

Mispricing of Index Options with Respect to Stochastic Dominance Bounds?

June 2017

Martin Wallmeier

*University of Fribourg, Bd de Pérolles 90, CH-1700 Fribourg, Switzerland. Email:
martin.wallmeier@unifr.ch*

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. However, we find that index options on the S&P 500, EuroStoxx 50 and DAX are priced almost perfectly in line with stochastic dominance bounds when adjusting for (a) the general level of option prices, (b) conditional volatility and (c) put-call parity in order to determine the appropriate (dividend-adjusted) underlying index level. Our results indicate that index option markets might be much more efficient than previous literature suggests.

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

JEL: G11, G14, G24

1. Introduction

European index options seem to provide an ideal setting for option valuation: their pay-off 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, “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 (henceforth: SD 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 to 1995 period, but then increase in 1997 to 2003, remaining at a high level until the end of the sample period. Observed deviations are large: scatterplots for 2000 to 2006 show quotes that are widely dispersed around the SD 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 OptionMetrics Database used for 1997 to 2006 provides more noisy data (end-of-day quotes) than the Berkeley Options Database used over the 1986 to 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).

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

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

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

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.

This paper reconsiders the question of whether index options violate SD bounds and provides new insight into the nature of potential mispricing. We show that index options on the S&P 500, EuroStoxx 50 and DAX are priced almost perfectly in line with SD bounds when (a) considering conditional volatility, (b) adjusting the bounds for the general level of option prices and (c) using put-call parity to estimate the dividend-adjusted underlying index level. Condition (b) means that the conditional volatility is adjusted such that the average at-the-money (ATM) implied volatility lies in the middle of the bounds range. Under conditions (a) to (c), more than 96% of option transactions lie within the bounds. The remaining cases can naturally be explained by a slightly different shape of the one-month index return distribution in times of market stress (e.g., after the bankruptcy of Lehman Brothers in Sept. 2008).

The SD bounds in our tests are affected by estimation errors. Therefore, the failure of finding bound violations does not imply that no dominating option trades exist. According to our results, if substantial mispricing is observed, it can be attributed to deviations from the above conditions (a) to (c), which reflect either estimation errors or “real” mispricing in the sense of irrational market behavior. Our results suggest that index option prices are consistent with put-call parity (condition (c)), but we do not address the question of whether the general level of option prices is appropriate (condition (b)).³ Following CJP, we only examine whether the shape of the skew fits into the SD bounds when the general level of option prices is taken as given.⁴ Thus, our analysis is related to the line of research that is

³Several 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).

⁴CJP, 1266, note: “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.”

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 out-of-the-money (OTM) puts are expensive compared to 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 address only the part of the CJP study that analyzes the Constantinides and Perrakis (2002) bounds. 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 violations of the Constantinides and Perrakis (2002) bounds. 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.⁵ 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 SD bounds of Constantinides and Perrakis (2002). 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

⁵Beare 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 SD bounds of Constantinides and Perrakis (2002) hold.

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.

The next section reports the SD bounds analyzed in this paper. Section 3 presents our analysis of transaction data for SPX, ESX and DAX options. Section 4 compares our results with CJP and shows the impact of differences in the study designs. Section 5 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 is⁶

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

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 the index. The upper boundary of the put price, which is generally less tight, is⁷

$$\bar{p}(S_t, t) = \frac{K}{R^{T-t}} + \frac{1-k}{1+k} \frac{\mathbb{E}[(K - S_T)^+ - K | S_t]}{R_S^{T-t}}, \quad (2)$$

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

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

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

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

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

and

$$\underline{p}(S_t, t) = \frac{\mathbb{E}[(K - S_T)^+ | S_t]}{R_S^{T-t}} \quad (4)$$

$$= \underline{c}(S_t, t) + \frac{K}{R^{T-t}} - \frac{S_t}{(1+d)^{T-t}}, \quad (5)$$

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 index trading instrument. For the main index 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)):

$$\bar{p}(S_t, t) = \bar{c}(S_t, t) + \frac{K}{R^{T-t}} - \frac{S_t}{(1+d)^{T-t}}. \quad (6)$$

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

$$\begin{aligned} & \bar{c}(S_t, t) - \underline{c}(S_t, t) \\ &= \bar{p}(S_t, t) - \underline{p}(S_t, t) = \frac{K}{R^{T-t}} - \frac{K}{R_S^{T-t}}. \end{aligned} \quad (7)$$

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

3.1. Data and methodology

3.1.1. Estimation of the strike price profile of implied volatilities

For a study examining option mispricing, it is crucially important to measure implied volatilities with great precision. Hentschel (2003, 788) describes the main source 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.” To ensure synchronicity, we rely on transaction prices. For SPX options, we use the concurrent S&P 500 index values reported by CBOE in the trade records files. For ESX and DAX options, we derive the appropriate index level from transaction prices of the corresponding index futures. We match each option trade with the previous futures trade and require that the time difference does not exceed 30 seconds. In fact, the median time span between matched futures 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 SPX and ESX options which are based on price indices, while 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 (2000), Hafner and Wallmeier (2007)). The matched index level S_{mt} at time m on day t is 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 value for all index levels observed during the day. Fig. 1 illustrates this adjustment for trades of SPX options with a time to maturity of 30 days on January 22, 2014 (for a similar example, see Hafner and Wallmeier (2000)). The left graph shows implied volatilities based on the unadjusted intraday index levels. In this graph, put-call parity appears to be violated, with put options (black crosses) trading at higher implied volatilities than call options (blue circles). However, when lowering all intraday index levels by a constant of 3.62 points, the recomputed implied volatilities line up as shown in the right graph. Note the conversion of three call options that initially appear to have negative time

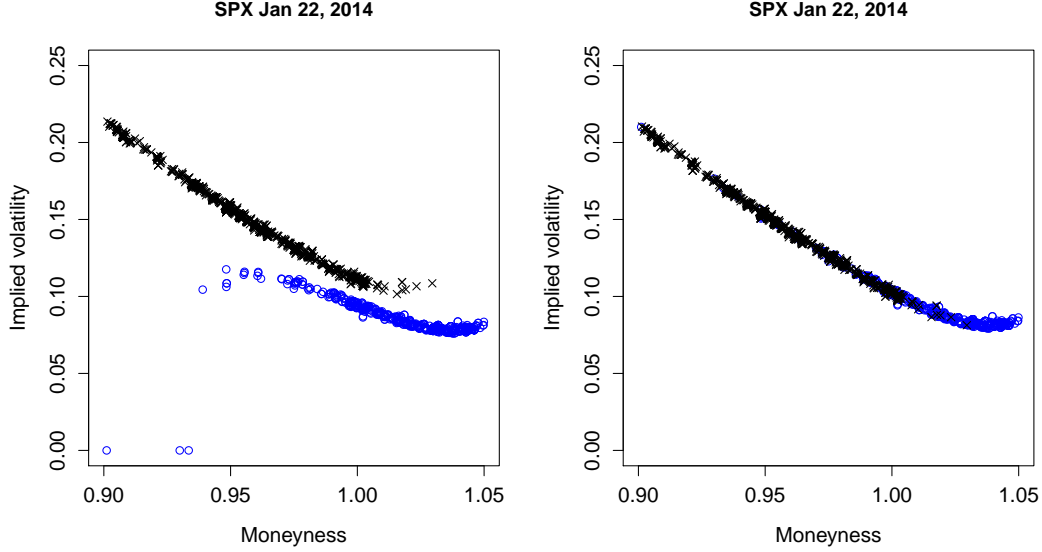


Figure 1: Smile profile and put-call parity

Description: The graphs show the smile for transactions in SPX options with a time to maturity of 30 calendar days on January 22, 2014. Black crosses: put options, blue circles: call options. Left graph: implied volatilities based on synchronized intraday index levels provided in the trade files of CBOE. Right graph: implied volatilities based on intraday index levels reduced by 3.62 index points.

Interpretation: The reduction by 3.62 reflects expected dividends until the maturity date. With dividend-adjusted index levels, the smile profiles of call and put options coincide, which is consistent with put-call parity.

values (shown with implied volatility of 0.00 in the left graph); after the adjustment, they perfectly fit into the smile profile. The situation is similar on all days of our sample period and for all three index options, i.e. the adjusted smile profiles of put and call options always coincide and negative time values are no longer observed.⁹

For SPX options, the adjustment is always negative, which is consistent with nonzero monthly expected dividend payments. Thus, without the adjustment, put options would always appear to be more expensive than call options. Our adjustments closely mirror the series of actual dividend payments, which corroborates our interpretation that the adjust-

⁹We ignore trades in SPX options in the first minute of trading (8.30 a.m. to 8.31 a.m.) each day because at this time, the index level sometimes appears to still include outdated stock prices. For example, on Dec. 17, 2014 options appear to be priced on the basis of an underlying index level that is 4.5 points above the reported index level.

ments capture anticipated dividend discounts.¹⁰ Also in line with this interpretation, we find that, typically, no adjustment is necessary for options on the performance index DAX. For ESX options, the adjustment is mostly negligible except in March and April. In these months, the option maturity months (April and May) are different from the next maturity date of the futures (June). Between the two maturity dates, most EuroStoxx 50 firms pay out dividends, which are therefore considered differently in options and futures prices. For this reason, the use of futures prices instead of index levels does not circumvent the problem of dividend adjustments.

On some days, at-the-money SPX options show an unusual implied-volatility pattern. In the Appendix, we give examples and discuss how we address the issue. The pattern arises from trades that are recorded with exactly the same option price—often an integer value—at different intraday levels of the underlying index. Our explanation is that these trades are part of a combined option strategy such as a collar (index futures plus long put plus short call).¹¹ In such a case, buyer and seller agree upon a price for the package (e.g., \$1 collar price) without specifying the component prices. The reporting system, however, allows only simple put and call option trades. Therefore, the collar price has to be decomposed. To simplify the entry, an integer value is often used for one price component. For example, a collar price of \$1 might be recorded as \$34 for the embedded long put and \$33 for the embedded short call. When the collar price decreases to \$0.50, the recorded put price might be kept constant at \$34 while the call price is adjusted to \$33.50 or the put price is reduced to \$33.50 while the call price is left at \$33. The bottom line is that the recorded prices are not informative if their connection is lost. The CBOE trade files do not allow for identifying combined trades. Therefore, we apply a simple identification rule that removes the pattern reasonably well (see details in the Appendix). For our study, this issue is of minor importance because the phenomenon is clearly visible on only a few days. Our results and conclusions remain the same without any attempt to remove these transactions.

