Review of Finance, 2016, 1–33
doi: 10.1093/rof/rfw061

OXFORD

What Drives Index Options Exposures?*

Timothy Johnson[1], Mo Liang[2], and Yun Liu[1]

[1]University of Illinois at Urbana-Champaign and [2]School of Finance, Renmin University of China

Abstract

This paper documents the history of aggregate positions in US index options and investigates the driving factors behind use of this class of derivatives. We construct several measures of the magnitude of the market and characterize their level, trend, and covariates. Measured in terms of volatility exposure, the market is economically small, but it embeds a significant latent exposure to large price changes. Out-of-the-money puts are the dominant component of open positions. Variation in options use is well described by a stochastic trend driven by equity market activity and a significant negative response to increases in risk. Using a rich collection of uncertainty proxies, we distinguish distinct responses to exogenous macroeconomic risk, risk aversion, differences of opinion, and disaster risk. The results are consistent with the view that the primary function of index options is the transfer of unspanned crash risk.

JEL classification: G12, N22

Keywords: Index options, Quantities, Derivatives risk

1. Introduction

Options on the market portfolio play a key role in capturing investor perception of systematic risk. As such, these instruments are the object of extensive study by both academics and practitioners. An enormous literature is concerned with modeling the prices and returns of these options and analyzing their implications for aggregate risk, preferences, and beliefs. Yet, perhaps surprisingly, almost no literature has investigated basic facts about quantities in this market. How much do investors actually use these contracts, and why?

This study addresses these questions by documenting historical patterns in the aggregate demand for (and supply of) S&P 500-based options. We establish primary properties about the level, trend, and the drivers of variation in index options exposure.

* We are grateful to Bernard Dumas and two anonymous referees for many constructive suggestions. We also thank Dimitri Muravyev for many useful comments and observations. This research is supported by Renmin University of China: the special developing and guiding fund for building world-class universities (disciplines), project No. 16XNKI005. Correspondence author: Mo Liang, moliang@ruc.edu.cn.

Although theories of option demand have been proposed for over 40 years, little is known about why index options are traded in practice.[1] In classical models, derivatives are redundant assets whose quantities are essentially indeterminate. Positive net demand arises with heterogeneous agents in the presence of frictions or market incompleteness. Examining demand patterns empirically may thus be informative about the relevant type and degree of heterogeneity, frictions, and incompleteness. Since index options effectively reference, the "market portfolio", these elements are potentially primitive features of the economy that may be fundamental in understanding asset pricing.

To guide our empirical investigation, Table I summarizes three potentially distinguishable families of theories of option demand. We identify these by the type of risk transfer that motivates the demand (first row) and the nature of the heterogeneity across agents that drives trade (second row). The hypotheses are not meant to be exhaustive. However, they represent distinct economic notions that may suggest differing predictions.

Column (1) describes theories in which options are used to transfer risk associated with the level of aggregate wealth. As discussed by Black (1975) and emphasized by practitioners, options may present a cost or efficiency advantage for such transfer when some agents face frictions like costly borrowing, short-sale constraints, or transactions costs in the underlying market. (This type of demand can exist even under the complete-market assumptions that support classical option pricing theories, since the theories only require that some agents face frictionless markets.) If these factors are important, one would expect variation in options usage to correlate with variation in measures of transaction costs. Moreover, since these costs are highest for individuals, options usage might be expected to vary with the level of retail participation.

Since at least Ross (1976) and Hakansson (1978) economists have appreciated that perfect markets fully spanned by the ability to trade in an underlying index asset is an abstraction, and that options may therefore play a real role in completing markets. Much work has focused on the case of jump-diffusion processes, emphasizing the role of options as hedges for downside jumps (i.e., crashes). But even a simple two-period economy (as in Grossman and Zhou, 1996) there will be unspanned risk with more than two market outcomes.

The second Column (2) describes theories of this type. There is a very broad range of models, both partial and general equilibrium, encompassing investors who differ on a number of different dimensions.[2] Most work considers heterogeneity in beliefs or preferences. But demand can also be driven by differences in such factors as background risk or portfolio constraints, which may be indistinguishable from preference differences.[3]

Column (3) describes models also concerned with the transfer of unspanned risk, but not of the market return itself, but of the return's second (or higher) moments. Models of unspanned stochastic volatility have been used in practice since the 1980s. More recently,

1 In contrast, a substantial empirical literature starting with Easley, O'hara, and Srinivas (1998) examines quantities of individual firm options, focusing on the role of private information. Recent contributions include Lakonishok *et al.* (2007) and Roll, Schwartz, and Subrahmanyam (2010). One would expect private information to play little role in index options markets.

2 Partial equilibrium models go back to Leland (1980) and Brennan and Solanki (1981). See also Carr, Jin, and Madan (2001). Recent general equilibrium contributions include Buraschi and Jiltsov (2006) and Bhamra and Uppal (2009).

3 See Franke, Stapleton, and Subrahmanyam (1998) and Vanden (2004).

Table I. Option demand hypotheses

	(1)	(2)	(3)
Risk:	Market risk (spanned)	Unspanned market risk (e.g., jump risk)	Higher moment risk (e.g., volatility risk)
Heterogeneity:	Trading costs; portfolio constraints	Beliefs; preferences; background risk; portfolio constraints	
Potential covariates:	Commissions; spreads; information technology; retail participation	Risk measures; volume; disagreement; wealth shares	
Likely preferred types of option:	High elasticity; at-the-money	low strike price	high volatility sensitivity

much work in asset pricing has explored the implications of nonstandard preference theories, such as those of Epstein and Zin (1989), in which agents' utility may directly depend on higher-order risks. As with Column (2), this type of unspanned risk could be combined with numerous types of heterogeneity to generate net option demand.[4]

For either type of unspanned risk transfer (Columns (2) and (3)), a natural hypothesis is that the quantity of options should fluctuate with the degree of risk.[5] Our empirical work will employ a number of recently developed proxies for different dimensions of economic risk. Similarly, measures that reflect changes in the degree of heterogeneity of the investor population may explain option usage. We have some candidate proxies for disagreement, sentiment (or risk-aversion). Moreover, with any type of heterogeneity, the degree of trade will vary with the relative wealths of the subsets of agents who are trading. We examine time variation in the composition of the US investor population (households versus different institutions) as a potential driver of options trends.

Finally, the last row of the table is motivated by the observation that different reasons for option usage will be reflected in different types of options being preferred by investors. For example, if cost considerations in gaining market exposure are an important factor, then investors may prefer options with maximum elasticity (bang for the buck) either on the long or short side. Alternatively, they may prefer options whose liquidity is highest, which are typically those at-the-money. In contrast, if investors are concerned with unspanned crash risk, demand would be expected to be concentrated in out-of-the-money puts. Whereas, if volatility risk transfer is an important motivation for trade, one might expect that most activity would be in long-dated options, whose volatility sensitivity (vega) is highest.

4 Recent models of heterogeneity in perception of (and/or aversion to) stochastic volatility or time-varying disaster risk include Bates (2008) and Chen, Joslin, and Tran (2012).

5 Since Rothschild and Stiglitz (1971), a large body of work has explored conditions under which demand for insurance increases with risk. See also Eeckhoudt and Gollier (2000). With both supply and demand effects, the sign of the effect on quantity is indeterminate.

To be clear, the goal of our study is to let the data speak to us on high-level issues, as suggested by the table. We do not aim to structurally test particular models, few of which have been developed to the point where they might reasonably be brought to the data. (And, as discussed above, many distinct models may, in fact, have similar empirical implications.) Instead, we view our contribution as laying the groundwork for future theoretical and empirical work on this topic.

For econometric clarity, our work proceeds by first documenting the scale of the index options market, then modeling its stochastic trend, and finally explaining variation from that trend. Our findings may be summarized as follows.

The overall size of the options market is small in terms of the aggregate value of contracts and their gross exposure to volatility risk. On the other hand, the market is large in the sense that market jumps have the potential to induce substantial exposures and economically large rehedging demands. Our findings about market scale appear to be robust to consideration of over-the-counter (OTC) derivatives, non-US contracts, and listed volatility derivatives.

Examination of options usage by type reveals that almost half of open interest is concentrated in out-of-the-money puts. This is strong evidence in favor of the unspanned crash-risk story for options usage. In contrast, volatility risk transfer would appear less important, as long-dated, at-the-money options (i.e., those with high volatility exposure) are a small component of open interest.

Over time, the scale of positions in index options is nonstationary. The outstanding number of contracts was quiet in the mid-1990s, rose substantially prior to 2007 and has since stagnated. The nonstationarity is well explained by the stochastic trend in stock market turnover, affirming the intuition above that trading volume impounds both the degree of investor heterogeneity and the evolution of financial transactions technology. On the other hand, we do not find that changing composition of the investor population can account for the evolution of options usage. Nor do the level of interest rates and information technology costs account for the trend in options positions.

The most significant determinant of fluctuations in detrended options positions is a negative response to risk, which is not driven by the risk-aversion component in measures of investor sentiment, nor by differences of opinion. Increases in belief dispersion and deterioration of investor sentiment induce significant positive responses of aggregate open interest. There is also evidence of a positive quantity response to some measures of tail risk.

The primary finding of a strong negative response to risk is suggestive of a story in which one set of agents (e.g., dealers) have tolerance for unspanned crash risk that is higher than that of another class (e.g., investors) in normal times, but which diminishes rapidly in the face of rising risk, possibly due to binding position constraints or wealth effects.

In all, we succeed to a large degree in answering the question in the paper's title: our empirical descriptions of aggregate options measures achieve high explanatory power. Fluctuations in equity market activity and different dimensions of risk drive most of the variation in options quantities.

The outline of the paper is as follows. The next section reviews the products with which the study will be concerned, addresses some conceptual issues in the measurement of options positions, and describes our aggregate exposure variables. Section 3 presents our basic findings for levels and economic magnitude of options quantities. Section 4 addresses the trend in the data via a cointegration analysis with potential nonstationary factors. Section 5 then investigates determinants of deviations from the fitted trend. A final section summarizes the findings and highlights issues for future research that our study raises.

2. Measuring Index Options Quantities

This section describes the options series that we construct and discuss some conceptual issues in aggregation. We also describe the most important sources of exposure that our measures do not capture.

Options on stock indexes have traded on exchanges in the USA since the early 1980s. Our goal is to study historical trends in this activity, and therefore we focus on a single underlying index—the S&P 500—to ensure consistency over time. We are thus omitting options on other indexes that have, at one time or another, enjoyed a degree of popularity.[6] Even within the set of products referencing the S&P 500, there has been a lot of variation in, and competition among, the products available for trade and their relative market shares. Our data consist of three classes of products, whose differing trade and settlement mechanism makes them not literally fungible, but which are essentially equivalent economically, allowing us to aggregate our measures across them. The products we use are the following:

- CBOE S&P 500 Index (SPX) Options.
- CME S&P 500 Futures Options.
- Options on SPDR exchange-traded funds (ETFs) (SPY).

Appendix A provides institutional details on each of these contract types and on the data we have for each.

As described in Section 1, our interest is in quantifying the extent of demand for index options as a distinct product class. At this point, it is worth considering conceptually what it is that we would ideally like to measure. There are two important issues.

First, options obviously are not interchangeable. So how should positions of different types be aggregated? In particular, call and put positions represent opposite exposures to index returns. Our view is that the economically interesting feature of options positions is the nonlinearity that is common to both products however. We are not attempting to measure directional exposure. Therefore, we will treat puts and calls interchangeably, and sum their gross positions. As a robustness check, and to gain further insight into demand characteristics, we will also report results for separate buckets of options classes (calls/puts; at/in/out of the money; short/long expiration horizon).

