santa-Clara, P. and Yan, S. (2010). Crashes, volatility, and the equity premium: Lessons from S&P 500 options, *Review of Economics & Statistics* **92**: 435–451.

Wallmeier, M. (2015). Smile in motion: An intraday analysis of asymmetric implied volatility, *Algorithmic Finance* **4**: 89–104.

Whaley, R. (1986). On valuing American futures options, *Financial Analysts Journal* (3): 49–59.

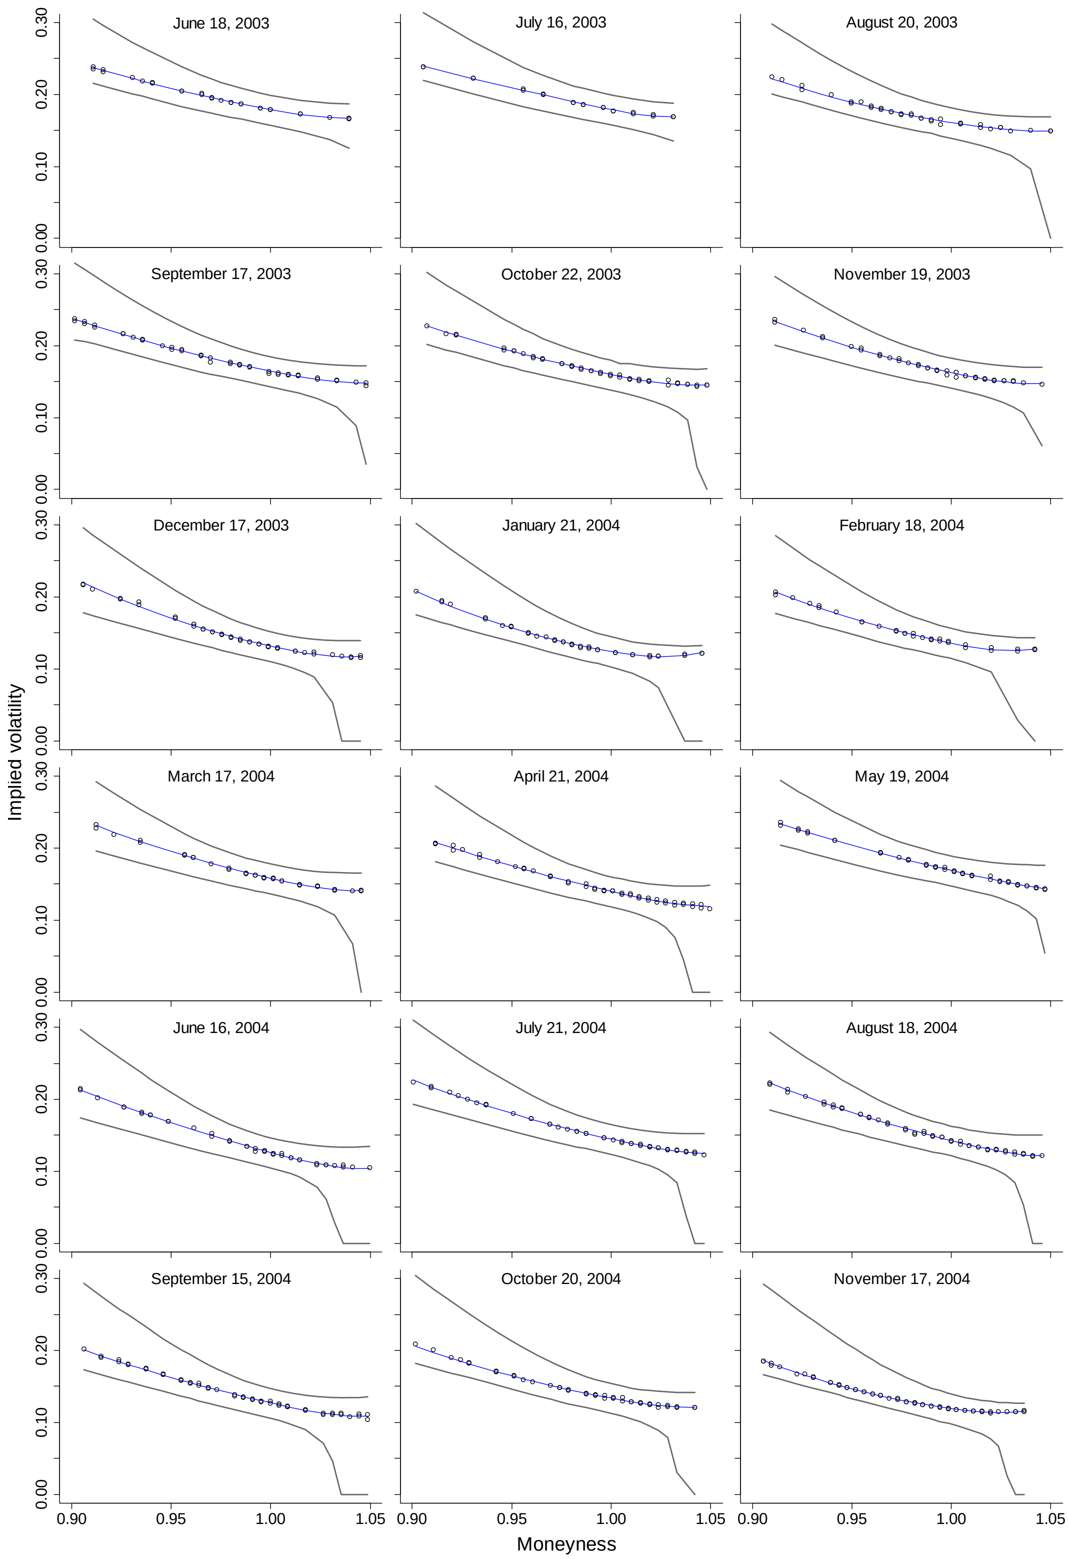

Figure 8: Smile patterns and stochastic dominance bounds of SPX options with a time to maturity of 30 days from June 2003 to November 2004. All puts and calls are included. Implied volatilies are based on mid quotes.