¹⁰For 2010 to 2014, the series of dividends for the S&P 500 index was: 23.12; 26.02; 30.44; 34.99; 39.44 (source: Bloomberg). Our cumulative index adjustments for the same years are (absolute values): 24.96; 27.96; 29.62; 30.20; 38.43.

¹¹A collar based on ATM options provides a riskless position. This trade can be used to exploit possible deviations of ATM options from put-call parity, thereby enforcing an appropriate parity relationship. In the following, we use the term “collar” for the embedded options without considering the index investment.

3.1.2. Study design

Following CJP, we consider options with a time to maturity of 30 calendar days.¹² 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 SPX and ESX options consists of 240 and 180 trading days, respectively. For SPX options, the underlying index values are missing in the trade files for May and June 2003 so that our final SPX sample includes 178 trading days.

The SD 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 to 2006. For ESX, we use the shorter period 1987 to 2006 because the index was introduced only in 1998 and calculated backwards up to 1987. The historical distribution 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 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. 2 (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 SD bounds indicated by the

¹²In 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).

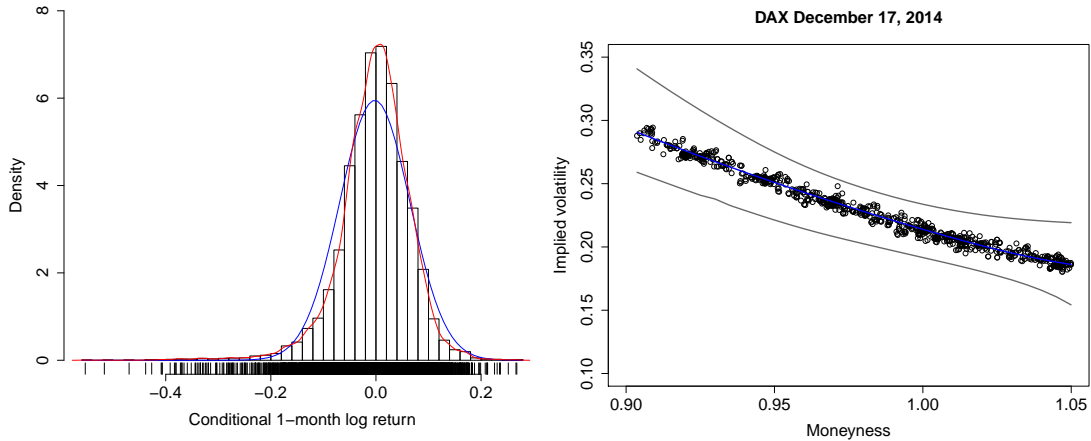


Figure 2: Identifying violations of stochastic dominance bounds

Description: The left graph shows the conditional one-month DAX return distribution. The right graph shows the corresponding stochastic dominance bounds and trades of DAX options with 30 days time to maturity on December 17, 2014.

Interpretation: On this day, all transactions lie within the bounds (no violations).

outer lines. The graph also shows the estimated regression line of the regression

$$IV = b_0 + b_1M + b_2M^2 + b_3DM^3, \quad (8)$$

where the moneyness measure M is defined as the logarithmic ratio of discounted strike price and contemporaneous index level, divided by the square root of time to maturity, b_i are regression coefficients and D is a dummy variable which is one for $M > 0$ and zero otherwise.¹³ 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 higher than 95% for SPX, ESX and DAX options. Because the regression line precisely reflects the smile profile, we will refer to its position 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 SD bounds as illustrated in the right graph of Fig. 2. We analyze this information in three steps. First, by pooling all sample days

¹³See Hafner and Wallmeier (2007). Note that the time-to-maturity adjusted moneyness measure M is only used in this regression. The smile graphs in this paper are based on moneyness defined as the ratio of discounted strike price and index level.

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.

3.2. Overview of results

Tables 1 to 3 report summary statistics for the pricing of SPX options (2000-2014), DAX options (1995-2014) and ESX options (2000-2014). In each case, 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 of Table 1 for SPX options, 305,164 transactions with moneyness between 0.9 and 1.05 are included for the sample period of 178 days. Puts are more often traded than calls (share of 57%). The vast majority of put and call transactions (96.1% and 97.1%) are located within the bounds. Among the remaining trades, lower bound violations occur slightly more often than upper bound violations. The mean of the lower deviations is 1.42 percentage point, corresponding to 5.3% of the lower bound implied volatility. The upper bound deviations tend to be even smaller.

Panels B to D show that trading in low moneyness options is heavily concentrated on puts (55,516 of 58,592 transactions in Panel B), while call option trades prevail at high moneyness levels (95,969 of 125,086 transactions in Panel D). OTM puts more often violate the bounds than OTM calls (4.0% vs. 2.5%). In the middle moneyness interval (Panel C), the proportion of trades inside the bounds is almost the same for puts (96.9%) and calls (96.4%).

The empirical results are very similar for DAX and ESX options, as seen in Tables 2 and 3. The DAX (ESX) sample includes 262,504 (288,065) transactions¹⁴ on 240 (180) days from 1995 to 2014 (2000 to 2014), of which 98.0% (97.6%) are located within the bounds. The size of the remaining bound violations is even smaller than for SPX options. In all, index

¹⁴Trading 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.

option prices generally appear to be very well aligned with SD bounds.

Due to the similarity of the index options in Europe (DAX and ESX), we hereafter omit the one with the shorter sample period, which is the ESX option. Thus, we report the following detailed results only for SPX and DAX options.

3.3. Timeline of violations

To illustrate the periods in which significant deviations from the SD 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 SD 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:

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

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

Fig. 3 illustrates the position of the smile pattern over time for DAX options (upper graph) and SPX options (lower graph). Both graphs show the measure $relPos(M^*)$ for $M^* = 0.9$ in the upper panel and for $M^* = 1.05$ in the middle panel. The bottom panel shows the ATM implied volatility as an indicator of the degree of uncertainty in the market.

The trajectories for DAX and SPX options are remarkably similar. Most of the time, $relPos(M^*)$ moves within the bounds of zero to one. Six times, the upper bound of OTM put options ($M^* = 0.9$) is violated or prices come close to the upper bound. These six events, which are marked by vertical lines in Fig. 3, 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;

Table 1: SPX option pricing with respect to stochastic dominance bounds, 2000-2014

Description: For put and call options in different moneyness classes, the table reports the proportion and size of bound violations.

Interpretation: 96.5% of all transactions lie within the bounds. The remaining deviations are small.

	Puts		Calls		All	
Panel A: All transactions						
	N	in %	N	in %	N	in %
Upper violation	4,880	2.8	1,390	1.1	6,270	2.1
Inside bounds	166,243	96.1	128,288	97.1	294,531	96.5
Lower violation	1,887	1.1	2,476	1.9	4,363	1.4
Sum	173,010	100.0	132,154	100.0	305,164	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0133	3.4	0.0144	4.0	0.0136	3.5
Lower deviation IV	0.0156	5.5	0.0131	5.2	0.0142	5.3
Panel B: $0.9 \leq \text{Moneyness} < 0.95$						
	N	in %	N	in %	N	in %
Upper violation	2,159	3.9	168	5.5	2,327	4.0
Inside bounds	53,311	96.0	2,793	90.8	56,104	95.8
Lower violation	46	0.1	115	3.7	161	0.3
Sum	55,516	100.0	3,076	100.0	58,592	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0147	3.3	0.0164	4.0	0.0148	3.4
Lower deviation IV	0.0075	2.7	0.0264	10.6	0.0210	8.3
Panel C: $0.95 \leq \text{Moneyness} < 1.0$						
	N	in %	N	in %	N	in %
Upper violation	2,330	2.6	573	1.7	2,903	2.4
Inside bounds	85,655	96.9	31,917	96.4	117,572	96.8
Lower violation	392	0.4	619	1.9	1,011	0.8
Sum	88,377	100.0	33,109	100.0	121,486	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0114	3.1	0.0129	4.0	0.0117	3.3
Lower deviation IV	0.0167	6.7	0.0193	8.9	0.0183	8.1
Panel D: $1.0 \leq \text{Moneyness} \leq 1.05$						
	N	in %	N	in %	N	in %
Upper violation	391	1.3	649	0.7	1,040	0.8
Inside bounds	27,277	93.7	93,578	97.5	120,855	96.6
Lower violation	1,449	5.0	1,742	1.8	3,191	2.6
Sum	29,117	100.0	95,969	100.0	125,086	100.0
	Mean	in %	Mean	in %	Mean	in %
Upper deviation IV	0.0172	5.5	0.0153	4.1	0.0160	4.6
Lower deviation IV	0.0155	5.2	0.0101	3.5	0.0126	4.3

Table 2: DAX option pricing with respect to stochastic dominance bounds, 1995-2014

Description: For put and call options in different moneyness classes, the table reports the proportion and size of bound violations.

Interpretation: 98.0% of all transactions lie within the bounds. The remaining deviations are small.

	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: $0.9 \leq \text{Moneyness} < 0.95$						
	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: $0.95 \leq \text{Moneyness} < 1.0$						
	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: $1.0 \leq \text{Moneyness} \leq 1.05$						
	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 3: ESX option pricing with respect to stochastic dominance bounds, 2000-2014

Description: For put and call options in different moneyness classes, the table reports the proportion and size of bound violations.

Interpretation: 97.6% of all transactions lie within the bounds. The remaining deviations are small.