Obviously, our measures will be of gross positions, since net options positions are always identically zero. This brings us to the second important issue in aggregation: hedging. Our data are comprehensive over time and across products, but they contain no information on net holdings of particular investors or investor class.[7] We have no way of

6 Today the most liquid US index products other than those referencing the S&P 500 are those on the NASDAQ 100. Other index products that have been important in the past include those on the Value Line index, the AMEX Major Market index, the Dow-Jones 30, the Russell 5000, and the S&P 100.

7 There are data sets of option trade by customer type covering some of our index contracts during some periods. These can shed direct light on questions of why specific investors trade specific types of contracts. See Garleanu, Heje Pedersen, and Poteshman (2009); Chen, Joslin, and Ni (2014); and Lemmon and Ni (2014) for recent contributions in this direction.

quantifying the extent to which participants may cancel the economic exposures of their options positions. For example, if a dealer buys a put and sells a call of the exact same expiration and maturity, and then goes long the underlying asset, he effectively has no position in either option. Similar remarks apply to options spread trades—being long and short similar but not identical options.[8]

Since clearly many options markets participants hedge to some degree, this means that our exposure measures may overstate actual economic exposures. However, this caveat itself requires a caveat. Not all kinds of hedging distort the interpretation of gross exposures. In particular, delta hedging—cancelling the directional risk of an options position—is not problematic for us as it leaves intact the essential "option-ness", or convexity, of the position. For example, if a market participant buys a zero-delta straddle, we do want to measure that as an open position.

Moreover, hedging by one party to an options trade does not by itself nullify the economic impact of the position. It may merely transfer it to whomever the hedging trade is done with. If a dealer buys a call (and sells the underlying) to hedge a short put position, certainly he is effectively out of the position. But unless the call and put position are both with the same counterparty, then someone has assumed the convexity of the original position by selling the put. And, again, we would want to measure that. In the extreme case, it is possible that all options market participants hedge all their convexity (i.e., with each other). But presumably this class of derivatives would not exist if there was not some underlying demand for the nonlinearity of the exposure they afford.

We can now describe the exact construction of the main measures that we will study. Our primary benchmark is simply gross open interest, which provides a non-monetary measure of the total amount of options outstanding. This is calculated as

$$\mathrm{OI}(t) = \sum_i \mathrm{CM}_i \, \mathrm{OI}_i(t), \tag{1}$$

where CM_i (contract multiplier) denotes the number of equivalent index units referenced by a single option of i's type, $\mathrm{OI}_i(t)$ (open interest) is the number of option of i's type outstanding on date t, and the sum runs over all option types. This measure is in index equivalent units, or "shares" of the S&P 500. Multiplying $\mathrm{OI}(t)$ by the level of the index on date t gives a measure of the total monetary value of shares underlying all outstanding options. The latter is effectively the definition of the "notional value" of these derivatives. A second basic measure of outstanding options is simply their aggregate market value, or:

$$\mathrm{MV}(t) = \sum_i \mathrm{CM}_i \, \mathrm{OI}_i(t) \, P_i(t). \tag{2}$$

where $P_i(t)$ is the closing price of the ith option-type. This measure is in units of dollars. It quantifies the gross investment of agents in this product. If all options were somehow voided (e.g., through a clearing-house default) or expired worthless this is the amount that would be lost. Since option values all scale with the level of the underlying, this measure,

8 Our data do, however, take into account literally off-setting positions. All the options in our study are centrally cleared. Reported open interest is netted by the clearing house at the account level. If a market maker buys an option from one customer and sells the exact same option to another, then, the recorded position is the same as if the two customers had traded with each other. This is typically not the case in OTC derivatives markets. We discuss OTC markets further below.

when compared with the monetary value of the index itself, yields a normalized average contract value.

To measure the gross risk characteristics of options positions, we next introduce two weighted average measures: aggregate vega and aggregate gamma of outstanding contracts.

A standard quantification of the volatility exposure of an option, "vega" is market parlance for the change in the option's value—according to a benchmark model—when the underlying annualized volatility is perturbed by 0.01 (one "vol"). If $\nu_i(t)$ represents the partial derivative of the ith option-type (underlying, strike, expiration) on date t, then our aggregate measure is:

$$\mathrm{AV}(t) = 0.01 \sum_i \mathrm{CM}_i \, \mathrm{OI}_i(t) \, \nu_i(t). \tag{3}$$

We follow market practice here, using benchmark Black and Scholes (1973) or binomial models for European and American options, respectively.[9] By definition, $\nu_i(t)$ measures the change in the option's value in response to a change in implied volatility. We note that this is not necessarily how much the option's value will change in response to a change in actual return volatility. However, AV does give a direct quantification of the gross amount of volatility risk transfer achieved in the index options market. Like MV it is in dollars.

Similarly, "gamma" is market parlance for the model-based sensitivity of the equivalent share exposure (the "delta") of an options position to a \$1 move in the underlying index. We convert this to an equivalent change per 1% move in the underlying via multiplying the partial derivative, $\Gamma_i(t)$, by $.01S(t)$ and then summing:

$$\mathrm{AG}(t) = 0.01 \sum_i \mathrm{CM}_i \, \mathrm{OI}_i(t) \, \Gamma_i(t) \, S(t). \tag{4}$$

This measure gives the total index equivalent units that would have to be traded (in absolute value) if all options positions were re-hedged to market neutrality in response to a 1% index change. It can be expressed as a turnover ratio by dividing by the total index equivalent units outstanding, defined to be simply the total monetary value of the index stocks (weighted by their index share) divided by the level of the index. As described in the next section, AG, can also be related to the total monetary exposure of options positions to the risk of a jump in the index.

Before turning to the data, it is worthwhile to consider briefly what can be said about stock index options not captured by our measures.

We have already alluded to various other US listed index options that have not been consistently successful. However, it is reasonable to wonder how important they might be collectively. The same question could be asked about index options products listed on non-US exchanges, referencing the Nikkei 225 or the FTSE 100 for example. Finally, there are active non-listed OTC index options markets that are not in our measures.

For aggregate information on all of the above products, we will compare our numbers to those reported in a semi-annual survey undertaken by the Bank for International Settlement (BIS) since 1996. The BIS numbers for exchange-traded options are computed from sources similar to ours, whereas the OTC numbers are estimated based on reports to

9 The computation is numerically challenging for large numbers of American options. Details of an efficient numerical procedure are given in the Appendix.

the BIS from member central banks.[10] There are issues in comparing OTC numbers to exchange-traded ones. In particular, OTC positions are rarely netted, and are typically cleared bilaterally. This means that amounts outstanding are greatly overstated compared with the equivalent numbers from exchanges. Moreover, the BIS does not correct for double counting of options positions open between reporting banks. This contributes further to the inflation of OTC statistics.[11] For these reasons, it would be inappropriate to literally combine our measures with the BIS numbers.

Finally, to the extent that our study is concerned with volatility risk transfer, a question arises as to what extent participants have utilized direct volatility derivatives to accomplish this. Investment banks have actively quoted market in variance swaps and other similar products for at least the last 20 years. Unfortunately, we do not know of any sources of data on quantities of OTC volatility derivatives. However, recent years have seen a rise in popularity of listed volatility derivatives, especially futures tied the CBOE implied volatility index, VIX. We have data on these contracts, and we discuss their magnitude in the next section.

3. The History of Index Options Quantities

Figures 1–4 plot the time series of our measures, showing the contributions of the three contracts separately. (The plots are in logarithms.) By all measures, the dominant class has been the CBOE SPX options. Futures options achieved nearly equal market share briefly in the mid-1990s but have since fallen back. The SPY contracts have grown rapidly and now rival the CME products in terms of open interest and gamma, although their monetary value and vega remains negligible.

All the series are clearly nonstationary. There has been a strong upward trend historically. However, there have also been extensive periods of stagnation, including the later 1990s and the period following the financial crisis. We attempt to identify exogenous drivers of these trends in the next section.

In terms of economic magnitudes, the numbers reveal an interesting contrast: the options market is small in monetary value, but large in terms of potential value as measured by the number of shares referenced. Table II reports the series values as of the end of the sample December 31, 2012, as well as on the highest recorded day of AV which was September 22, 2008. For reference, the first and fourth columns give the size of the S&P 500 in (respectively) billions of dollars and billions of equivalent shares outstanding. (The latter is defined as the former divided by the nominal index level.)

From the second column, the total value of all S&P 500 options is less than $100 billion, about the scale of a single mid-sized company. The capitalization of the market itself is two orders of magnitude larger. In terms of risk transfer, the vega numbers in the third column indicate that, again, investors have not insured much exposure. An aggregate vega value of $2 billion means that a doubling of return volatility, for example, from 20% to 40%,

10 The survey originally covered banks in the G10 countries, and currently covers thirteen countries. The coverage within countries depends on the reporting requirements of each central bank.

11 To give a numerical example, suppose Dealer A sells a customer a quantity Q of an option, and buys an off-setting position from Dealer B. Then the BIS survey will see this as open positions of 3Q–2Q reported by A and 1Q by B. In a centrally cleared market the open interest would be recorded as Q.

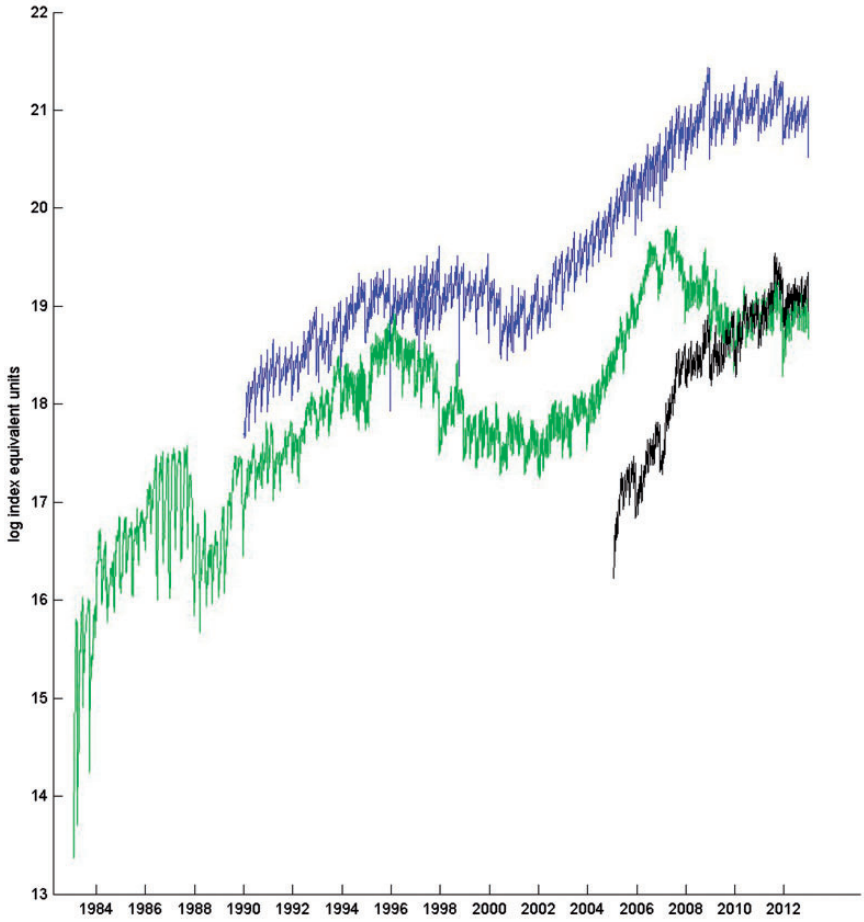

Figure 1. Total open interest.
The top line is the total open interest of all outstanding SPX options listed on the CBOE. The middle line is the open interest of all CME-listed options on S&P 500 futures. The lowest line is the open interest of all options on SPDR ETFs. The units are logarithms of index equivalent units.

would transfer an economically negligible \$40 billion from option writers to option purchasers.