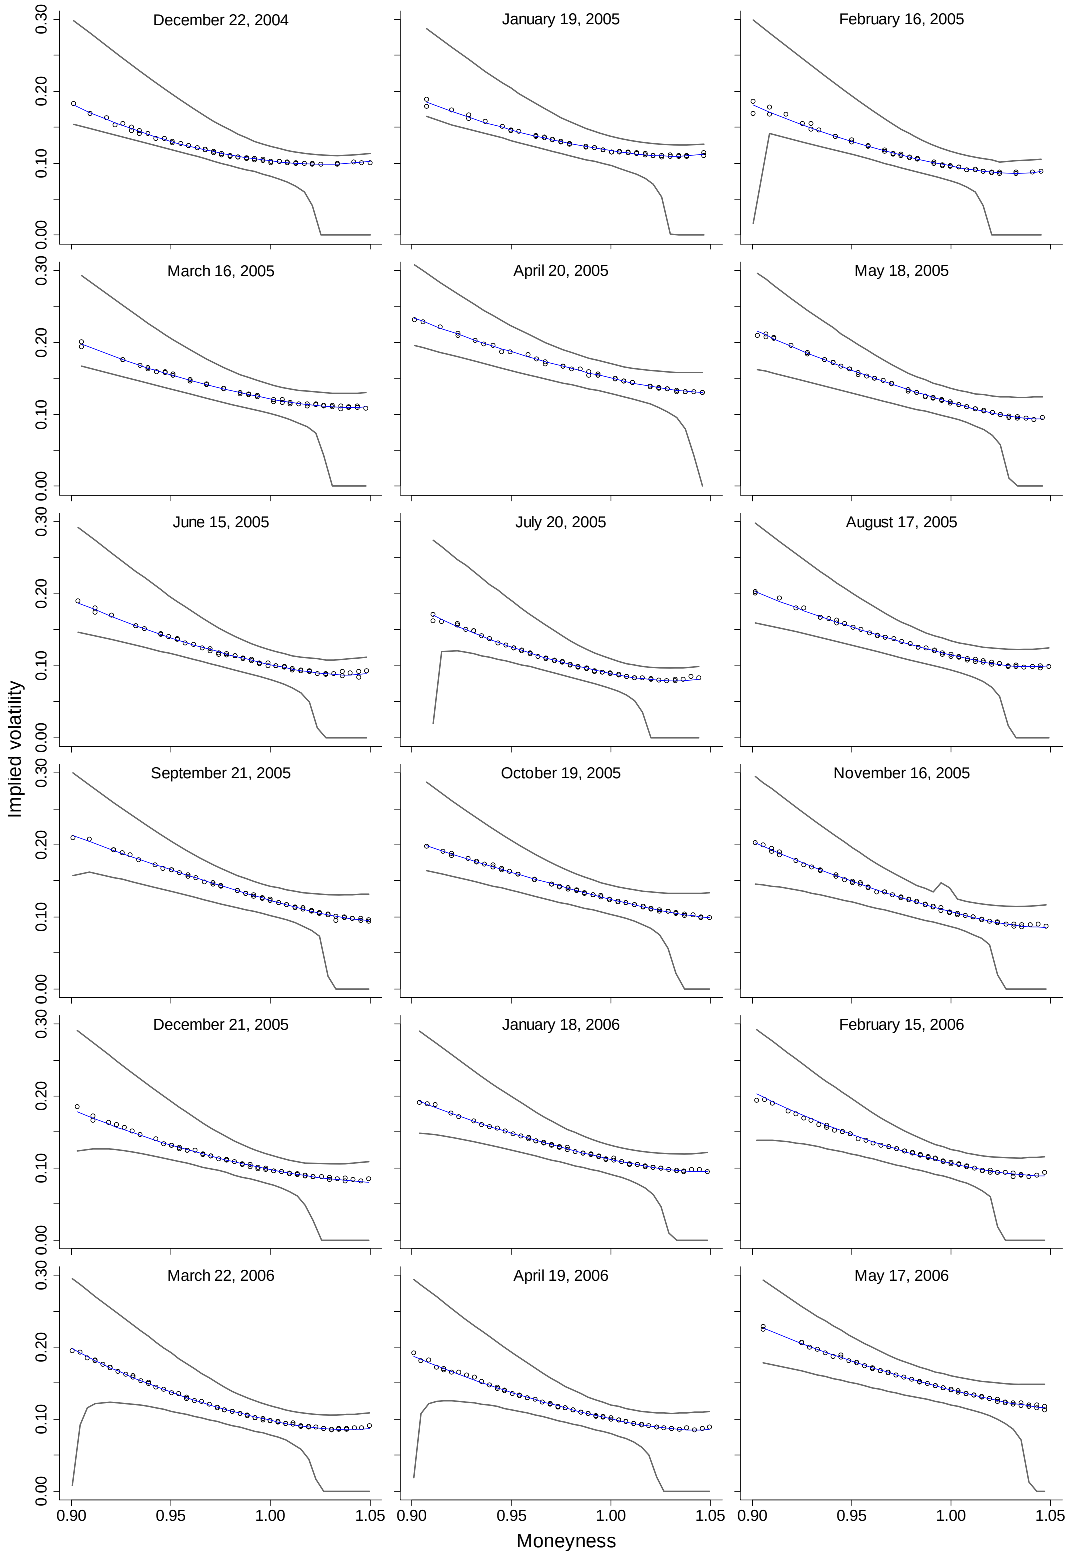

Figure 9: Smile patterns and stochastic dominance bounds of SPX options with a time to maturity of 30 days from December 2004 to May 2006. All puts and calls are included. Implied volatilies are based on mid quotes.