	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: $0.9 \leq \text{Moneyness} < 0.95$						
	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: $0.95 \leq \text{Moneyness} < 1.0$						
	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: $1.0 \leq \text{Moneyness} \leq 1.05$						
	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

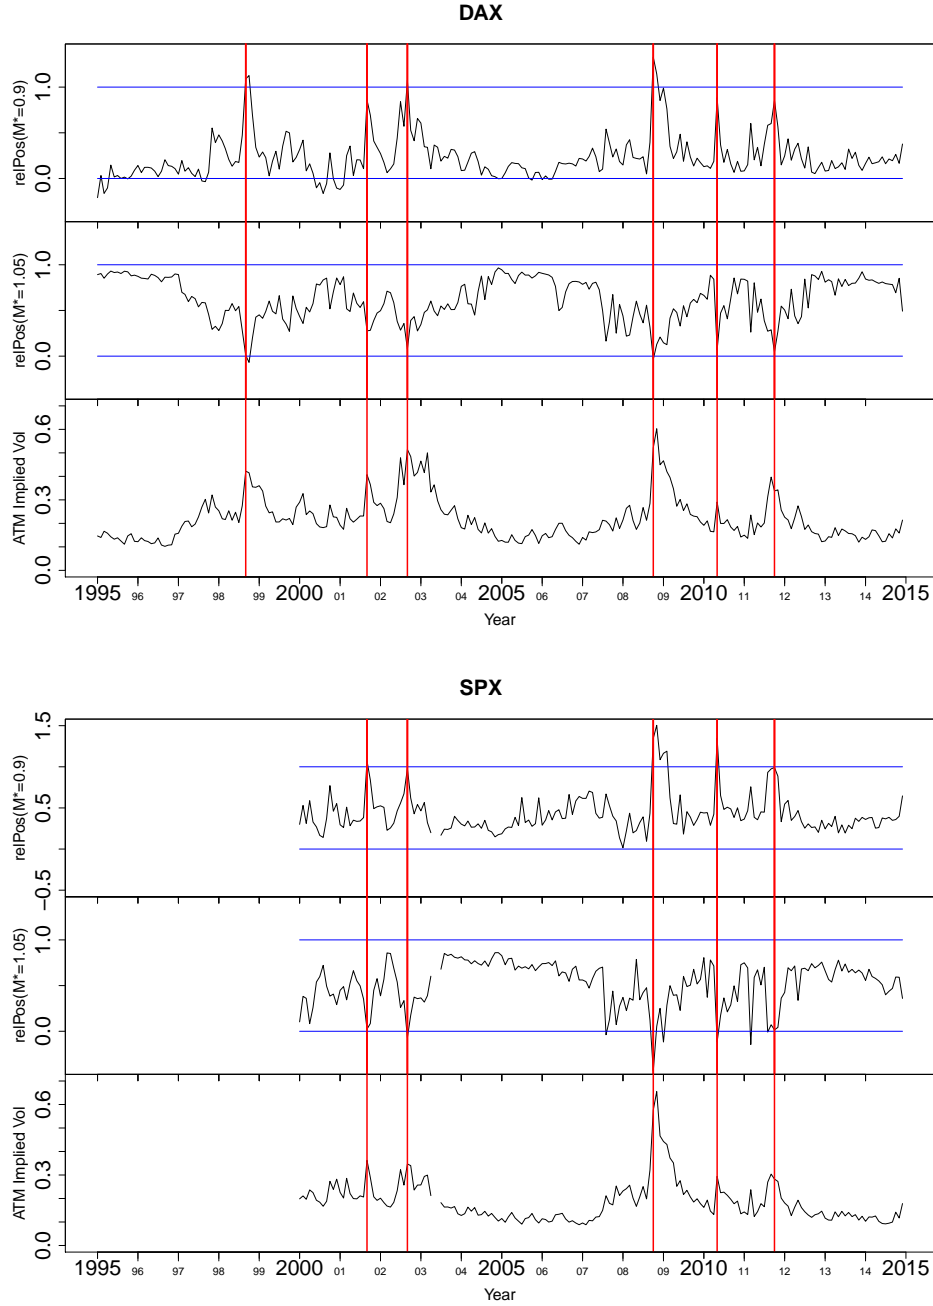


Figure 3: Position of the smile of DAX and SPX options with respect to stochastic dominance bounds over time

Description: For each option, the upper two panels show the position of the smile regression with respect to stochastic dominance bounds at moneyness 0.9 and 1.05. The bottom panel shows the ATM implied volatility. The vertical lines indicate crisis events. Data are monthly, with one sample day per month. The options have a time to maturity of 30 days.

Interpretation: Most of the time, the smile profile is located within the bounds. In times of market stress, the skew tends to become steeper so that it may breach the upper bound at a low moneyness level and the lower bound at a high moneyness level.

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. If we compare both panels, it becomes obvious that the skew profile 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 SD 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 (DAX) or near the middle of the range (SPX), 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 DAX smile is almost flat so that the lower bound is slightly violated.¹⁵ 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 related to the six crisis events presented in section 3.3, more than 98% of all transactions (SPX and DAX) lie within the SD bounds and the remaining transactions deviate by less than 0.7 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. 4 (SPX) and 5 (DAX). The left graph in each row refers to the month prior to the crisis, the right graph to the crisis month itself. The four rows represent the Russian crisis of 1998 (only DAX), 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

¹⁵We 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.

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 times of market stress, skewness and kurtosis are presumably different than on average.¹⁶

It is also important to note that 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. Fig. 6 illustrates the intraday changes of the SPX smile pattern for the first sample day after the bankruptcy of Lehman Brothers (Oct. 22, 2008, upper panel) and the last day of our sample period (Dec. 17, 2014, lower panel). For Oct. 22, 2008, the graph on the right picks out the transactions between 2 p.m. and 3 p.m. and highlights transactions with strikes 850 and 900. Implied volatilities in this hour were much higher than average implied volatilities during the day so that most trades lie outside the SD bounds representing the average situation of the day. The highlighted observations for a constant strike price are upward sloping. For a given strike, increasing moneyness reflects a falling index level, which in turn is associated with higher implied volatilities. Typically, the intraday shifts of the smile pattern are almost parallel (see Wallmeier (2015)). This can also be seen in the lower right graph for Dec. 17, 2014, which depicts all daily transactions and highlights strikes from 1,850 to 2,050 in steps of 50. Again, the upward sloping patterns for a given strike indicate parallel shifts of the smile in an inverse relationship to the index level. The observed violations would mostly disappear when adjusting the bounds to the changing intraday volatility level. Given these considerations, we interpret the empirical evidence as almost perfectly in line with SD 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,

¹⁶Kozhan et al. (2013) show that skew risk is closely related to variance risk.

as Fig. 4 and 5 illustrate, 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. Thus, in our analysis, observed violations are far too small and too rare to be able to devise a profitable trading strategy.

3.4.2. *Periods without a (pronounced) skew*

In the second half of 2000 until February 2001, the *lower* bound of OTM puts is violated for DAX options (see upper panel in Fig. 3). Fig. 7 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. Comparison with Constantinides et al. (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 OptionMetrics. We consider only 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% do not seem to fit into familiar smile patterns. Third, there are many cases of arbitrage violations in which implied volatility could not be computed (marked on the horizontal axis).

We find that these three properties disappear when (1) put-call parity is considered, and (2) the bounds are adjusted to conditional volatility. In the following, we give further details

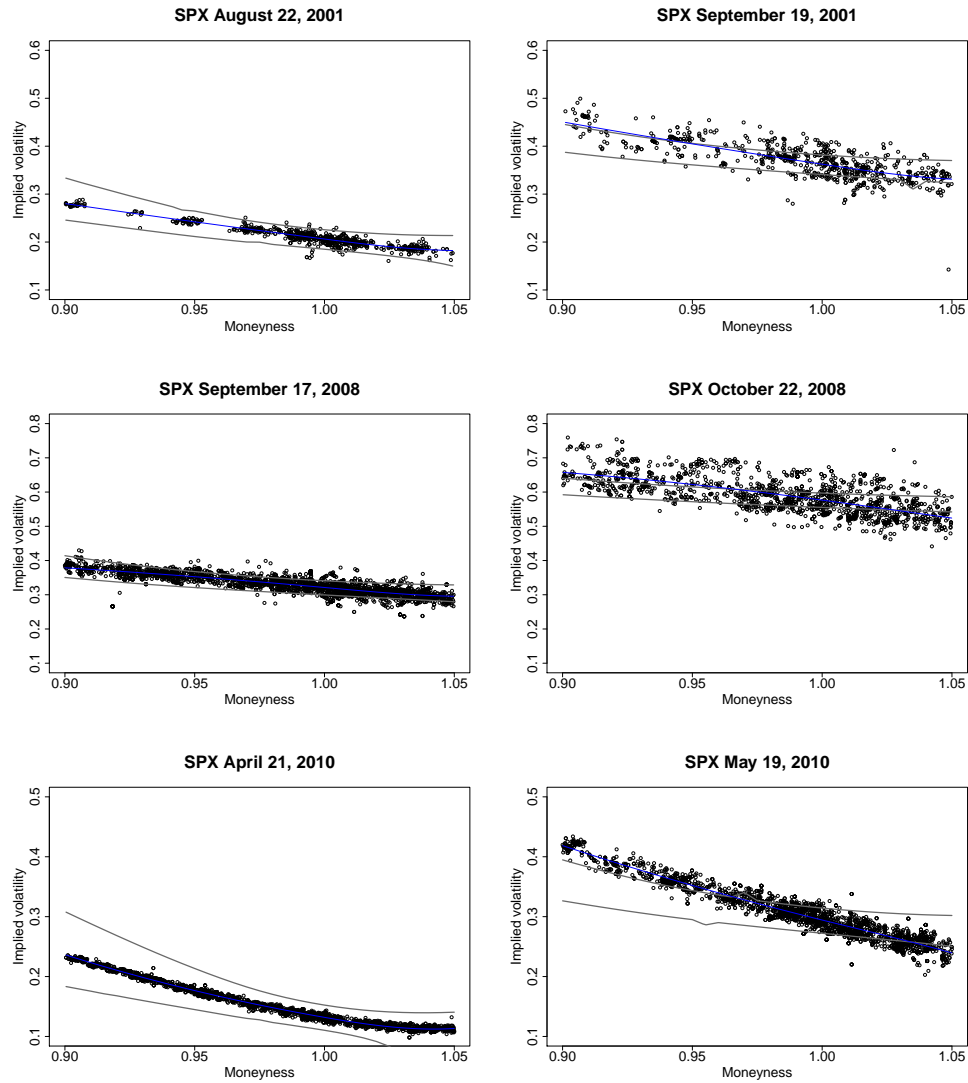


Figure 4: Smile profiles of SPX options before and after crisis events

Description: The left graphs show the smile profiles in the month before the crisis event, the right graphs the first smile profile affected by the crisis event. The three rows represent the 9/11 attacks, the collapse of Lehman Brothers and events related to the European debt crisis.