On the other hand, from the fifth column, the numbers of shares referenced by index options are a nonnegligible fraction of the total equity of the market. The small monetary value and vega, then, are indications that most option positions are short-term and/or out-of-the-money. The sixth column indicates that aggregate gamma is high in the sense that a 1% move in the market could induce rehedging demands that correspond to approximately 1% of shares outstanding. This may not seem large until one recognizes that total turnover of the US equity markets averages well less than 1% of shares outstanding on any given day. Thus, options positions are quite large economically in terms of their potential to transfer risk. Although we know it is not the case, as emphasized above, that long options positions and short option positions are held by disjoint sets of agents, even if only a

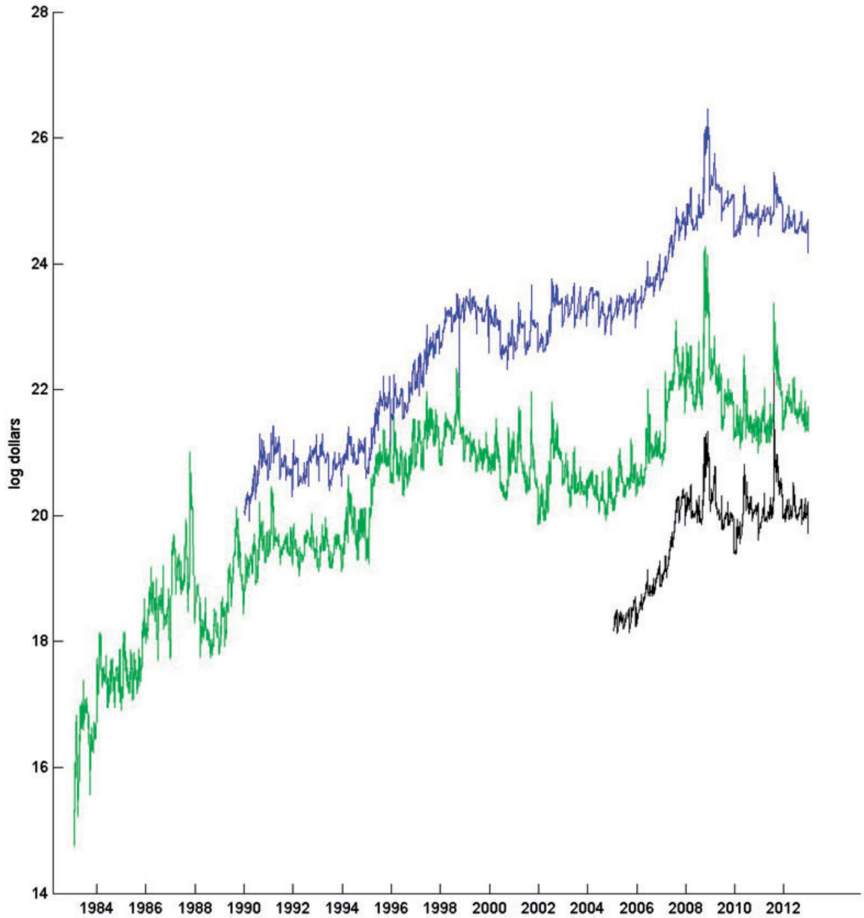

Figure 2. Market value.
The top line is the total market value of all outstanding SPX options listed on the CBOE. The middle line is the market value of all CME-listed options on S&P 500 futures. The lowest line is the market value of all options on SPDR ETFs. The units are logarithms of dollar value.

fraction induce rehedging, the response to, say, a 10% index move could strain market liquidity.[12]

Table III shows the average fraction of open interest represented by twelve types of option: calls and puts, separately classified as short or long term depending on whether the expiration is under or over 40 days away, and classified as in-, at-, or out-of-the-money based on Black–Scholes delta cutoffs (in absolute value) of 0.375 and 0.625. (The table restricts attention to the dominant SPX class of options.)

12 Of course, if long and short gamma players both maintain delta-neutral positions, then their rehedging demands would offset, that is, they could always trade with each other. In practice, market-makers—who tend to be short gamma—are more likely to rehedge than are end users.

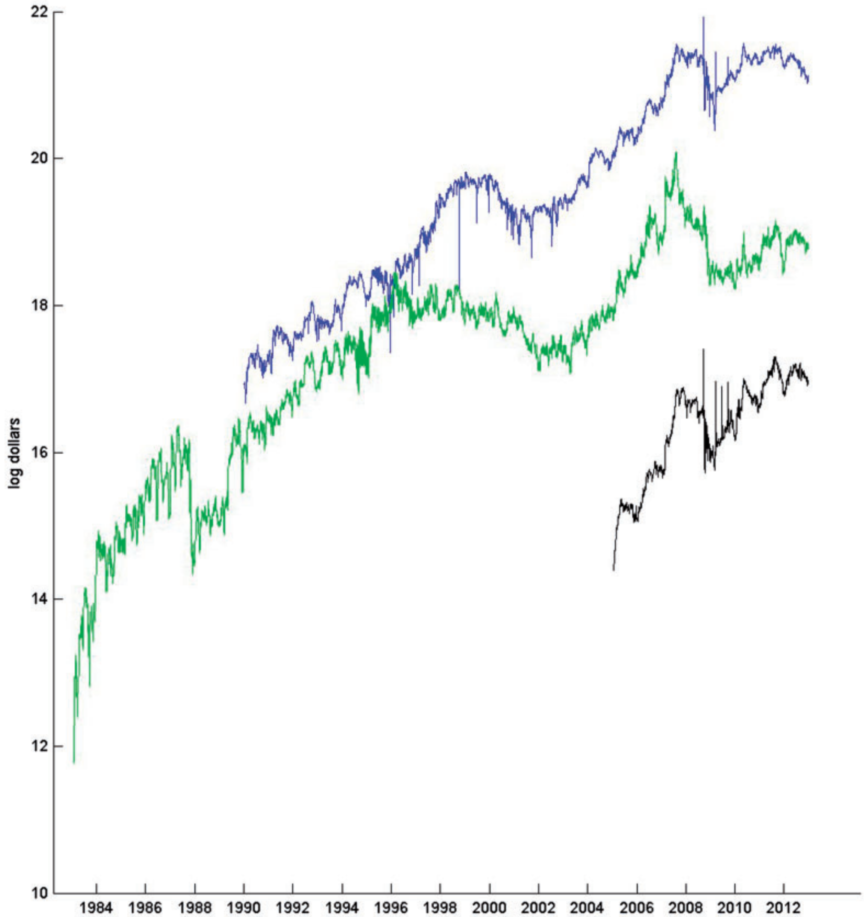

Figure 3. Aggregate vega.
The top line is the total volatility sensitivity of all outstanding SPX options listed on the CBOE. The middle line is the volatility sensitivity of all CME-listed options on S&P 500 futures. The lowest line is the volatility sensitivity of all options on SPDR ETFs. The units are logarithms of dollar value per 100 basis point change in annualized volatility.

Investor positions are concentrated in out-of-the-money puts. As of the end of the sample these constituted more open interest (51%) than all other types combined. Note that, by put-call-parity, the exposure transferred by low strike calls is essentially the same as that of low strike puts.

While longer term options represent about two-thirds of open interest, short-term options are traded about twice as often relative to their open interest. This is shown in the second panel, which compares trading volume to the (absolute value) of open interest changes. A value of 3, for example, indicates that a third of trades alter investors' net exposure.

The picture of a market dominated by low strike price put positions is consistent with index options being used primarily for crash insurance. Indeed, the amount of crash insurance can be quantified with our aggregate gamma numbers. For any position hedged with

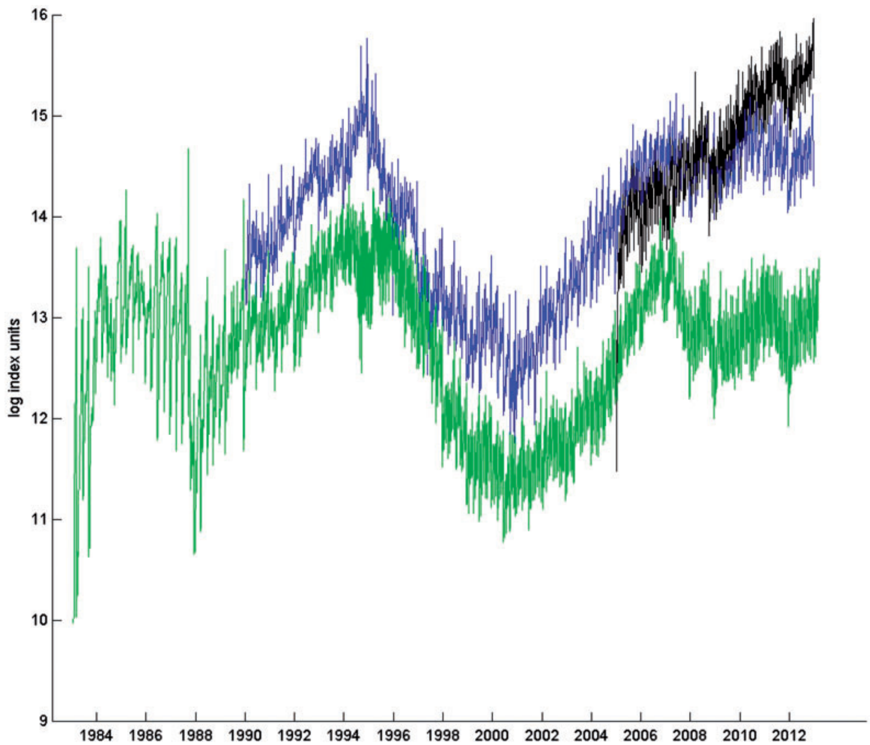

Figure 4. Aggregate gamma.
The top line is the total delta sensitivity of all outstanding SPX options listed on the CBOE. The middle line is the delta sensitivity of all CME-listed options on S&P 500 futures. The lowest line is the delta sensitivity of all options on SPDR ETFs. The units are logarithms of index equivalent units per 1% index change.

respect to market direction, the leading term in a Taylor series expansion for the change in position value in response to a jump of size j is

$$\frac{1}{2}\Gamma j^2,$$

where Γ is the second partial derivative of the position value with respect to the underlying. Aggregating this quantity as of December 2012, Figure 5 shows the potential monetary transfer as a function of the jump percentage. For moderate jumps, the number is in the hundreds of billions of dollars and approaches trillions for a severe crash. In fact, in the case of a downward jump this number is an understatement since we know from above that most open positions are out-of-the-money puts. The graph also shows the change in position value when the aggregate position is approximated as a three-month put with strike chosen to match the aggregate market value and gamma.[13]

13 In fact, both approximations are still understatements for a large move (in either direction) since such a move would also substantially raise implied volatilities. We thank a referee for suggesting this analysis.

Table II. S&P 500 options quantities

The table reports option quantity measures on two dates. For comparison, the first and fourth columns give the capitalization of the S&P 500 (in billions of dollars) and the effective number of index shares outstanding (in billions of index units) defined as the capitalization divided by the index level. Market value (MV) and aggregate vega (AV) are in billions of dollars. Open interest (OI) and aggregate gamma (AG) are in billions of index-equivalent units.