Interpretation: In times of market stress, the smile profile shifts upwards and becomes steeper.

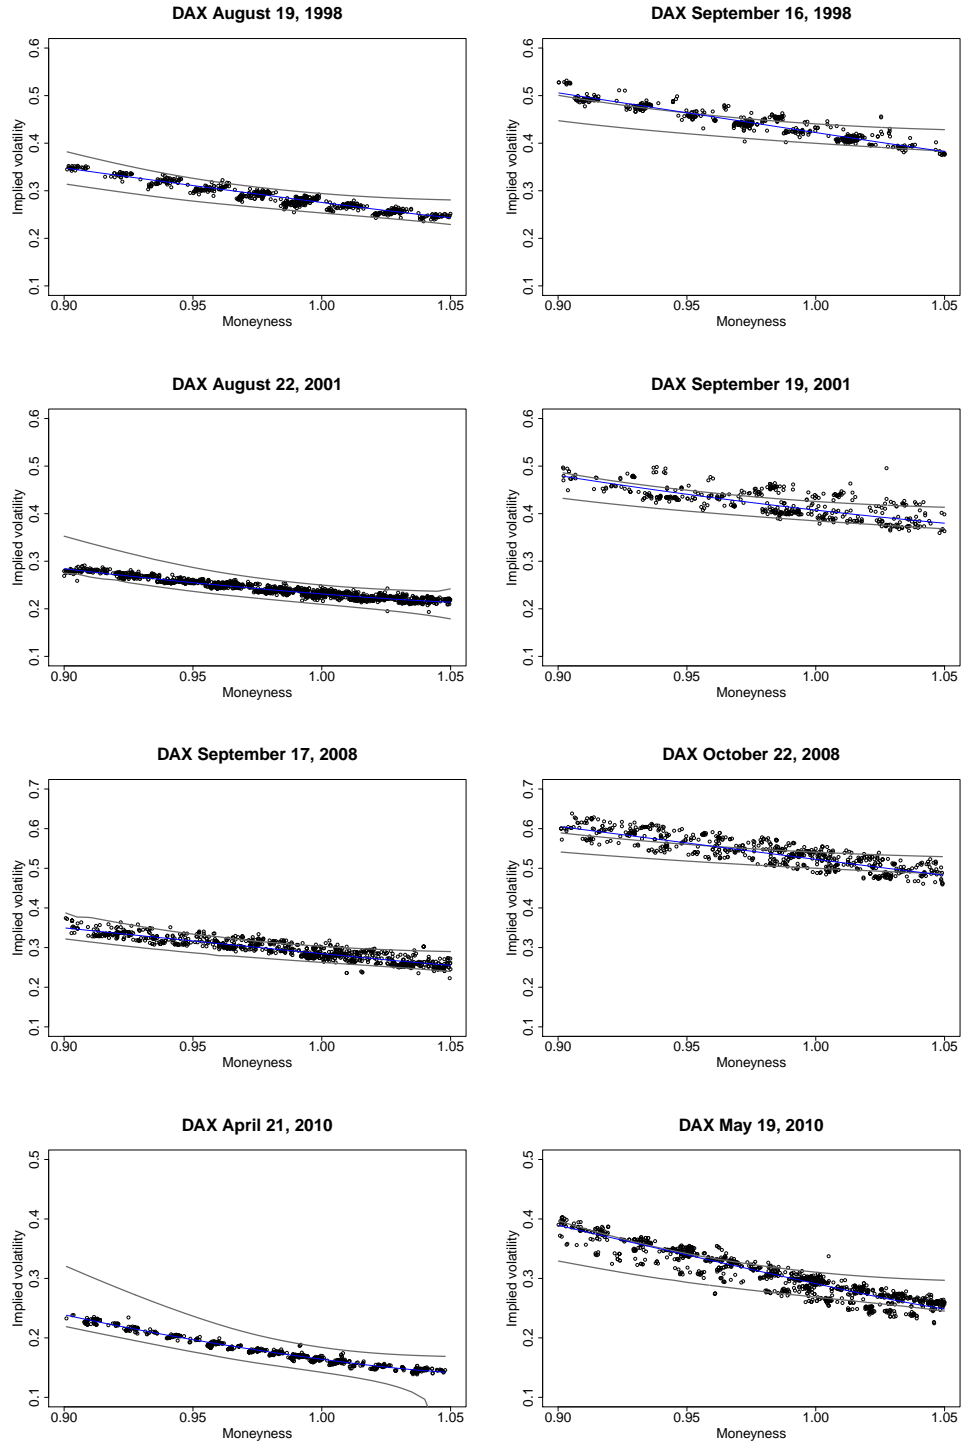


Figure 5: Smile profiles of DAX options before and after crisis events

Description: The left graphs show the smile profiles in the month before the crisis event, the right graphs the first smile profile affected by the 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.

Interpretation: In times of market stress, the smile profile shifts upwards and becomes steeper.

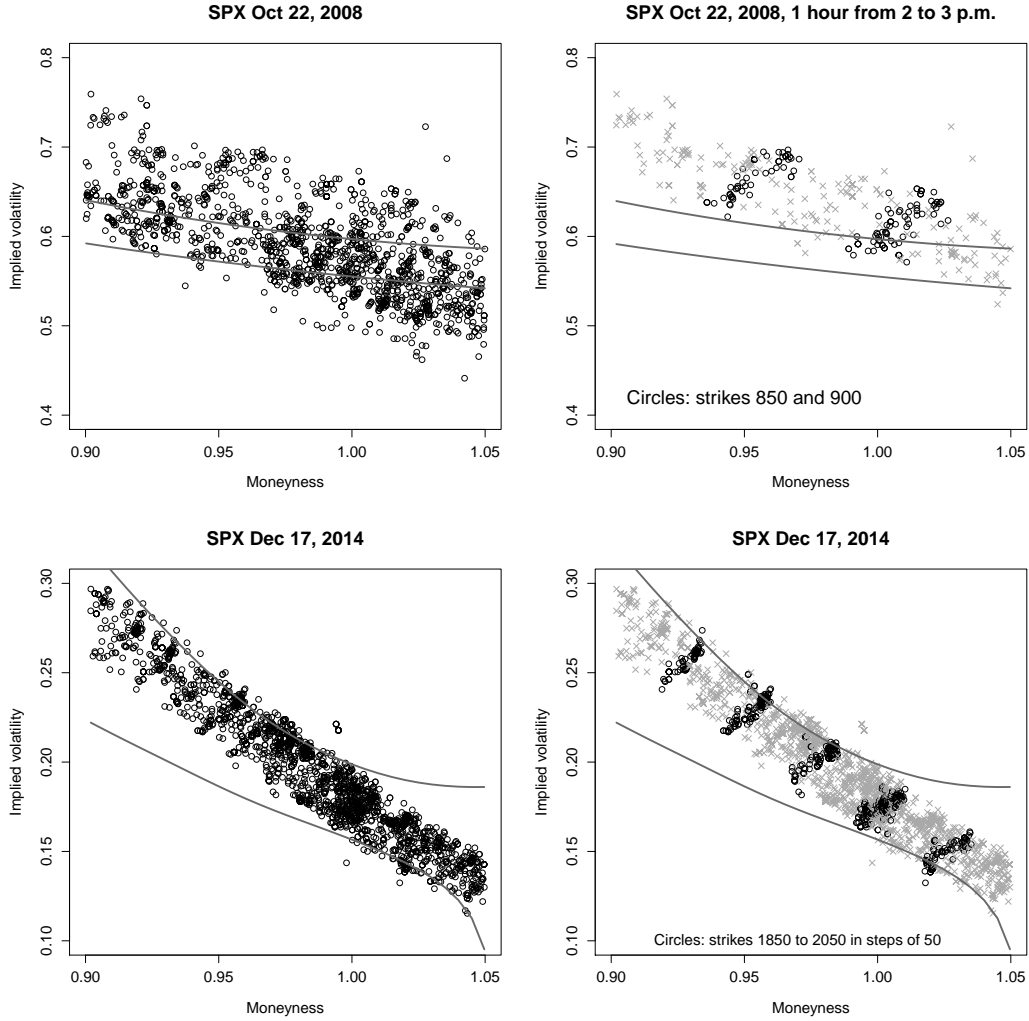


Figure 6: Illustration of intraday movements of the smile pattern

Description: The upper graphs show the smile profiles and stochastic dominance bounds for SPX options on October 22, 2008, which is the first sample day affected by the bankruptcy of Lehman Brothers, and the lower graphs show the smile on December 17, 2014, which is the last day of our sample period. The right upper graph picks out one hour on Oct. 22, 2008, and highlights trades with two particular strike prices. The right lower graph highlights trades on Dec. 17, 2014, with strike prices from 1,850 to 2,050 in steps of 50.

Interpretation: For a given strike, implied volatilities are upward sloping with respect to moneyness. This indicates that the smile shifts upwards when the index falls, and vice versa. Adjusting the bounds to the intraday volatility level would further reduce bound violations.

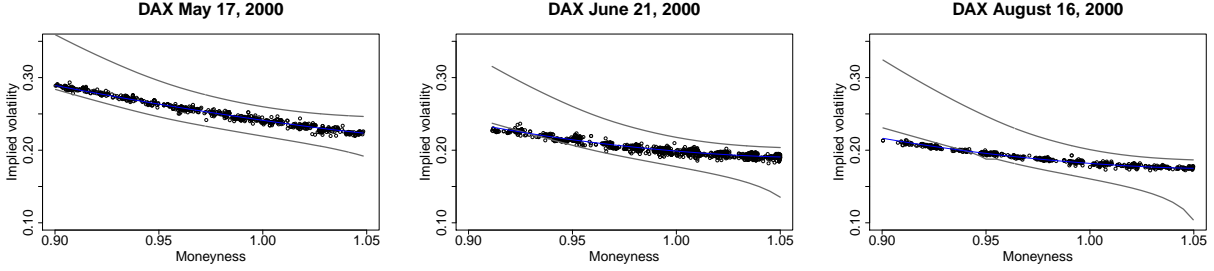


Figure 7: Flattening smile of DAX options in the second half of 2000

Description: The graphs show three months with an almost flat smile profile.

Interpretation: Unlike the typical situation in most other months, the risk-neutral index return distribution is almost symmetrical. As a consequence, the lower bound is slightly violated at low moneyness levels.

on these two differences of our analysis compared to CJP.

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.¹⁷ 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, e.g., Ait-Sahalia and Lo (1998), Fan and Mancini (2009), van Binsbergen et al. (2012), Chen and Xu (2014)). CJP, however, attempt to determine the *interest rate* based on put-call parity.¹⁸ We argue in favor of an implied index level rather than an implied interest rate because measurement error in the index level is much more likely to occur (owing to timing mismatches and imprecise dividend estimates) than measurement error in the appropriate interest rate. In addition, the impact of measurement errors in the index level on estimated implied volatilities is much larger than the impact of errors in the interest rate.

¹⁷CJP 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).

¹⁸For 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 OptionMetrics Database, the authors “cannot arrive at a consistently positive interest rate implied by option prices [...] and use T-bill rates instead” (p. 1274). In a more recent paper studying SPX options from 1986 to 2012, Constantinides et al. (2013) write: “Since we believe that put-call parity holds reasonably well in this deep and liquid European options market, we use the put-call parity-implied interest rate as our interest rate in the remainder of the paper and for further filters” (Appendix B, p. 253).

Fig. 8 illustrates our approach for the last day of the sample period of CJP, which is May 17, 2006. The scatterplots are similar to Fig. 1 for transaction data. The left graph shows the smile pattern based on the closing index level (1,270.32), the middle graph shows implied volatilities provided by OptionMetrics,¹⁹ and the right graph shows the smile for the underlying index level that is consistent with put-call-parity (1,264.10). The differences are large, especially when only call options are included, as in CJP.²⁰ The situation is similar on the other sample days: when the index level is determined such that put-call parity holds for ATM options, the implied volatilities of put and call options coincide over the full range of the smile profile.²¹

With our approach to put-call parity, we obtain modified versions of Figures 3 and 4 in CJP, which are shown in the upper two graphs of Fig. 9. 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. 9) with Figure 4 in CJP reveals that the irregularities and arbitrage violations have disappeared.²²

Conditional volatility Following the bounds analysis in CJP, we adjust the bounds to the implied volatility level so that the test is on the *shape* of the skew instead of its *level*:

¹⁹The separation of put and call options in the middle graph of Fig. 8 is noteworthy because OptionMetrics actually assumes that put-call parity holds (see OptionMetrics Ivy DB File and Data Reference Manual Version 2.5, Rev. 5/5/2005, p. 28: “For dividend-paying indices, ... put-call parity relationship is assumed, and the implied index dividend is calculated ...”). However, OptionMetrics uses two simplifying assumptions: 1) compound interest is linearized; 2) the dividend yield is assumed to be constant over the whole range of option maturities available. For the 1-month options considered here, assumption 2) introduces a non-negligible error if the expected dividends for the next month do not correspond to the average expected dividend yield up to the longest option maturity. Whenever the expected dividends over the next month are above (below) average, the implied volatilities of puts will be higher (lower) than those of calls. This bias could easily be avoided by applying put-call parity to each option maturity separately and, in this way, allowing for a time-changing dividend yield.

²⁰The distorted patterns for calls in the left and middle panels of Fig. 8 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 Fig. 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 Fig. 8.

²¹More 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.

²²In this respect, our study is similar to Battalio and Schultz (2006) who find that most of the apparent violations of put-call parity in Internet stocks in the 1999 to 2000 period disappear when carefully analyzing high-quality option data.

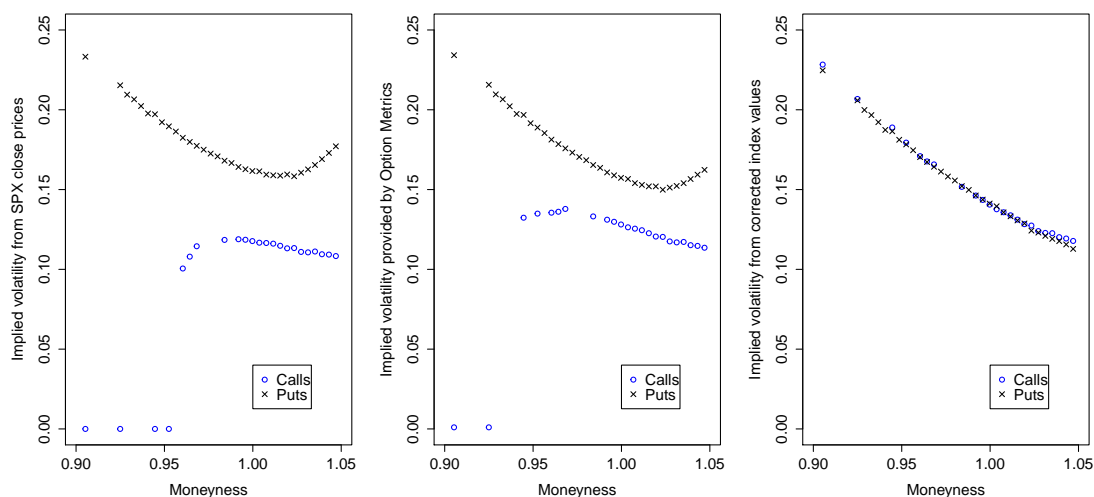


Figure 8: Smile of S&P500 options on May 17, 2006 (time to maturity: 30 days)

Description: Option values are midpoints of bid and ask quotes. Left panel: implied volatilities based on the closing index level of 1,270.32. Middle panel: implied volatilities provided by OptionMetrics (Option Price Files). Right panel: implied volatilities based on an adjusted underlying index value of 1,264.10; the adjustment of -6.22 corresponds to -0.49% of the closing index level.

Interpretation: An index adjustment is necessary to account for dividends and a potential timing mismatch. With this adjustment, the smile profiles of put and call options coincide, which is consistent with put-call parity.

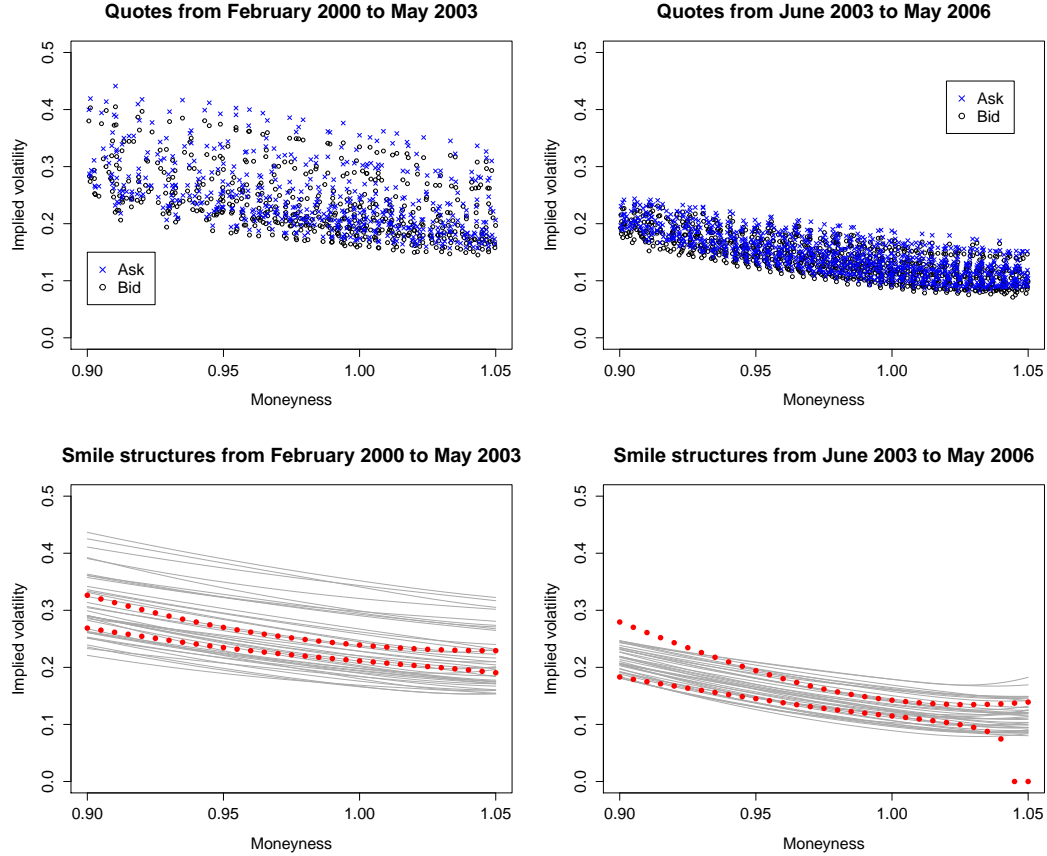


Figure 9: Replication of Figures 3 and 4 in CJP

Description: 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 corresponding 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.

Interpretation: With our approach, negative time values disappear and the smile profiles become very regular. It is crucial to consider conditional volatility with daily updating.

“Since the bounds are adjusted by the implied volatility, irrespective of whether this volatility is rational or not, we can draw inferences about the shape of the skew but not about the general level of option prices” (CJP, 1266). However, while CJP adjust conditional volatility on a monthly basis in another part of the paper, in this part on SD bounds, conditional volatility is set equal to the *average* implied volatility over subperiods of up to three years. Within these subperiods, the conditional volatility and the SD bounds are assumed to be constant.

This assumption induces many apparent bound valuations because volatility varies considerably over such an extended period, which can be seen in the two lower graphs of Fig. 9. The lines show the estimated smile regressions for each month in the sample period (left graph: 2000-2003, right graph: 2003-2006).²³ The dotted (red) lines show bounds computed as in CJP (based on average volatility). It is apparent from the almost parallel skew profiles that the volatility level shifted substantially during the subperiods, particularly from 2000 to 2003. Because the bounds used in CJP are not adjusted accordingly, many upper and lower deviations are observed. CJP state that “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.) However, the higher incidence of apparent violations over 2000 to 2003 is mainly because volatility varied more strongly over this period so that the smile moved out of the average bounds range (see, e.g., the high implied volatilities after the September 11 attacks).