	Index Cap	MV	AV	Index shares	OI	AG
End of sample—December 31, 2012						
SPX		34.75	1.456		0.909	0.0264
Futures		1.85	0.142		0.136	0.0046
SPY		0.38	0.023		0.183	0.0675
Total	13,121	36.98	1.621	9.200	1.229	0.0985
Peak size—September 22, 2008						
SPX		102.93	3.316		1.452	0.0250
Futures		4.56	0.213		0.204	0.0041
SPY		1.09	0.036		0.147	0.0325
Total	10,836	108.57	3.566	8.980	1.804	0.0615

Table III. S&P 500 options quantities by type

The top panel table reports the average fraction of aggregate open interest in CBOE-listed SPX options comprised of calls and puts; long (over 40 days) and short time to expiration; and strike prices below, near, or above the current index level. The strike price breakpoints are determined by the Black–Scholes delta of the respective options, with low strike puts having delta greater than −0.375, low strike calls having delta greater than 0.625, high strike puts having delta less than −0.625, high strike calls having delta less than 0.375. The bottom panel shows time series averages of each month's ratio of volume to absolute change in open interest. The sample period is January 5, 1990 to December 31, 2012.

	Low strike	At-the-money	High strike
Fraction of open interest			
Calls, long-term	0.0742	0.0660	0.1067
Calls, short-term	0.0538	0.0206	0.0762
Puts, long-term	0.2939	0.0597	0.0345
Puts, short-term	0.1668	0.0196	0.0281
Average volume over change in open interest			
Calls, long-term	1.11	1.74	1.67
Calls, short-term	2.95	2.98	3.23
Puts, long-term	1.78	1.90	1.23
Puts, short-term	3.17	3.10	2.78

Our findings on the economic magnitude of the index option market are not altered by the consideration of the broader numbers in the BIS survey. Bearing in mind the difficulties in comparison with OTC numbers previously described, we report in Table IV the "notional" amounts from the BIS of each class of derivative alongside our own aggregated

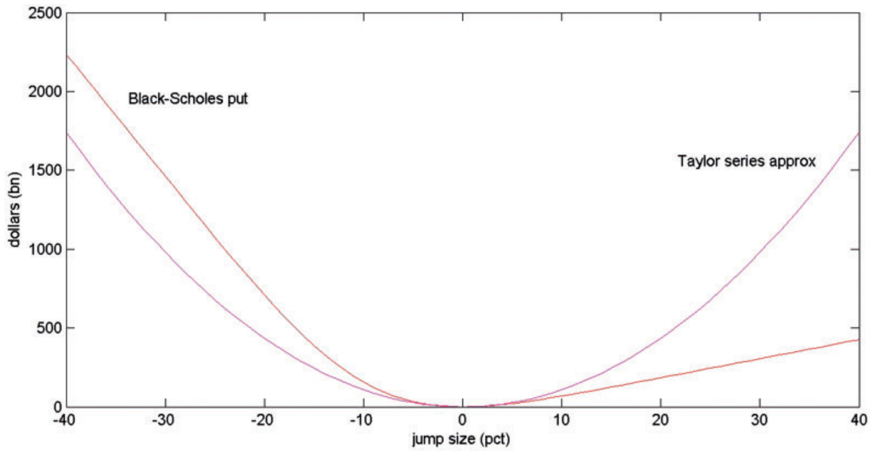

Figure 5. Jump risk.
The figure plots the change in aggregate option value in response to a jump in the index using the December 2012 aggregate gamma. The plot shows both the Taylor series approximation $\frac{1}{2}\,\Gamma\, j^2$ and the value computed from the Black–Scholes model when the total open position is approximated by a 3-month put chosen to match the aggregate market value and gamma. Both approximations hold implied volatility fixed.

Table IV. All index option markets

The table reports outstanding amounts of equity index options as of December 2012. The first column aggregates the contracts used in this study, and is a daily average over the month. The remaining columns are from the BIS semi-annual derivatives surveys: http://www.bis.org/statistics/derstats.htm. For the second and third columns, the US numbers include North American exchanges. The OTC columns count all equity options not just those referencing indices.

	Notional amount outstanding (USD billion)				
	Listed S&P 500	US listed	non-US listed	US OTC	non-US OTC
December 2012	2,320.0	1,934.8	2,005.8	1,266.7	2,940.2

number as of December 2012. Notional amounts are the value of the underlying equity index referenced by the outstanding options and thus correspond to our series OI(t) multiplied by $S(t)$, the value of the underlying index on each date.[14]

The numbers in the table show that our study is not missing any significant US listed index options.[15] Similarly, given the overcounting of OTC positions, the US OTC index

14 In practice, OTC participants may report the notional amount as the value of the underlying stock on each option at the exercise price. For example the right to buy $100 million dollars of stock at price K would be equivalent to a listed option on $100 million-divided-by-$K$ units which we would record as a notional amount of (S/K) times $100 million, whereas OTC participants may report $100 million as the notional value.

15 The fact that our contracts appear to exceed the size of a set that strictly includes them may be due to timing differences in data collection.

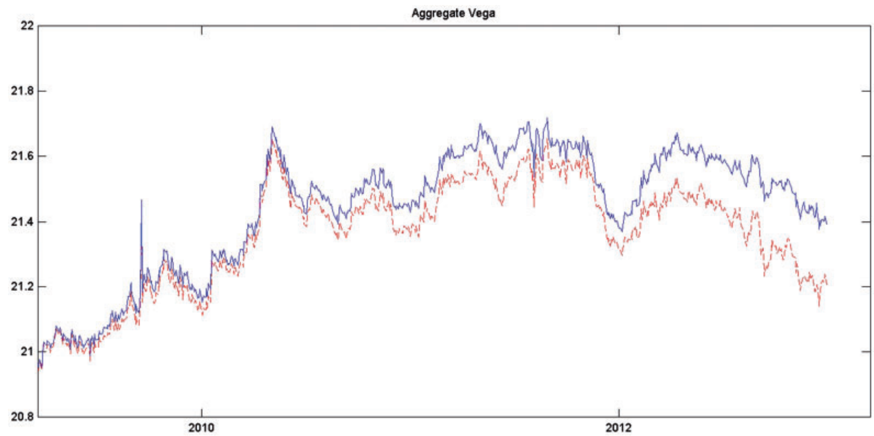

Figure 6. Listed volatility futures.
The lower dashed line is the aggregate vega of listed S&P 500 options. The solid line includes the contribution of listed futures on the VIX index. The units are logarithms of dollar exposures.

exposure not captured in our data, while substantial, is at least not larger than the ones we cover. Last, non-US index options, listed and OTC, are comparable in size to our data set.

The previous section also noted the emergence of listed derivatives on stock market volatility as a related and possible substitute product class. Although volatility and variance swaps have been actively traded in OTC markets for at least 10 years, the first exchange-traded products were cash-settled futures tied to the VIX index introduced in 2004.[16] For unclear reasons, investor interest in these products—and their liquidity—took off in the last few years. By the design of these contracts, their open interest is directly comparable to our AV measure of volatility exposure in the index options market.

Figure 6 shows our aggregate vega series both alone (solid line) and with the additional total vega of the futures (dashed line) since 2009. As of the end of the sample, the difference in the series is around 0.25, meaning that the futures account for about a fifth of the total volatility exposure. We conclude that this exposure is comparatively substantial (and rapidly growing), but is still small in economic magnitude.

4. The Stochastic Trend in Options Exposure

We now turn to the question of characterizing the trend driving the evolution of options exposure. An initial observation about the four series is that the stochastic component appears to be common among them. If we express the gamma, vega, and market value series on a per-contract basis by dividing by open interest, and also normalize the latter two series by also dividing by the index level,[17] the resulting series (shown in logs in Figure 7) appear stationary. Standard tests for unit roots reject the null of nonstationarity, consistent with the visual impression.

16 We are unaware of data series on quantities of OTC volatility derivatives.

17 Option values and vegas scale mechanically with the level of the underlying. With our unit definition of gamma, that series does not scale with the underlying.

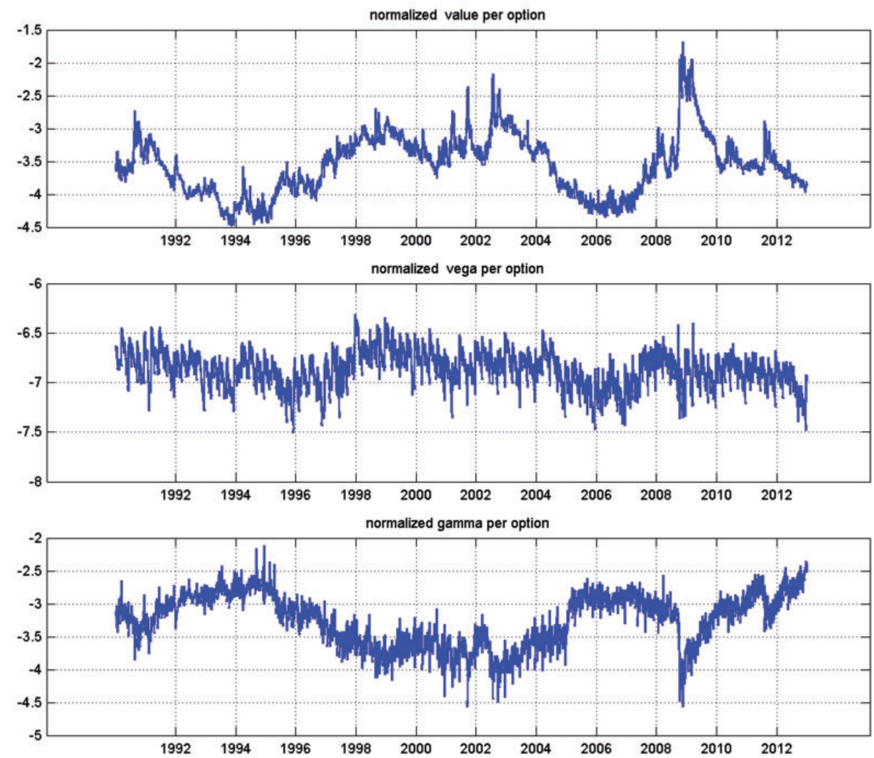

Figure 7. Re-scaled options exposures.
The figure shows scaled versions of (from top down) MV, AV, and AG. Each series is expressed on per-contract basis by dividing by OI. MV, AV are divided by the level of the S&P 500 index. The series are shown in logarithms.

From this analysis, it follows that the composition of the options outstanding in terms of their convexity and market value is relatively stable over time. We can, therefore, focus our attention on a single quantity series—aggregate open interest—in modeling the nonstationary component of options positions. Note that OI, measured in index-equivalent units, is not mechanically related to the level of the stock market or its volatility.

In seeking to explain the stochastic trend in OI, we refer to the hypotheses outlined in Section 1 and consider possible trends in investor heterogeneity and/or market frictions. Here, a natural prior hypothesis is to link options market trends to trends in other measures of financial activity that are likely to be driven by the same forces. In particular, equity market trading volume should also reflect differences in preferences and beliefs, as well as trends in trading costs and technology in the financial industry. To formally assess the hypothesis, we perform cointegration tests using the log of median daily turnover, TO, of stocks whose primary listing is on the NYSE and American Stock Exchange.[18]

18 The volume measures encompass trading consolidated across all exchanges, trading platforms, dark pools, etc. The turnover series is computed as a trailing 21-day average, dropping days adjacent to holidays.

Although detrending OI with equity market turnover will turn out to be successful, it does raise the question of what drives turnover. That topic is itself an important open question that is beyond the scope of the present work. However, it is reasonable to ask whether the evolution of options exposures can be associated more directly with measures of population heterogeneity or investment frictions. We propose a few ideas.

One hypothesis is that an important dimension of heterogeneity is of the type of investors or institutions that are active in the market. Rather than attempt to find proxies for unobservable preferences or beliefs, this approach posits that these characteristics may be common to a class of participants. Demand for and supply of options may then be associated with changes in the relative size of different types of participants. We assess this hypothesis using data from the Federal Reserve's quarterly $Z.1$ reports on ownership of US corporate equity. From these, we construct the market share of five classes of participants: households, mutual funds and ETFs, pension and insurance funds, foreign institutions, and brokerage firms and securities dealers. Shares of equity ownership are, of course, not direct measures of the population of index options market participants. But it is certainly plausible that trends in the two are closely related. Our tests assess whether the trend in OI is explained by the trend share of any one of the five types.