Results Combining our approach to put-call parity with our conditional volatility estimate, we obtain the smile profiles shown in Fig. 10 and 11 for the three-year period from 2003 to 2006 (last subperiod in CJP). The implied volatilities are based on midpoints of bid and ask quotes, and the graphs include call and put options. We also include the smile regression line according to equation (8). As in Section 3, we use the conservative assumption of zero transaction costs and assume a risk premium of 6%. Again, the unconditional distribution is the smoothed historical distribution of S&P 500 returns for 1972 to 2006. As before in our

²³The 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.

transaction analysis, we choose conditional volatility such that the ATM implied volatility lies in the middle of the bounds range.²⁴ This estimate guarantees that we reproduce the general level of option prices, which is in the spirit of minimizing the incidence of violations. Actual violations will be more frequent if options are generally too expensive, which is an issue we do not analyze.

We find that almost all smile profiles fit perfectly into the bounds range. In fact, Fig. 10 and 11 reveal that not a single violation is observed in 36 consecutive months from 2003 to 2006 applying our approach on the CJP data. In all, from January 1997 to May 2006 (113 months), we find violations of SD bounds in 5 months affected by severe market stress (August, September, October, December 1998 and September 2001; see similarly the first two events in Fig. 5 for transaction data).

Bound violations in Constantinides et al. (2011) In a subsequent study, Constantinides et al. (2011) (CCJP) address the concern that the reported violations in CJP do not account for potential errors in the estimation of the SD bounds.²⁵ The authors examine the significance of bound violations for *American*-type S&P 500 *futures* options from 1983 to 2006. Conditional volatility is updated on a daily basis. In one model, it is adjusted to the implied volatility level so that the general level of option prices is considered. Even in this setting, bound violations are still frequent; in particular, 30.5% of all call bid quotes in the moneyness range 1.01 to 1.03 and 15.2% of call bid quotes with moneyness between 0.99 and 1.01 violate the upper SD bound (see CCJP, Table II, p. 1419).

This result appears to contradict our previous findings for SPX options. An obvious question is whether the difference can be explained by the fact that the options are not exactly the same. We argue that this is not the case. It is true that the option studied in CCJP differs in two characteristics from the SPX option: (1) the underlying asset (S&P 500 futures contract vs. S&P 500 index) and (2) the exercise type (American vs. European). However, the first difference should be irrelevant for the question at hand because the futures value is tightly linked to the index level. Moreover, the valuation effect of the second difference—the early exercise option—is known to be small for short-term options (see Whaley (1986) and Ra-

²⁴This conditional volatility is closely related but not identical to the ATM implied volatility.

²⁵The authors are the same with the inclusion of M. Czerwonko whose research assistance is acknowledged in CJP.

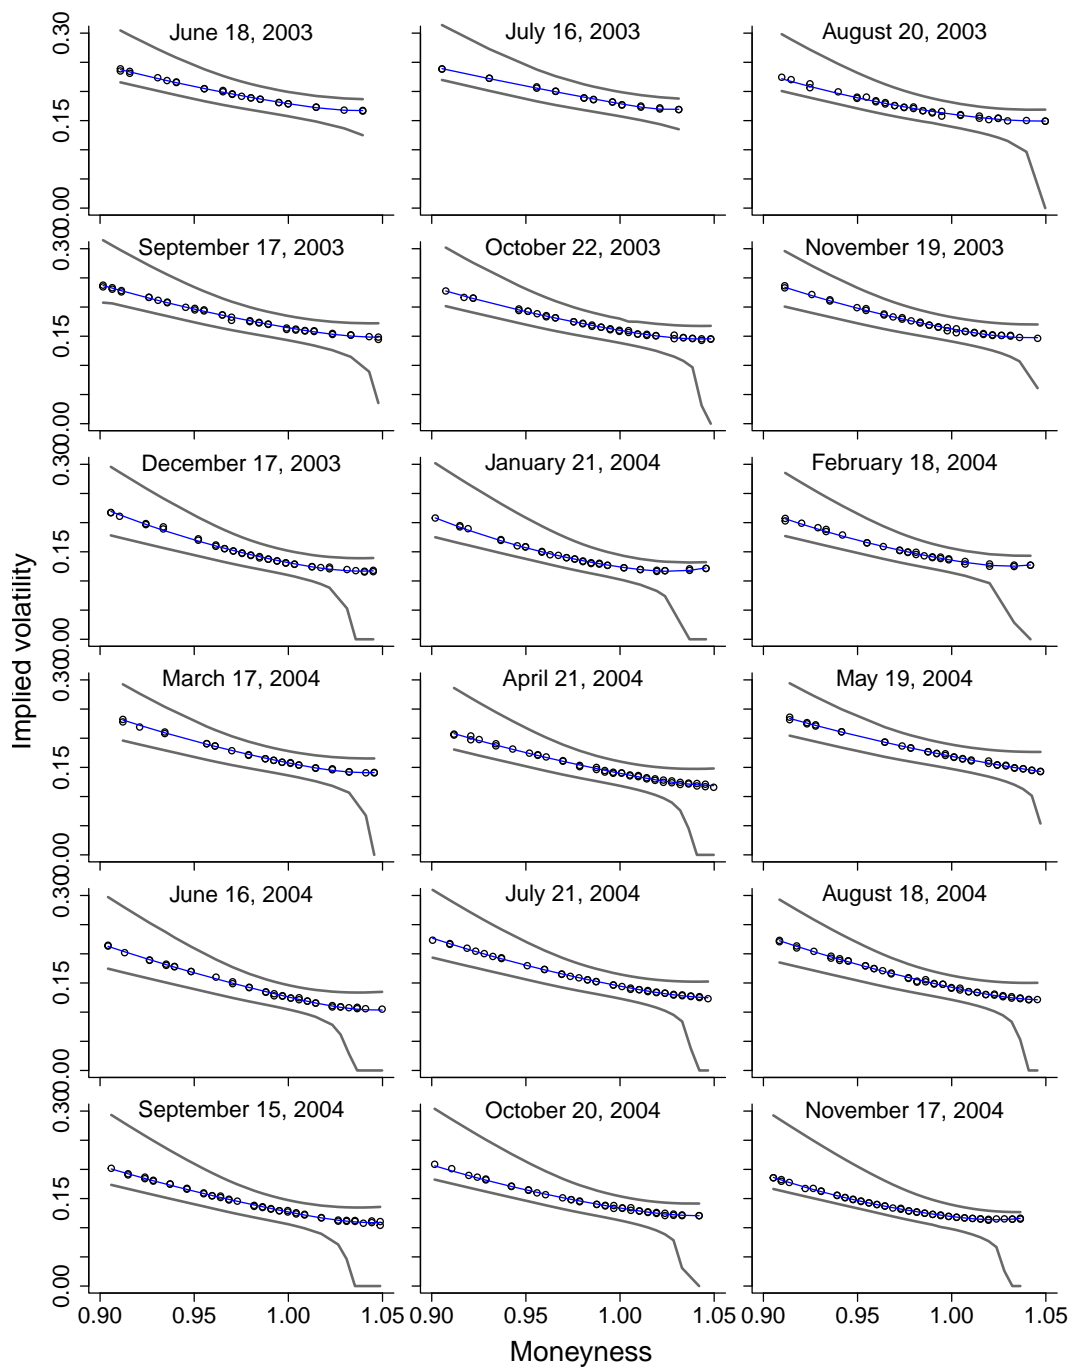


Figure 10: Smile patterns and stochastic dominance bounds of SPX options with a time to maturity of 30 calendar days from June 2003 to November 2004

Description: All puts and calls are included. Implied volatilities are based on mid quotes.

Interpretation: No bound violations occur. SPX option prices are perfectly in line with stochastic dominance bounds when adjusting for (a) the general level of option prices, (b) conditional volatility, and (c) put-call parity.

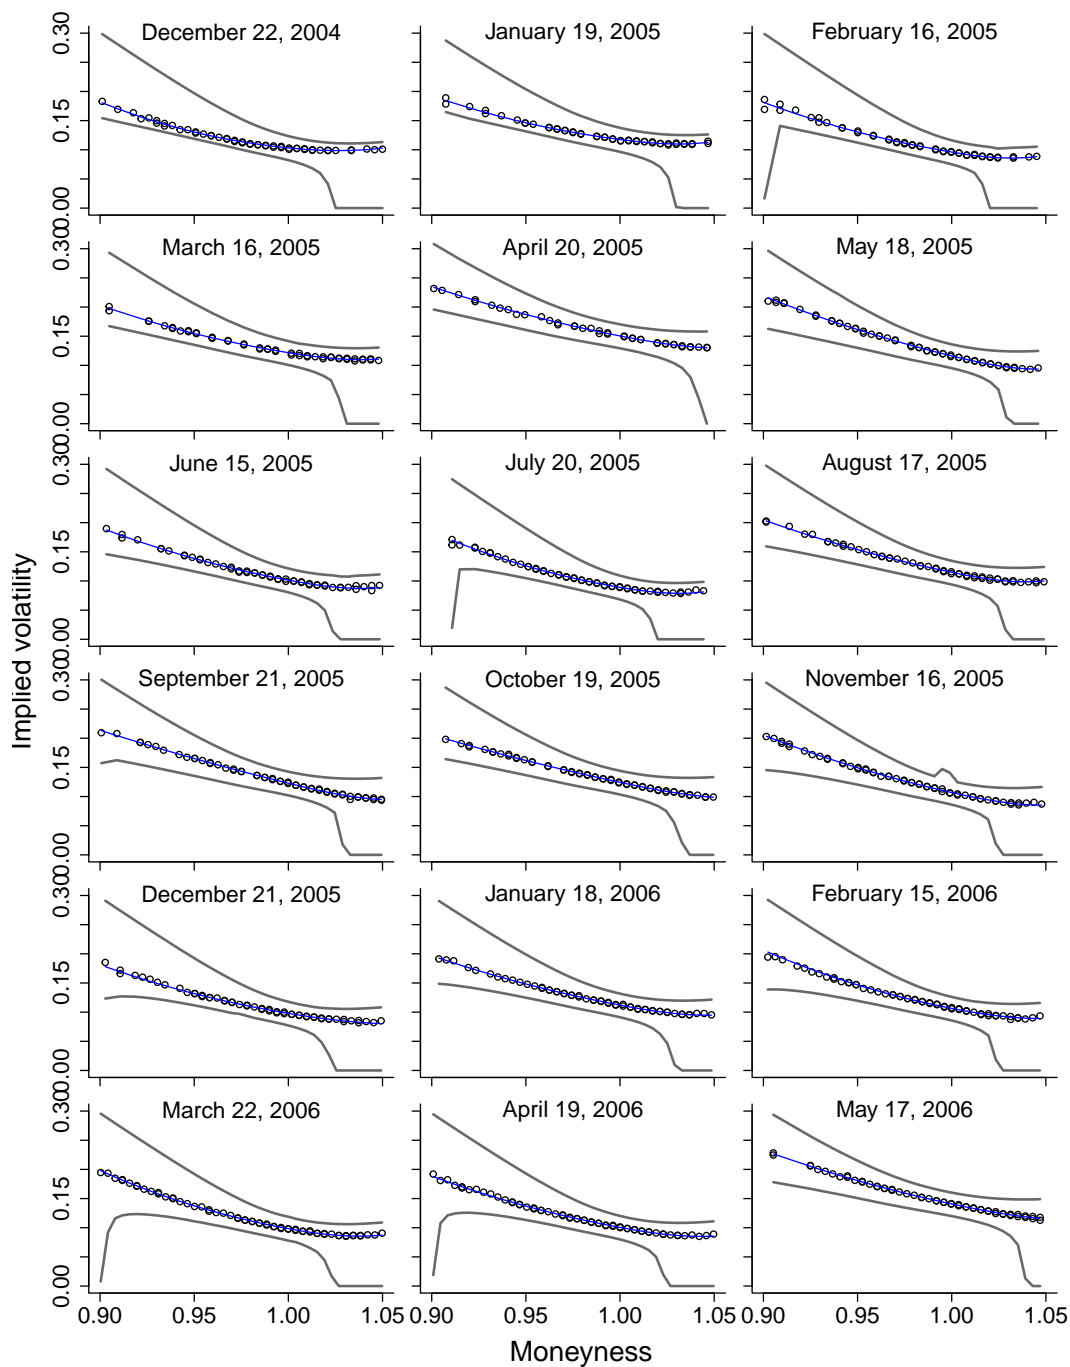


Figure 11: Smile patterns and stochastic dominance bounds of SPX options with a time to maturity of 30 calendar days from December 2004 to May 2006.

Description: All puts and calls are included. Implied volatilities are based on mid quotes.

Interpretation: No bound violations occur. SPX option prices are perfectly in line with stochastic dominance bounds when adjusting for (a) the general level of option prices, (b) conditional volatility, and (c) put-call parity.

maswamy and Sundaresan (1985)). In economic terms, the two options are extremely similar, which will be reflected in similar smile patterns. We confirm this expectation empirically as illustrated in Fig. 12. It shows the overlay of implied volatilities of SPX options (CBOE)²⁶ and S&P 500 futures options (CME)²⁷ for the last four months of the sample period in CJP (settlement data).²⁸ The smile profiles are almost indistinguishable. We verify that this is generally the case at least since 2000. This means that the prices of the American-type futures options fit into the same bounds as SPX option prices. Thus, they also fit into the much wider SD bounds for American options.²⁹ We conclude that our former results carry over to the American-type futures options: no substantial violations of SD bounds are found when considering conditional volatility, the general level of option prices and put-call parity.

5. Conclusion

For S&P 500 options, CJP report widespread violations of the SD 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 on potential mispricing based on a comprehensive database of index options on the S&P 500, EuroStoxx 50 and DAX index. Our main finding is that the three index options are priced almost perfectly in line with SD bounds when (a) considering conditional volatility, (b) adjusting the bounds for the general level of option prices and (c) using put-call parity to estimate the dividend-adjusted underlying index level. Our results strongly suggest that index option prices are consistent with put-call parity, but we do not address the question of whether the general level of option prices is appropriate. Under conditions (a) to (c), more than 96% of option transactions in the period 1995-2014 (DAX)

²⁶Data source: OptionMetrics.

²⁷Data source: “Historical DataMine – EOD (F/O) S&P 500 - Complete History”, provided by Chicago Mercantile Exchange (CME).

²⁸We use the implied volatilities provided by CME. For two months, we verify that these are almost identical to our own calculations based on the Barone-Adesi and Whaley (1987) approximation of American option values with an underlying asset value consistent with put-call parity. While the 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.

²⁹See 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.)

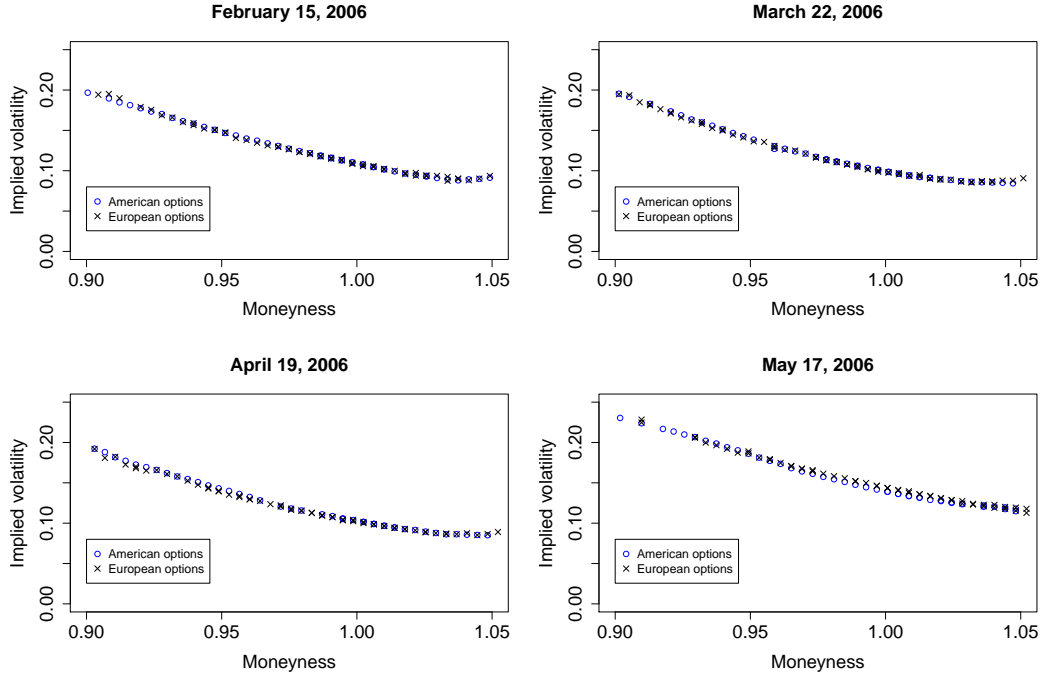


Figure 12: Comparison of smile patterns of American-type S&P 500 futures options with European-type S&P 500 index options (SPX)

Description: The graphs illustrate the smile pattern of both options in the last four months of the sample period of CJP. 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. The options have a time to maturity of 30 calendar days.

Interpretation: The two options are so similar that their smile profiles are almost indistinguishable. If the European index options are priced in line with stochastic dominance bounds, we can safely conclude that this is also the case for the American index future options.

and 2000-2014 (SPX and ESX) lie 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. In all, our results indicate that index option markets might be much more efficient than previous literature suggests.

Appendix: “Radiation” phenomenon

On some days, a number of observations appears to radiate from the center of the smile profile as shown in the left graph of Fig. 13 (SPX option on Aug. 20, 2008). A close inspection of the data reveals that this pattern arises from option prices that remain constant at different index levels. For example, the middle graph of Fig. 13 shows trades in call (blue circles) and put (black crosses) options with strike price 1,265 and an option price of exactly 30. Because the strike price is constant, the variation of moneyness reflects changes in the index level. To offset these changes and remain at a constant option price, implied volatility has to move along the “radiation” trajectories seen in the left and middle graph.

The “radiation” phenomenon is observed since 2006. It is particularly pronounced in 2008 and 2009 and only occurs in options with moneyness very close to 1.00 and a maturity date of the quarterly cycle March, June, September, and December. Our explanation is that these trades are part of a combined trading strategy involving put and call options as well as index futures contracts, which are available only for the quarterly maturity cycle. For example, the option part might comprise a long put and a short call as in a collar trade. In this trade, buyer and seller agree upon a price for the package of the options (e.g., \$1 collar price) without specifying the prices of the individual components. The CBOE reporting system, however, allows only entries for simple put and call options. Therefore, the collar price has to be decomposed. To simplify the entry, an integer value is typically used for at least one of the two price components. For example, a collar price of \$1 might be recorded as \$34 for the embedded long put and \$33 for the embedded short call. When the collar price decreases to \$0.50, the recorded put price might be kept constant at \$34 while the call price is adjusted to \$33.50 or the put price might be reduced to \$33.50 while the call price is left at \$33. In any case, the recorded prices are not informative if their connection is lost.