Next, we consider potential measurement of investment frictions. Over the sample period, the cost of transacting in the stock market declined dramatically for most investors due to a combination of cheaper information processing, data availability, and competition among intermediaries. Commissions and direct trading costs have trended strongly downward while market liquidity has trended upward. Borrowing costs, another potential constraint on trading activity, also trended strongly downward during the sample period.

Note that the theoretical prediction here would be that options demand would decline with interest rates if investors use options precisely because of their embedded leverage. Likewise, to the extent that the underlying markets for index stocks become cheaper to trade, it becomes easier to replicate any nonlinear payoff achievable with options, suggesting a decline in demand. On the other hand, clearly cheaper options trading itself would be expected to increase options usage. Similarly improvements in trading infrastructure to lower technological barriers to trading options should promote usage.

With these considerations in mind, we include in our analysis a measure of average trading costs in the options market, a general information processing proxy, and the three-month eurodollar interest rate. The trading cost proxy is an average of closing percentage bid/ask spreads for the CBOE SPX options. The technology variable is an index of the cost of household information technology, hardware, and services compiled by the Bureau of Labor Statistics (BLS).

For each of the variables described above, we undertake a series of bivariate cointegration tests with the time series of the log of aggregate open interest. Table V shows the resulting test statistics for the null of no cointegration. Rejections are thus evidence of a common stochastic trend.

The tests show positive evidence of a common stochastic trend between OI and TO, but no evidence that any of the other variables shares a common trend with OI. [19] Evidently,

19 The Johansen statistics for the bid-ask series and the pension and insurance share are misleading because this test considers arbitrary combinations of the two variables. Here, the estimated weighting on OI in the cointegrating vector is actually zero. The test is simply saying that the other two series on their own look stationary. In contrast, the other two test statistics are based on

Table V. Cointegration analysis of aggregate open interest

The table reports results of three cointegration tests for bivariate systems, the first element of which is the log of the aggregate open interest in list S&P 500 options and the second element is the variable shown in the first column. TO is the median daily turnover for the prior month of stocks whose primary listing is the NYSE or AMEX. The investor composition variables, s, are the relative shares of each category in total US equity ownership as of the end of the previous quarter. bid/ask is the open-interest weighted average closing percentage bid/ask spread of CBOE SPX options. CPI_{info} is the real household cost of information technology, hardware, and services compiled by the BLS. $r^{\text{ED}}_{\text{3mo}}$ is the three-month eurodollar interest rate. The second column gives the test statistic for the test of Phillips–Perron (1988); the third column gives the test statistic for the Engle–Granger (1987) test; the fourth column shows the Johansen (1991) trace test statistic. One, two, and three asterisks denote rejection of the null of no cointegration at the 90%, 95%, and 99% thresholds, respectively. All tests use three months of daily lags. The sample period is January 5, 1990, to December 31, 2012.

Variable:	Phillips–Perron:	Engle–Granger:	Johansen:
Financial activity			
log TO	−4.44***	−3.63**	14.70*
Investor composition			
$s^{\text{household}}$	−1.57	−1.57	6.41
s^{funds}	−1.46	−1.19	5.37
$s^{\text{pens+ins}}$	−2.03	−2.01	14.74*
s^{broker}	−1.12	−1.04	10.01
s^{foreign}	−2.09	−1.90	9.71
Frictions			
bid/ask	−1.52	−1.12	26.28***
CPI_{info}	−1.96	−1.43	11.80
$r^{\text{ED}}_{\text{3mo}}$	−0.69	−0.43	7.44

the remaining variables fail to capture relevant dimensions of heterogeneity or costs that affect options activity. The explanatory power of turnover cannot, then, be attributed to these factors.[20]

Although it remains an important and interesting topic to explain what exactly equity market turnover is capturing, for present purposes the conclusion is that the same trend driving it also drives options usage. Is reverse causality a possibility? We previously documented that the scale of options position was large enough to potentially induce substantial equity trade (via rehedging of delta-neutral positions). We are not able to measure actual options hedging trades in the stock market. However, we can measure trading that is not plausibly attributable to index options: namely, activity in non-index stocks. When we construct our turnover measure using exclusively these stocks, the cointegration results are stronger than those reported in the table.[21] It is thus reasonable to conclude that the exogenous component of equity market activity drives options positions.

regression residuals that take OI as the dependent variable. The remaining variables examined in the table all pass standard unit root tests.

20 This conclusion is robust to consideration of multivariate cointegrating relationships. Results are omitted for brevity.

21 All three tests statistics reject at the 99% level. Results are available upon request.

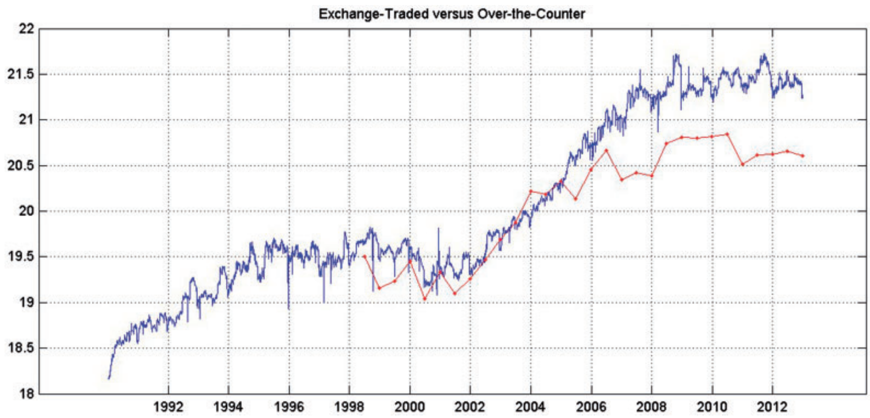

Figure 8. Aggregate open interest.
The figure plots aggregate open interest for exchange-traded and OTC US equity index options. The OTC series is from the BIS and is converted from net notional value into S&P 500 equivalent units. Both series are in logarithms.

To close this section, we consider whether our conclusions about the driving trends in options exposure are likely to be sensitive to the omission of data from the OTC market. Figure 8 compares the evolution of our aggregate open interest series with that of the OTC series from the BIS, when the latter is expressed in terms of equivalent S&P 500 units. (Both series are in logs.) The two series track each other quite closely from 1998 to 2006, coincidentally almost being equal in magnitude. However, there has been a pronounced erosion in the OTC market share in recent years. In raw terms, the OTC market now accounts only half as much net notional exposure as the exchange-traded products.[22]

The graph suggests that combined index options positions from the two venues may have declined more since 2007 than our exchange-traded series alone indicates. On the other hand, the two series do move more-or-less in parallel from 2008 onward. Thus, we conclude that our conclusions about the broad trends since 1990 are unlikely to be biased by the absence of OTC data.

We have now answered an important part of the question in the paper's title. There are noteworthy findings, both positive and negative. We do not find support for theories that posit that options usage is driven by borrowing constraints or relative transactions costs vis a vis the underlying market. Nor are we able to associate options quantities with a particular class of investor participation. Instead, we find the stochastic trend that drives options positions is well described by that of equity market turnover. We have posited that the latter trend reflects trends in heterogeneity in beliefs or preferences, as well as the cost of trading. Substantiating that conjecture remains a topic for future research.

22 The gap between the series dates from the first half of 2007, prior to the financial crisis or to any subsequent regulatory developments. A possible explanation may be the industry's own emphasis on more efficient trade netting and compression of OTC positions that had arisen from the soaring trade in credit derivatives during the expansion preceding the crash.

5. Risk and Index Options Quantities

We now turn to the topic of modeling options quantities in terms of their deviation from the stochastic trend documented in Section 4. As discussed in Section 1, theoretical considerations imply that the amount of risk transfer should vary with the degree of risk. We therefore consider here a number of measures of different types of risk.

The dependent variable in this section is the stochastically detrended log open-interest series. Specifically, the full-information maximum-likelihood (FIML) estimator for the cointegrating coefficient in the Johansen (1991) estimation is approximately 1.70. So the finding above is that $\log \mathrm{OI} - 1.7 \log \mathrm{TO}$ is stationary. It is this difference that we study here.[23] Our analysis reports results of levels regressions. These are consistent under the null of stationarity.[24] As a robustness check, in Appendix A, we report results of nonstationary regressions in which turnover is on the right-hand side.

To start, we consider measures of investor perception of risk. We have two convenient time series that characterize the most important features of the S&P implied volatility surface since 1990. The well-known VIX series (the so-called new version) provides a model-free estimate of risk neutral expected volatility for the next 30 days. More recently, the CBOE has also constructed an analogous model-free estimate of the skewness of implied volatility.[25] Since this skewness is typically negative (i.e., the return distribution is left-skewed) the CBOE value has a minus sign in the definition. Moreover, it is a measure of standardized skewness. That is, the variance has been scaled out of it: the correlation of VIX and SKEW is -0.05. This is by construction. A SKEW value of 120, for example, indicates intuitively that there is 20% more mass in the left tail than the right tail. It does not indicate an absolute amount of probability in the left tail. We therefore interact VIX and SKEW to capture the latter information as well.

Table VI presents regressions of detrended log open-interest on these variables. The first row shows that there is a significantly negative response of options quantities to VIX. This effect is economically large. All variables are in logs, and the standard deviation of VIX is approximately 0.35. So, using the smallest coefficient in the first column, a one standard deviation increase in VIX implies a 15% decline in open interest, and this alone explains over 20% of the variability in OI.

When SKEW is included, in the second column, it too enters with a negative response, although insignificantly. However, this becomes very significant when the interaction term is included (third column).The interaction of VIX and SKEW is positive and highly significant, meaning that for large values of either variable, the marginal response to the other will be positive. Figure 9 plots the two-dimensional response surface over a range of approximately ± 3 standard deviations of both independent variables. From the scale of the vertical axis, the effects are economically large. For reference, the standard deviation of the dependent variable is 0.39.

23 The series still contains a strong seasonal component, notably shrinking over the fourth quarter of each year and rebounding in the first quarter. We remove this component in the regression analysis below by subtracting the full-sample average growth rate for each quarter. There is also a small seasonal related to the option expiration cycle. The results below are not sensitive to removing this component.

24 Estimation error in the cointegrating relationship adds further noise to *yt*, which works against our ability to detect significant covariates.

25 Both series are available from http://www.cboe.com/micro/vix/historical.aspx.

Table VI. Risk effects in index option quantities

The table shows OLS regressions of aggregate index option quantities on volatility statistics measures of risk. The dependent variable is the log of aggregate open interest detrended as described in the text. VIX and SKEW are model-free measures of the second and third moments of the option-implied distribution of 30-day ahead index returns computed by the CBOE. σ_{VIX} is the daily estimated volatility of volatility. BBD is a series from Baker, Bloom, and Davis (2013) of intensity of economic policy news stories. MM is an index of uncertainty news intensity used in Manela and Moreira (2015). All variables are in logarithms. Newey–West (1987) *t*-statistics (shown in parentheses) are computed with 1 year of lags. The sample period is January 5, 1990, to December 31, 2012.

	(1)	(2)	(3)	(4)	(5)
VIX	−0.5561	−0.5586	−1.6351	−1.6143	−1.4452
	(4.00)	(4.04)	(7.41)	(8.17)	(6.15)
SKEW		−0.3452	−20.5192	−18.9961	−10.3758
		(0.47)	(6.49)	(6.58)	(2.94)
VIX* SKEW			6.8515	6.3512	3.5318
			(6.64)	(6.70)	(2.98)
σ_{VIX}				0.3224	0.3849
				(3.12)	(4.91)
BBD					−0.3055
					(3.90)
MM					0.7652
					(4.82)
R^2	0.24	0.25	0.32	0.36	0.51
N	5,776	5,776	5,776	5,776	5,042

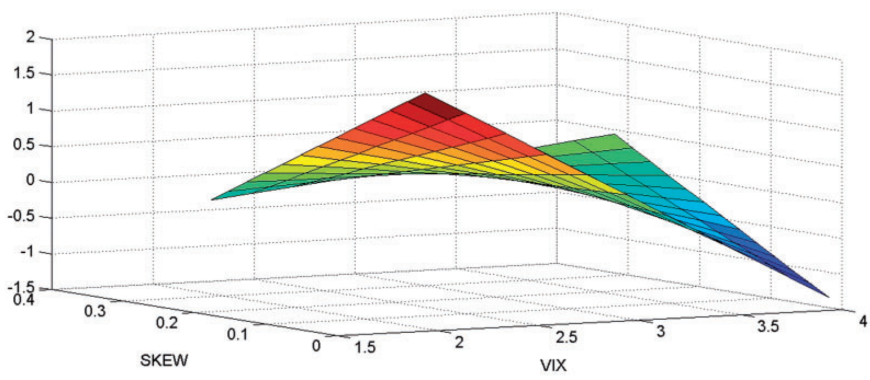

Figure 9. Effect of risk on options quantities.
The figure shows the response surface estimated in Table VI of open interest to the second and third moment of the risk-neutral distribution of index returns. All variables are in logs.

We will discuss the economic interpretation of these responses below. Intuitively, they tell us that, under most market conditions, increases in variance and left-tail risk elicit negative quantity responses. However, the effect may be reversed when both VIX and SKEW are above their means, which can be viewed as times of high market stress.

The third column of the table adds another dimension of risk: the volatility of volatility. The motivation for this variable is the theory that a primary purpose of index options may be volatility risk transfer. We estimate this series by fitting an EGARCH model (Nelson, 1991) to daily changes in log of VIX itself.[26] We find an independent and statistically significant positive response to this variable. High volatility of volatility again indicates conditions of high market stress, in which both tails of the return distribution expand.

The rightmost column augments the options-derived measures of risk perception with two others recently introduced into the literature. Baker, Bloom, and Davis (2013) compile the frequency of articles combining terms related to uncertainty with terms related to government policy from 10 US newspapers. The authors show that their index is increasingly correlated with measures of stock market implied volatility as the forecast horizon increases. In other words, it may be seen as capturing long-term volatility as opposed to the 30-day horizon used in the construction of VIX. Manela and Moreira (2015) similarly use a machine-learning algorithm to identify uncertainty-relevant language appearing on the front pages of the Wall Street Journal and tabulate its intensity.[27] Based on extensive analysis of drivers of variation in their index, these authors argue that the component that is not related to stock market uncertainty is strongly correlated with the probability of rare disasters.

Our regressions find independent and economically large explanatory power for both variables in explaining option quantities. The regression R^2 now exceeds 50%. The signs of the effects are opposite, despite the similarity of construction. Intriguingly, interpreting each of the indexes as proposed by their creators, the message again seems to be of a negative response to (long-horizon) variance, with a positive response to extreme stress (disasters).

To get a closer look at quantity responses to variance and skewness, we run similar regressions for open interest of different types of options. Table VII shows the dynamic responses to VIX and SKEW of open interest for each of the twelve buckets defined in Section 3. To avoid mechanical effects, these regressions are run in differences, holding the basket breakpoints fixed across successive days.[28] Open interest changes for each basket are expressed as a fraction of the trailing three-month average level of open interest in that basket. The table shows the sum of the contemporaneous response and one-lag response coefficients.

The table shows that the unconditional negative response to VIX documented above is driven by low strike price options, with high strike and at-the-money options actually

26 To perform this estimation on the longest possible time series, we create a new volatility proxy using the CME data, which goes back to 1983. This series is constructed using a kernel density weighting of CME option implied volatilities, wherein the weighting on moneyness and expiration are chosen to maximize the fit to VIX in the post-1990 period. Details are available upon request.

27 We are grateful to the authors for making this series available to us.

28 Specifically, at date t we examine the change in open interest from $t-1$ to t for options whose strike and maturity place them in basket k at date t. The regressions also include a lag of the dependent variable.

Table VII. Risk responses by option type

The table reports regressions of the change in open interest for options classified by type as described in the caption to Table III. The independent variables are the contemporaneous change in VIX and SKEW and one daily lag of each as well as of the dependent variable. The sum of the response coefficients to VIX and SKEW are reported. One, two, and three asterisks denote rejection of the null that the coefficient sums are equal to zero at the 90%, 95%, and 99% thresholds, respectively. The sample period is January 5, 1990, to December 31, 2012.

	Low strike	At-the-money	High strike
VIX			
Calls, long-term	−0.1314***	0.0080	0.0852***
Calls, short-term	−0.1683***	−0.0272	0.0409**
Puts, long-term	−0.0489***	0.2112***	0.0944***
Puts, short-term	0.0213	0.3323***	0.1414***
SKEW			
Calls, long-term	0.1060	0.1885**	−0.1536**
Calls, short-term	0.1343	0.2333**	−0.2095***
Puts, long-term	−0.0872	0.1019	−0.0033
Puts, short-term	−0.0509	0.1357	−0.3374***

responding positively. As we saw in Table III, long-term, low strike puts comprise about 30% of all open interest. By put-call-parity, low strike calls are effectively equivalent to low strike puts with respect to crash risk. These represent an additional 7–13% of open interest. Chen, Joslin, and Ni (2014) document that closing of open positions in far out-of-the-money index puts predicts (positive) stock market returns and attribute this to market makers reacting to tightening VaR constraints by shedding risky positions in general (i.e., including stocks themselves). The pattern we document is consistent with covering of crash risk positions by protection sellers in response to increased market variance. However, it is well known that VIX does not predict future stock returns. So the position covering isolated by Chen, Joslin, and Ni (2014) is likely driven by something other than dealer responses to increases in market risk.

Low strike put positions also contract in response to increased left tail risk (although now there is a positive response in the low strike calls). Perhaps surprisingly, positions exposed to right-tail risk—high strike puts and calls—also contract on an increase in SKEW. In contrast, at-the-money positions, which have the highest gamma, expand.

Risk measures derived from options prices reflect both risk and risk aversion (marginal utility) in investor perception of future states. That is, they reflect the “risk neutral” distribution of stock returns, rather than the “physical” or “true” distribution. This raises the question of whether option quantity responses in Table VI are reflecting the endogenous risk aversion component or the exogenous risk component. To get at this issue, we do three things. First, we decompose VIX into a statistical forecast of true market return variance and a residual that can be interpreted as the variance risk premium (VRP). Second, we include a gauge of exogenous economic risk not derived from options prices or stock returns. Third, we include proxies for investor sentiment (or risk aversion).

For the variance return forecasts, we fit an EGARCH model to daily futures returns since 1983 that includes asymmetry effects and leptokurtosis (via *t* distributed innovations). The model delivers one-day-ahead expected return variances every day. We iterate the forecasts to 30-day horizon for comparability with VIX.

The exogenous risk measure, taken from Jurado, Ludvigson, and Ng (2013), is a formal estimate of the conditional second moment of a broad cross-section of macroeconomic time series at the monthly level. The authors estimate prediction equations for 132 variables[29] that explicitly allow for time-varying first and second moments. They then average the forecast standard deviations to different time horizons. We employ their three month measure $\bar{\mathcal{U}}^y(3)$.

For risk aversion, we employ the University of Michigan consumer confidence index, cns.cnf, as well as the lagged quarterly return on the S&P 500 index itself, r^m_{1Q} to proxy for possible extrapolative expectations.

Table VIII shows regression results with these controls.[30] The first column establishes that the negative response of OI to VIX is driven by the true variance component. There is also some evidence of a negative response to the VRP, but the statistical significance is not robust across specifications. Column (2) confirms this conclusion by showing that the negative response to the Jurado, Ludvigson, and Ng (2013) measure of "fundamental" risk is even stronger than the response to forecasted stock market volatility.

Moreover, the risk aversion proxies in Equations (3) and (4) do not appreciably change the uncertainty effect. Interestingly, when both these variables and $\bar{\mathcal{U}}^y(3)$ are included (in Equation (4)), we see a substantially larger coefficient on exogenous risk and a significant negative response to the consumer confidence measure.[31] If this measure is indeed picking up investor risk aversion, and if risk aversion was partially responsible for the negative variance response, then one would have expected a positive response to cns.cnf (since lower values indicate pessimism).

Column (5) in the table includes the BBD measure that we interpreted above as potentially capturing longer horizon variance, but which is a measure of investor concerns and thus also reflects risk and risk aversion. This variable again enters significantly negatively and more strongly than in Table VI. It also strengthens the negative response to consumer sentiment, suggesting that the BBD response is due to the fundamental, exogenous component of risk in this index.

Finally, since Table III shows that out-of-the-money puts comprise, on average, almost half of open interest, we repeat specification (5) using only these options as the dependent variable.[32] Our specifications have less overall explanatory power for the low-strike puts.

29 The forecasts use information in 279 time series: the macro series plus 147 financial predictors. The uncertainty measures are taken from http://www.econ.nyu.edu/user/ludvigsons/jlndata.zip.

30 For brevity, the table focuses on second-moment proxies and omits the ones from Table VIII that we interpreted as higher measures of tail risk.

31 The finding that there is no significant response of index options positions to market returns is robust to the return horizon: lagged one-week and one-month returns also have no significant effects.

32 When using out-of-the-money puts, the regressions also control for contemporaneous market returns because of the mechanical effect: a fall in the market automatically leads to fewer out-of-the-money strike prices, etc.

Table VIII. Risk effects: physical versus risk-neutral uncertainty

The table shows OLS regressions of aggregate index option quantities on measures of physical and risk-neutral uncertainty, and risk-aversion proxies. $E[\sigma^2]$ is the 30-day-ahead forecast variance from an EGARCH model of market returns. VRP is the difference between the VIX index (squared) and this forecast. $\bar{\mathcal{U}}^y(3)$ is the index of macroeconomic uncertainty at the three-month horizon created by Jurado, Ludvigson, and Ng (2013). cns.cnf is the University of Michigan index of consumer confidence. r^m_{1Q} is the trailing three-month return on the market index. Column (6) repeats the specification in Equation (5) with open interest in out-of-the-money puts as the dependent variable. All the risk variables are in logarithms. Newey–West (1987) *t*-statistics (shown in parentheses) are computed with one year of lags. The sample period is January 5, 1990, to December 31, 2012.

	(1)	(2)	(3)	(4)	(5)	(6)
$E[\sigma^2]$	−0.3623	−0.1961	−0.3724	−0.1882	−0.1225	−0.1105
	(4.50)	(2.75)	(4.08)	(2.39)	(1.96)	(2.25)
VRP	−0.1824	−0.0827	−0.1903	−0.0705	−0.0359	0.0658
	(1.86)	(0.98)	(1.94)	(0.89)	(0.46)	(1.21)
BBD					−0.4724	−0.4402
					(4.06)	(5.02)
$\bar{\mathcal{U}}^y(3)$		−1.5572		−2.1098	−1.4390	−0.7622
		(3.14)		(4.52)	(3.37)	(2.14)
cns.cnf			−0.0016	−0.0076	−0.0165	−0.0166
			(0.45)	(3.03)	(4.17)	(5.45)
r^m_{1Q}			−0.0284	−0.2164	−0.3419	−0.5868
			(0.07)	(0.55)	(1.02)	(1.76)
R^2	0.29	0.40	0.29	0.46	0.58	0.49
N	5,789	5,548	5,789	5,548	5,548	5,548

The responses to the uncertainty variables are somewhat weaker, but still statistically significant.