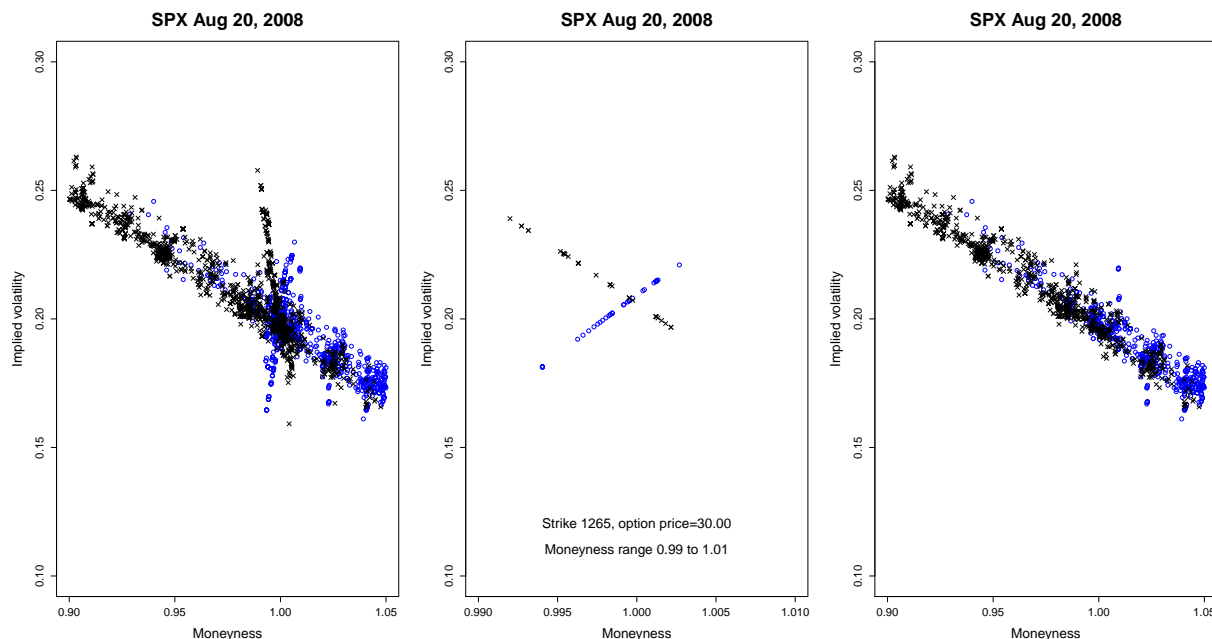


Figure 13: Radiation phenomenon in SPX smile profile

Description: Black crosses: put options, blue circles: call options. The left graph shows an example of a particular pattern of trades radiating from the center of the smile profile on August 20, 2008. The phenomenon occurs with varying intensity since 2006 when the maturity date falls into the quarterly cycle of maturities of index futures. Our explanation is that the pattern is caused by combined trades that have to be decomposed for recording in the CBOE reporting system. The artificial component prices are apparent from constant integer option prices in the middle graph. We apply a simple identification rule to remove (a part of) the corresponding trades. The result is shown in the right graph.

Interpretation: The way of recording combined option strategies in the CBOE reporting system can produce artificial deviations from the regular smile profile. A simple identification rule allows us to remove an important part of these trades.

The CBOE trade files do not allow for identifying combined trades.³⁰ Therefore, we develop a simple identification rule and remove the corresponding observations from our database. Our rule is to remove records which fulfill all of the following four conditions: (a) trade year 2006 or later, (b) maturity date of the quarterly cycle, (c) moneyness between 0.99 and 1.01, and (d) integer value of the option price. The right graph of Fig. 13 shows that on this particular day, our rule effectively removes the artefact. In general, the rule is, of course, imprecise for two reasons. First, many trades that fulfill the four criteria will be “regular” trades (not part of a complex strategy). For our study, the reduced quantity is not a concern because there is an abundance of daily trades in ATM options. Moreover, removing these trades is not likely to bias our results because there is no apparent reason why regular trades that fulfill conditions (a) to (d) might be more or less prone to bound violations than other regular trades. Second, the rule will eliminate only one “leg” of the combined trade, except for cases with an integer collar price. This means that the other “leg” will still artificially increase the number of bound violations. In this sense, our correction rule is conservative. It could be improved in different ways, for example by considering the distance to the smile. We do not follow this direction in order to minimize the risk that real bound violations are removed.

Fig. 14 illustrates our correction with three further examples (left graph: before correction, right graph: after correction). The upper panel shows one of the most pronounced radiation artefacts in our data. Apparently, it is partially removed. The middle and lower panels show examples where the radiation phenomenon is less pronounced but still stronger than on almost all the other days in our sample. We follow from these examples and a visual inspection of all other smile profiles that our correction works reasonably well.

³⁰The CBOE Support Team (Market Data Express) confirmed in writing that no details are available for identifying combined trades. I am not aware of a discussion of this phenomenon in the literature.

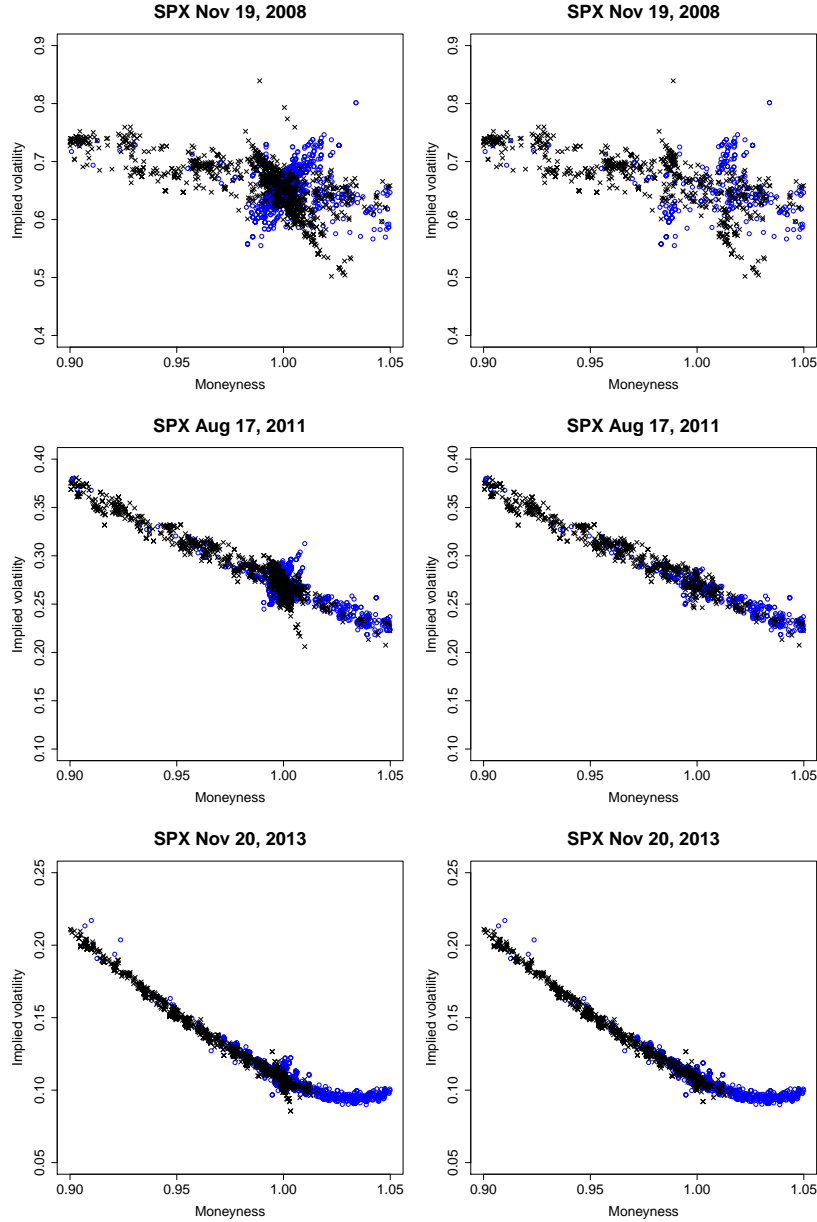


Figure 14: Further illustrations of radiation phenomenon and our correction

Description: Black crosses: put options, blue circles: call options. The left graphs show examples of a particular pattern of trades radiating from the center of the smile profile of SPX options. The phenomenon occurs with varying intensity since 2006 when the maturity date falls into the quarterly cycle of maturities of index futures. Our explanation is that the pattern is caused by combined trades that have to be decomposed for recording in the CBOE reporting system. We apply a simple identification rule to remove (a part of) the corresponding trades. The result is shown in the right graphs.

Interpretation: The way of recording combined option strategies in the CBOE reporting system can produce artificial deviations from the regular smile profile. A simple identification rule allows us to remove an important part of these trades.

References

- Aït-Sahalia, Y., Lo, A., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499–547.
- Aït-Sahalia, Y., Lo, A., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.
- Barone-Adesi, G., Whaley, R., 1987. Efficient analytic approximation of American option values. *Journal of Finance* 42, 301–320.
- Battalio, R., Schultz, P., 2006. Options and the bubble. *Journal of Finance* 61, 2071–2102.
- Beare, B., 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., Johannes, M., 2009. Understanding index option returns. *Review of Financial Studies* 11, 4493–4529.
- Carr, P., Wu, L., 2009. Variance risk premiums. *Review of Financial Studies* 22, 1311–1341.
- Chen, S., Xu, Z., 2014. On implied volatility for options - some reasons to smile and more to correct. *Journal of Econometrics* 179, 1–15.
- Chernov, M., 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 Economics* 56, 407–458.
- Constantinides, G., Czerwonko, M., Jackwerth, J., 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., Perrakis, S., 2008. Option pricing: Real and risk-neutral distributions. In: Birge, J., Linetsky, V. (Eds.), *Handbooks in OR & MS Vol. 15*. pp. 565–591.
- Constantinides, G., Jackwerth, J., Perrakis, S., 2009. Mispricing of S&P 500 index options. *Review of Financial Studies* 22, 1247–1277.
- Constantinides, G., Jackwerth, J., Savov, A., 2013. The puzzle of index option returns. *Review of Asset Pricing Studies* 3, 229–257.
- Constantinides, G., 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., Perrakis, S., 2007. Stochastic dominance bounds on American option prices in markets with frictions. *Review of Finance* 11, 71–115.
- Driessen, J., Maenhout, P., 2013. The world price of jump and volatility risk. *Journal of Banking & Finance* 37, 518–536.

- Fan, J., Mancini, L., 2009. Option pricing with model-guided nonparametric methods. *Journal of the American Statistical Association* 104, 1351–1372.
- Hafner, R., Wallmeier, M., 2000. The dynamics of DAX implied volatilities. Working Paper University of Augsburg.
- Hafner, R., 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., 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., Schneider, P., 2013. The skew risk premium in the equity index market. *Review of Financial Studies* 26, 2174–2203.
- Ramaswamy, K., Sundaresan, S., 1985. The valuation of options on futures contracts. *Journal of Finance* 11, 1319–1340.
- Rosenberg, J., 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., Saretto, A., 2009. Option strategies: Good deals and margin calls. *Journal of Financial Markets* 12, 391–417.
- Santa-Clara, P., Yan, S., 2010. Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *Review of Economics & Statistics* 92, 435–451.
- van Binsbergen, J., Brandt, M., Koijen, R., 2012. On the timing and pricing of dividends. *American Economic Review* 102, 1596–1618.
- 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* 42 (3), 49–59.