Another important distinction in the measurement of risk is between uncertainty and disagreement. It seems plausible that differences of opinion between market participants are correlated with fundamental risk. This can be investigated using measures of survey dispersion. As discussed in Section 1, there is a strongly grounded prediction for a positive relationship between measures of heterogeneity—like differences of opinions—and quantities of options positions. We have, however, documented the opposite relationship to volatility measures. This implies that, controlling for uncertainty, the regressions can isolate the belief differences as components of survey dispersion.

We have two dispersion measures. The first (called BOS) is from Bachmann, Elstner, and Sims (2013), and is constructed from the dispersion in responses to the Business Outlook Survey conducted monthly by the Philadelphia Federal Reserve.[33] This is a regional survey of executives of manufacturing firms and concerns their operational climate. In contrast, the Survey of Professional Forecasters (SPF), also compiled by the Philadelphia Federal Reserve, is a quarterly assessment of the national economic outlook as gauged by

33 This series and the BBD measure are from http://www.aeaweb.org/articles.php?doi=10.1257/mac.5.2.217.

Table IX. Risk effects: uncertainty versus disagreement

The table shows OLS regressions of aggregate index option quantities on measures of survey dispersion and other controls. BOS and SPF are constructed from, respectively, the Business Outlook Survey and the SPF. Both surveys are compiled by the Federal Reserve Bank of Philadelphia. Column (5) repeats the specification in Equation (4) with open interest in out-of-the-money puts as the dependent variable. Newey–West (1987) *t*-statistics (shown in parentheses) are computed with one year of lags. The sample period is January 5, 1990, to December 31, 2012.

	(1)	(2)	(3)	(4)	(5)
BOS	0.4863	1.1526	1.3196	0.8377	0.5162
	(0.71)	(2.24)	(2.95)	(2.67)	(1.80)
SPF	−0.1417	0.1163	0.0442	0.0321	−0.1437
	(0.94)	(1.20)	(0.48)	(0.47)	(2.72)
$\bar{\mathcal{U}}^y(3)$		−2.5481	−3.1221	−2.1458	−1.0752
		(5.35)	(6.86)	(4.55)	(2.72)
cns.cnf			−0.0099	−0.0173	−0.0187
			(4.18)	(4.91)	(7.61)
BBD				−0.4249	−0.4149
				(4.80)	(6.15)
R^2	0.03	0.42	0.50	0.59	0.49
N	5,548	5,548	5,548	5,548	5,548

economists. We employ the interquartile range of forecasts for real GDP growth in the current quarter.[34]

Table IX shows that there is no evidence that these dispersion measures relate to options positions on their own (first column) but that, controlling for macroeconomic uncertainty (second column) the predicted positive role for disagreement does emerge, at least using the BOS survey.[35] The conclusion is reinforced by inclusion of the indexes for consumer confidence and policy uncertainty (in Equations (3) and (4)). Specification (5) uses only the out-of-the-money put open interest, which yields less positive dispersion responses, comparable to those in Equation (1).

Although the incremental explanatory power for index option positions of the BOS measure is small in our regressions, it is perhaps remarkable to find an association at all given that the topic of the survey (local manufacturing outlook) is several steps removed from the direct concerns of the index options markets.

34 These measures only capture disagreement about the first moment of economic activity (i.e., expected growth), and do not speak to difference in beliefs about higher moments.

35 Our findings here accord with those of Buraschi and Jiltsov (2006) who build a model of options trade driven by differences of belief about the economic state and calibrate it using a survey-based measure of disagreement. Overidentifying tests fail to reject the model's restrictions on open interest changes for a cross-section of S&P 500 options during 1986–1996. Regression evidence is moderately supportive of a positive association between option volume levels and belief differences.

Summarizing the results in this section, we highlight the following conclusions:

1. The most significant determinant of fluctuations in OI is a negative response to risk, which is not driven by the risk-aversion component in measures of investor sentiment, nor by differences of opinion. For example, regressions including only $E[\sigma^2]$ and $\bar{\mathcal{U}}^y(3)$ achieve an R^2 of 40%.
2. Increases in belief dispersion and deterioration of investor sentiment induce significant positive responses of OI. Even with limited proxies for these, explanatory power approaches 60% (e.g., specification (4) in Table IX).
3. There is also evidence of a positive quantity response to some measures of tail risk (the VIX-SKEW interaction, volatility of volatility, and MM index). Unreported regressions combining these with the variables above yields an R^2 over 70%.

The negative risk response is consistent with the model of Chen, Joslin, and Ni (2014) in which one class of agents (e.g., dealers) have tolerance for unspanned crash risk that is higher than that of another class (e.g., investors) in normal times, but which diminishes rapidly in the face of rising risk, possibly due to binding VaR constraints (that do not apply to the investors). The contrasting positive responses we document present a more complex challenge to theorists. One possibility is that these measures are all, to some extent, picking up subjective measures of investor sentiment or marginal utility, to which dealers are less susceptible.

Whatever the correct interpretation, empirically we have succeeded to a significant degree in describing the factors driving index options positions. As Figure 10 illustrates, the fitted specification (model (4) of Table VIII is shown) tracks the historical data series closely for over 20 years.

6. Conclusion

An enormous and still-growing body of research studies stock index option prices. This literature recognizes the unique ability of this class of derivatives to identify investor preferences and risk assessment. We contribute to this research by presenting the first comprehensive look at quantities of these options. A complete understanding of the role that these options play in the economy should aim to explain both the price and quantity dimensions. Our work thus fills an important gap in the empirical options literature.

Measured in terms of monetary value or gross volatility exposure, the index options market is small. But it is large in terms induced of exposure to jumps. Out-of-the-money puts constitute the predominant position type, suggesting that the primary function of this market is the transfer of unspanned crash risk.

Guided by theories of derivative usage, we seek to explain the time-variation in options exposures by variables potentially capturing frictions, types of heterogeneity, and risks.

The stochastic trend in options usage is well-described by exogenous equity market activity. We interpret the explanatory power of stock market turnover as proxying for heterogeneity (and possibly frictions). Direct evidence on the nature of this heterogeneity is still lacking. It does not appear to be linked to the participation shares, or retail (household) investors or that of mutual funds, pensions, or banks and brokers.

We also do not find evidence linking options quantities to proxies for trading costs, information technology, or interest rates. While the available proxies are not ideal, this

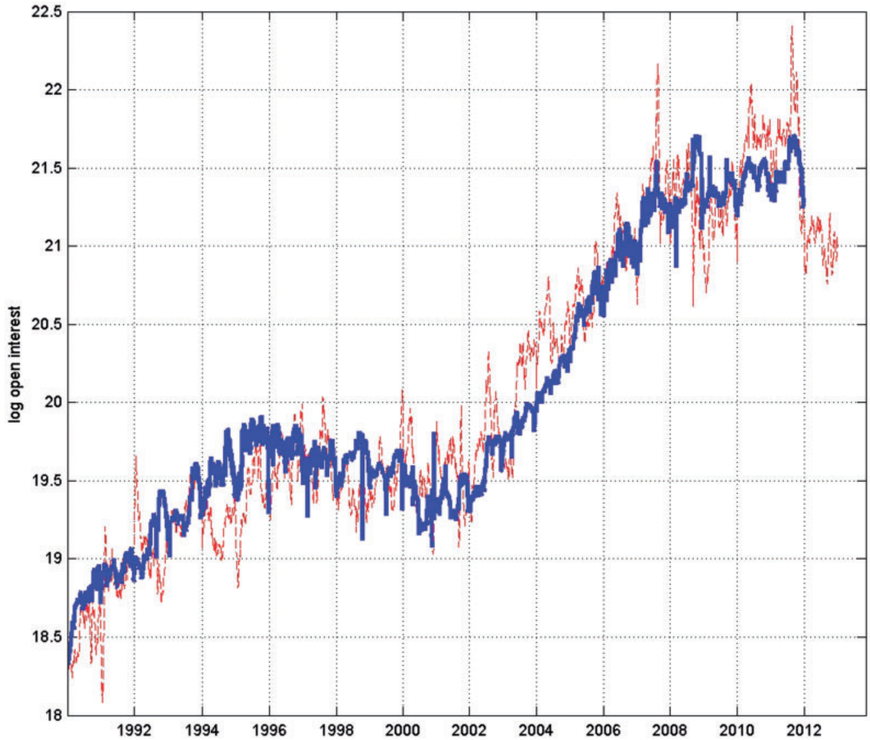

Figure 10. Index option quantities: empirical fit.
The figure shows fit from specification (4) of Table VIII (dashed line) and the raw open interest series (bold line). The stochastic trend and seasonal growth components have been added back to the regression fit.

finding speaks to a widely held view that options offer cost advantages to some investors, and that this drives demand.

Finally, using a rich collection of uncertainty proxies, we distinguish distinct responses to exogenous macroeconomic risk, risk aversion, differences of opinion, and disaster risk. Together these can explain most of the variation in detrended open interest. The single most important relationship is a strong negative response to increases in variance. This is consistent with a negative supply effect driven by sellers of crash risk.

Overall, these findings lay the groundwork for further development and testing of theoretical models of multi-agent economies.

References

Bachmann, R., Elstner, S., and Sims, E. R. (2013) Uncertainty and economic activity: evidence from business survey data, *American Economic Journal: Macroeconomics* 5, 217–249.

Baker, S. R., Bloom, N., and Davis, S. J. (2013) Measuring economic policy uncertainty. Working Paper, Stanford University.

Bates, D. S. (2008) The market for crash risk, *Journal of Economic Dynamics and Control* **32**, 2291–2321.

Bhamra, H. S. and Uppal, R. (2009) The effect of introducing a non-redundant derivative on the volatility of stock-market returns when agents differ in risk aversion, *Review of Financial Studies* **22**, 2303–2330.

Black, F. (1975) Fact and fantasy in the use of options, *Financial Analysts Journal* **31**, 36–41.

Black, F. and Scholes, M. (1973) The pricing of options and corporate liabilities, *Journal of Political Economy* **81**, 637–654.

Brennan, M. J. and Solanki, R. (1981) Optimal portfolio insurance, *Journal of Financial and Quantitative Analysis* **16**, 279–300.

Broadie, M. and Detemple, J. (1996) American option valuation: new bounds, approximations, *Review of Financial Studies* **9**, 1211–1250.

Buraschi, A. and Jiltsov, A. (2006) Model uncertainty and option markets with heterogeneous beliefs, *The Journal of Finance* **61**, 2841–2897.

Carr, P., Jin, X., and Madan, D. B. (2001) Optimal investment in derivative securities, *Finance and Stochastics* **5**, 33–59.

Chen, H., Joslin, S., and Ni, S. X. (2014) Demand for crash insurance, intermediary constraints, and stock return predictability. Working paper, M.I.T.

Chen, H., Joslin, S., and Tran, N.-K. (2012) Rare disasters and risk sharing with heterogeneous beliefs, *Review of Financial Studies* **25**, 2189–2224.

Easley, D., O'hara, M., and Srinivas, P. S. (1998) Option volume and stock prices: evidence on where informed traders trade, *The Journal of Finance* **53**, 431–465.

Eeckhoudt, L., and Gollier, C. (2013) The effects of changes in risk on risk taking: a survey, in: Dionne, G (ed), *Handbook of Insurance*, Springer, New York, pp. 123–134, http://link.springer.com/chapter/10.1007/978-1-4614-0155-1_5.

Engle, R. F. and Granger, C. W. J. (1987) Co-integration and error correction: representation, estimation, and testing, *Econometrica* **55**, 251–276.

Epstein, L. G. and Zin, S. E. (1989) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework, *Econometrica* **57**, 937–969.

Franke, G., Stapleton, R. C., and Subrahmanyam, M. G. (1998) Who buys and who sells options: the role of options in an economy with background risk, *Journal of Economic Theory* **82**, 89–109.

Garleanu, N., Heje Pedersen, L., and Poteshman, A. M. (2009) Demand-based option pricing, *Review of Financial Studies* **22**, 4259–4299.

Grossman, S. J. and Zhou, Z. (1996) Equilibrium analysis of portfolio insurance, *The Journal of Finance* **51**, 1379–1403.

Hakansson, N. H. (1978) Welfare aspects of options and supershares, *Journal of Finance 33*, 759–776.

Johansen, S. (1991) Estimation and hypothesis testing of cointegrating vectors in Gaussian vector autoregressive models, *Econometrica* **59**, 1551–1580.

Joshi, M. (2007) The convergence of binomial trees for pricing the American put. Available at SSRN 1030143.

Jurado, K., Ludvigson, S. C., and Ng, S. (2013) Measuring uncertainty. Discussion paper, National Bureau of Economic Research.

Lakonishok, J., Lee, I., Pearson, N. D., and Poteshman, A. M. (2007) Option market activity, *Review of Financial Studies* **20**, 813–857.

Leland, H. E. (1980) Who should buy portfolio insurance?, *The Journal of Finance* **35**, 581–594.

Lemmon, M and Ni, S. X. (2013) Differences in trading and pricing between stock and index options, *Management Science* **60**, 1985–2001.

Manela, A. and Moreira, A. (2015) News implied volatility and disaster concerns. Working paper, Washington University.

Nelson, D. B. (1991) Conditional heteroskedasticity in asset returns: a new approach, *Econometrica* **259**, 347–370.

Newey, W. K. and West, K. D. (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**, 703–708.

Park, J. Y. and Phillips, P. C. B. (1988) Statistical inference in regressions with integrated processes: Part 1, *Econometric Theory* **4**, 468–497.

Park, J. Y. and Phillips, P. C. B. (1989) Statistical inference in regressions with integrated processes: Part 2, *Econometric Theory* **5**, 95–131.

Phillips, P. C. B. and Perron, P. (1988) Testing for a unit root in time series regression, *Biometrika* **75**, 335–346.

Roll, R., Schwartz, E., and Subrahmanyam, A. (2010) O/s: the relative trading activity in options and stock, *Journal of Financial Economics* **96**, 1–17.

Ross, S. A. (1976) Options and efficiency, *The Quarterly Journal of Economics* **90**, 75–89.

Rothschild, M. and Stiglitz, J. E. (1971) Increasing risk ii: its economic consequences, *Journal of Economic Theory* **3**, 66–84.

Tian, Y. (1993) A modified lattice approach to option pricing, *Journal of Futures Markets* **13**, 563–577.

Vanden, J. M. (2004) Options trading and the CAPM, *Review of Financial Studies* **17**, 207–238.

Appendix A

Contract Description

This appendix provides details on the three classes of options that the study amalgamates.

CBOE S&P 500 Index (SPX) Options

The Chicago Board Options Exchange (CBOE) began trading options based on S&P 500 index with the ticker symbol SPX in 1983. The SPX options are European style, and are cash-settled. The SPX options have a large notional size with a multiplier of 100 of the underlying index value. The CBOE has introduced numerous variants of its basic products over the years, including those with smaller denomination, end-of-quarter and end-of-week expirations, longer term (LEAPS), and closing-price settlement.[36] The most liquid versions have consistently been the standard third-Friday expirations on the March, June, September, December cycle augmented with the nearest two non-cycle monthly expirations. The CBOE options were exclusively pit traded until 2012 when trading in a parallel version was launched on the CBOE's electronic exchange, known as C2.

Our data on CBOE options comes from two sources: the CBOE itself and OptionMetrics. The latter is a commercially available database going back to 1996. The CBOE itself has stored data from 1990 onward, which we acquired. Our data include most of the minor variants listed above, as well as the flagship liquid product. The data include closing price, traded volume, and open interest for every available strike-expiration pair of puts and calls.

CME S&P 500 Futures Options

The Chicago Mercantile Exchange (CME) began trading futures contracts referencing the S&P 500 in 1982, and introduced options on them in 1983. The underlying for each option

36 Currently the standard products are settled by an average of opening prices of the underlying index stocks on expiration day. Closing-based settlement procedures were used when the options were first introduced.

is one cash-settled futures contract; the original futures contract unit size was 500 times the index value. This was lowered to 250 in 1997. The options are American style.[37] The futures contracts have the same quarterly, third-Friday settlement calendar as the main CBOE products. The CME options may have expiration date the same as, or prior to, the future settlement. As with the CBOE, the most liquid products are the nearest few months (with standard third-Friday expiration). Also like the CBOE, the CME has experimented with variants like end-of-week and end-of-month expirations with limited investor interest.

In contrast, an extremely successful variation has been the introduction of smaller denomination contracts. The CME began trading so-called "e-mini" S&P 500 futures in 1997 with a contract size of 50 times the underlying index value. In 1997, the CME listed options based on the e-mini contract, again with one option on one futures contract. In every respect other than denomination, the original (big) contract and the e-mini are the same: the options are American and the primary expiration dates are the same cycle. The trading mechanisms for the big and little contracts are different, however. Both futures and options on the former are primarily pit-traded, while e-mini futures and options are traded on Globex, the CME's electronic limit order book.[38]

Our data on futures options were purchased from the CME and goes back to contract inception in 1983. We have the same basic end-of-day statistics as for the CBOE. Coverage includes the minor contract variants, as well as the main big and e-mini options.

Options on SPDR ETFs (SPY)

Introduced in 1993, State Street's Standard & Poor's Depositary Receipts (SPDR—ticker symbol SPY) has become an extremely liquid ETF that is designed to closely track the S&P 500 stock market index. In 2005, the CBOE introduced options based on SPDRs. SPDRs are traded like common stock and the SPDR options have the same features as standard listed US stock options: American-style exercise and physical settlement (not cash). Each SPDR options contract is on 100 units of the underlying ETF. However, the denomination of the ETF is 1/10th of the value of the S&P 500 index itself. Thus, effectively, one option on SPY references 10 index-equivalent units.

Our data on SPY options come from OptionMetrics, which consolidates data across multiple option exchanges.

Efficient Computation of American Option Exposures

This appendix describes the computational method used to handle the large number of American-style options on futures traded on the Chicago Mercantile Exchange. Our study requires us to aggregate the volatility exposures (vegas) across all options with any open interest on a given day. This, in turn, requires an "implied volatility" for each such option. Because the options are American and early exercise is sometimes optimal, closed-form solutions are not available. Numerical methods are required for two separate steps. First, a calibration step is required to deduce an implied volatility for each option. Second, the volatility sensitivity of the resulting valuation is computed. We need to do this for approximately 4 million option observations with nonzero open interest.

37 Note that these options are not cash-settled: exercise results in a long or short position in the futures contract.

38 The big contracts are also eligible for trade on Globex.

Our approach utilizes binomial trees as the basic valuation model. While trees can be made quite fast, the primary numerical challenge to a brute-force approach is in the implied volatility step, which would require an inefficient search over the space of trees.

To deal with this, we construct highly accurate, multidimensional look-up tables (LUTs). We precompute option values on a four-dimensional grid of expiration, interest rate, moneyness, and volatility. The two steps described above can then be performed through a combination of implicit and explicit interpolation within these tables, as we describe below.

For each node on the grid we evaluate call and put prices on a third-moment tree as per Tian (1993), with analytical smoothing (via Black's formula) in the last step as suggested by Broadie and Detemple (1996). This procedure has among the best convergence properties of the methods studied by Joshi (2007).

Implied volatility

Given an option price, the first step is to do an implicit interpolation in our LUT to deduce the volatility which would yield that price. That is, given the time to expiration, interest rate,[39] and moneyness, we linearly interpolate to obtain the option values, $\hat{o}(\sigma_i)$, at each node on our volatility grid, $\{\sigma\}_i$. The volatility grid values are then treated as a function, and its value is interpolated at the observed price of the option over the grid defined by $\hat{o}$, to obtain the implied value $\hat{\sigma}$.

The procedure fails when the option price is low enough that it is in (or below) the early exercise region, where the volatility surface is not invertible. In this case, we can assign the option an implied volatility of zero without loss of accuracy because the ultimate object of interest is the volatility sensitivity of the option, which must necessarily be zero for such an option. (These zero implied volatility values are not themselves used for anything else.)

Sensitivity

Given the implied $\hat{\sigma}$, we obtain the volatility sensitivity via numerically differencing the values in our LUT along the volatility dimension, and then explicitly interpolating the differenced LUT at the observation's expiration, interest rate, moneyness, and implied volatility. Our grid step size was chosen with the specific criterion of achieving acceptable accuracy of these first differences.

This step is quite efficient since four-dimensional linear interpolation can be handled by vectorized MATLAB routine interpn. A single call to this routine suffices to handle all the put data; a second processes all the calls.

Nonstationary Regressions

The inferences in the tables in Table V are all made under the null hypothesis that the estimated trend from the previous section correctly removes the nonstationarity from OI. However, it is well known that stock market turnover (which drives the fitted trend) is positively correlated with volatility. Therefore, it is worth asking whether inference about the trend is affected by the presence of uncertainty variables, and vice versa. As a robustness check, we therefore run the regressions on the raw (not detrended) OI series, with TO on the right-hand side.

39 We use the 30 day eurodollar rate from the Federal Reserves H15 survey as the riskless rate in pricing calculations.

Table AI. Nonstationary open interest regressions

The dependent variable in the regressions is the raw (non-detrended) aggregate open interest series. The independent variables are as described in the captions to Tables VI and VIII. Finite sample *t*-statistics, shown in parentheses, are computed under the null of cointegration between OI and TO with a cointegrating vector $(-1, 1.70)$.

	(1)	(2)	(3)	(4)
VIX	−0.3942	−0.9827	−0.9305	−0.9380
	(3.15)	(5.21)	(4.74)	(4.82)
SKEW		−9.1617	−8.3063	−8.3863
		(3.20)	(2.88)	(2.96)
VIX $*$ SKEW		3.4607	3.1707	3.2227
		(3.59)	(3.27)	(3.30)
σ_{VIX}			0.2379	0.1999
			(3.30)	(2.89)
FDISP				0.8105
				(1.40)
BBM				0.0014
				(1.77)
$\bar{\mathcal{U}}^y(3)$				−1.0891
				(1.58)
TO	1.6017	1.5973	1.6074	1.6389
	(0.84)	(0.82)	(0.76)	(0.50)
N	5,776	5,776	5,776	5,526

One needs to be careful econometrically with estimating such specifications, since it requires inference about the joint effects of stationary and nonstationary variables. Ordinary least squares (OLS) is not appropriate for such inference due to the spurious regression problem. However, the consistency of OLS under the null of cointegration has been shown by Park and Phillips (1988, 1989) who also derive the nonstandard asymptotic theory. Here, we adopt the null hypothesis of a correctly specified cointegrating relationship (as per the FIML results) and we report finite-sample standard errors computed via simulation under this null.[40] This specification includes no seasonal growth term. The results, shown in Table AI indicate that there little evidence of misestimation of the turnover coefficient due to the omission of risk variables. The risk responses themselves are somewhat attenuated in this specification, but remain economically large.

40 Specifically the simulation assumes that the nonstationary independent variables, differenced, together with the stationary independent variables follow a stationary AR(1) process, which is independent from the cointegrating residuals, which are also an AR(1) process.